

Using Orthogonal Defect Classification to Characterize NoSQL Database Defects

João Agnelo[†], Nuno Laranjeiro[†], Jorge Bernardino^{*†}

[†]*CISUC, Department of Informatics Engineering,
University of Coimbra, Portugal*

^{*}*ISEC – Polytechnic of Coimbra, Portugal*

jagnelo@student.dei.uc.pt, cnl@dei.uc.pt, jorge@isec.pt

Abstract—NoSQL databases are increasingly used for storing and managing data in business-critical Big Data systems. The presence of software defects (i.e., bugs) in these databases can bring in severe consequences to the NoSQL services being offered, such as data loss or service unavailability. Thus, it is essential to understand the types of defects that frequently affect these databases, allowing developers take action in an informed manner (e.g., redirect testing efforts). In this paper, we use Orthogonal Defect Classification (ODC) to classify a total of 4096 software defects from three of the current top NoSQL databases: MongoDB, Cassandra, and HBase. The results show great similarity for the defects across the three different NoSQL systems and, at the same time, show the differences and heterogeneity regarding works carried out in other domains and types of applications, emphasizing the need for possessing such unique information. Our results expose the defect distributions in NoSQL databases, provide a foundation for selecting representative defects for NoSQL systems, and, overall, can be useful for developers for verifying and building more reliable NoSQL database systems.

Keywords—*Defect Analysis, NoSQL, Orthogonal Defect Classification, Software Defect, Software Fault*

I. INTRODUCTION

In recent years, we observed a noticeable growth in the use of NoSQL Database Management Systems that nowadays compete with the older and much more mature, Relational Databases. This evolution derives from the needs of the Big Data era, namely the need for supporting high data volumes, high velocity, and also variety (i.e., different forms of data), along with easier horizontal scalability [1]. In these systems, losing data or being unable to use the storage service has negative consequences not only for the user, but also for the provider, that sees its reputation affected. Such events are many times the result of the presence of a defect in the software [2], [3].

Software defects (i.e., bugs) can range from simple spelling or grammar mistakes in user messages, to security vulnerabilities which, once exploited, may result in disclosure of private information or in infrastructure damage [4], for example. There are many cases of software defects that have led to disastrous consequences, from misguided space rockets, to crashing military airplanes, and ultimately, human deaths [5]. Independently of the severity of a defect, for the sake of dependability, any defect found should be immediately fixed. Otherwise, the product reputation (and also the company behind the product, or the one using it) may collapse to an irrecoverable level. In the case of the most popular NoSQL databases, the identification and fixing of bugs is largely supported by their large communities, which actively find software defects, report them in an issue-tracking platform, and trigger their correction.

Understanding the nature of software defects (e.g., by using structured methods for defect classification [6]–[8]) can provide extremely useful information for developers, who can better focus their software verification efforts or, for instance, concentrate their development efforts in certain components or code areas known to be prone to a particular software defect type [6]. This kind of effort has been undertaken for many and very different types of software that operate in distinct domains, including browsers, games, operating systems, satellite software, business-critical applications, build systems, or machine learning software, just to name a few [9]–[12]. Many times, the general hypothesis placed is that the specificities involving the system (e.g., programming language, project nature and requirements, type and experience of the application development team) make it necessary for such specific analysis to take place. In this context, the analysis of defects affecting storage intensive systems (i.e., NoSQL databases, in the case of this work) has been disregarded so far.

In this paper, we analyze the defects of three top NoSQL databases: MongoDB [13], Cassandra [14], and HBase [15]. We begin by training the application of Orthogonal Defect Classification (ODC) [6] against a total of 300 defects affecting these databases. We discard the training dataset and then analyze a set of 4096 defects extracted from the databases' issue-tracking systems and manually classify them with ODC using the previously trained ODC researcher. We internally double-check the classification result for 20% (i.e., 820 bugs) of the defects and then ask for two external and ODC-knowledgeable researchers to produce independent ODC classifications for an additional 20% of the total defects (10% per researcher).

Finally, we analyze the resulting dataset of ODC-classified defects according to the following four perspectives: i) we analyze the distribution of values concerning six of the eight existing ODC attributes (e.g., type of defect, conditions that trigger the defect, impact of the defect); ii) we form pairs of attributes and analyze the value distributions, as in related work (e.g., [6], [9], [16]); iii) we

analyze the defect type distribution in the top 3 most affected components per each database; and iv) we analyze our results in perspective with previous work.

We observed a huge variation in the distribution of defect types found in related work. Our results have confirmed this heterogeneity, which in practice means that we cannot assume, *a priori*, the existence of a certain defect distribution. This implies that this kind of study must be carried out whenever it is important to know which are the representative defects for a certain type of system, built in specific conditions. In this work, we actually observed a unique distribution of values (i.e., in terms of the relative popularity of the types of defects found), that is not present in any of the related work analyzed. The results suggest that the overall nature of the system may be a factor that influences the distribution of defects.

We also observed great similarity across the three databases, when inspecting the results for the individual ODC attributes. Sometimes, a single value dominated the distribution. When crossing pairs of attributes, three of the main observations include the fact that testing activities are more than two times more frequently associated with reliability defects than with capability defects, also, checking defects (e.g., input validation) are more frequent when the impact is reliability and are more likely to be associated with the “missing” qualifier. We also noticed disparities in the distribution of defect types in different system components, which may be justified with the nature of the component (e.g., a replication component holding Timing/Serialization defects), and, finally, we found certain types of defects being consistently associated with longer times to fix across all three databases (e.g., Function/Class/Object).

Overall, our results (available in detail, along with supporting code, at [17]) confirm the need for understanding the defect trends in NoSQL systems and bring insight regarding their defect distribution. The main contributions of this paper are the following:

- We use Orthogonal Defect Classification to provide a detailed view over a relatively large dataset of reported bugs in three popular NoSQL databases, including a view on the most affected database components;
- We show our results in perspective with the very heterogeneous related work, signaling differences and also similarities found;
- We provide an open dataset holding 4096 bug reports classified using ODC and freely available for future research.

The analysis provided in this paper can help uncover weak spots in these databases (e.g., poor focus on design, a tendency for having a certain type of defect in a certain database component) and this represents essential information for practitioners (e.g., NoSQL database developers) and researchers. The data allow practitioners to obtain knowledge regarding the reliability – or lack thereof – of their systems, but it could also help to improve the quality of the development processes, for instance by directing verification efforts to the appropriate areas (e.g., design, algorithm) or selecting certain verification techniques (e.g., testing, code inspections). Researchers may benefit from this type of results and use them as basis for creating new testing techniques for this kind of systems; for carrying out fault injection campaigns using the defect information obtained from this work; or even for tailoring development processes to specifically consider the specificities of this kind of systems and the nature of the defects that typically affects NoSQL databases.

The remainder of this paper is organized as follows. Section II presents background on ODC and discusses the related work. Section III describes our approach for classifying the extracted defects and Section IV presents and analyses the results. Section V discusses the main findings of this work and Section VI presents the threats to the validity of the work. Finally, Section VII concludes the paper.

II. BACKGROUND AND RELATED WORK

In this section, we briefly go through the main software defect classification schemes used nowadays, with focus on the main concepts regarding the Orthogonal Defect Classification, and we then discuss the related work.

A. Background on Orthogonal Defect Classification

Software Defect Classification Schemes are powerful tools which developers and researchers can use for different purposes such as development process improvement, product quality enhancement or empirical analysis. The three most popular software defect classification schemes are Hewlett-Packard’s Defect Origins, Types and Modes (DOTM) [8], IEEE’s standard 1044-2009 (IEEE-1044) [7], and Orthogonal Defect Classification [6].

Hewlett-Packard’s Defect Origins, Types and Modes (DOTM) [8] contains three base categories to classify defects i) *Origin* – the first activity in the defect’s lifecycle in which it have been detected, ii) *Type* – the area, of a particular Origin, that is responsible for the defect (e.g., “Logic” or “Standards” for Origin “Code”, or “Software Interface” for Origin “Design”), and iii) *Mode* – classification of why the defect occurred. IEEE 1044-2009 - IEEE Standard Classification for Software Anomalies [7] is a relatively complex standard, mostly due to its large number of attributes. Still, this allows for less subjectivity in its application, as the large number of attributes covers a wider variety of possibilities.

Orthogonal Defect Classification is a set of analytical methods used for software development and test process analysis to characterize software defects that consists of eight orthogonal attributes [6]. ODC was created with the intention of bridging two

methods commonly used for analyzing defects, namely Statistical Defect Modelling (e.g., software reliability statistical models) [18] and Root Cause Analysis [19]. It allows the defect classification process to be faster (an advantage of the former method) and have better accuracy in categorizing issues (similarly to the latter method) [6]. ODC is widely regarded as very popular means to characterize defects and, as discussed in this section, has been used by researchers and practitioners in several contexts. Its documentation [6], [20] is easily accessible and detailed, leaving little space for doubts regarding its use. ODC is based on the definition of eight attributes grouped into two sections: open-report and closed-report [20]. *Open-report* refers exclusively to the attributes that can be classified when the defect is found (i.e., independently of a later possible correction of that particular defect) and corresponds to the following attributes:

- **Activity:** indicates the activity being performed at the time the defect was disclosed (e.g., system testing);
- **Trigger:** describes what caused the defect to surface, i.e., the required condition that allowed the defect to manifest itself;
- **Impact:** refers to either the impact a user experienced when the defect surfaced (in case of a user-reported defect), or the impact a user would have hypothetically suffered, had the defect surfaced (when a developer finds a dormant defect, which is yet to be triggered).

The *closed-report* section refers to the attributes that can be classified when a defect has been corrected and the correction information becomes available. This section’s attributes are the following:

- **Target:** the entity that was corrected (e.g., source code);
- **Defect Type:** the nature of the change that was performed to fix the defect (e.g., algorithm);
- **Qualifier:** this attribute complements the defect type and describes the state of the implementation prior to having been corrected (e.g., whether it was missing, incorrect, or present but unnecessary);
- **Age:** refers to the point in time in which the defect was introduced (e.g., introduced during the correction of another defect);
- **Source:** whether the defect was introduced by an external component (i.e., outsourced), or is something developed by the team itself (i.e., in-house).

At the time of writing, the most recent ODC documentation is at version 5.2 [20] and it describes the ODC attribute set, their possible values and example cases. This ODC documentation should be used as a guideline and software development teams are allowed to create their own set of values for each attribute, in case there is a better fit for their specific software system. In the context of this work, we opted to use the sets of attribute values presented in [20], without further customization.

It is worth mentioning that the defect type and qualifier attributes are only applicable when a defect’s target is either “Design” or “Code”, which means that for all other target values, both attributes must be suppressed [20]. Additionally, there is a mapping between the activity and trigger attributes, where certain activities directly map to certain triggers [20], this mapping is shown in Table I.

TABLE I. ACTIVITY-TRIGGER MAPPING

Activity	Design Review	Code Inspection	Unit Test	Function Test	System Test
Trigger	<ul style="list-style-type: none"> • Design Conformance • Logic/Flow • Backward Compatibility • Lateral Compatibility • Concurrency • Internal Document • Language Dependency • Side Effects • Rare Situation 		<ul style="list-style-type: none"> • Simple Path • Complex Path 	<ul style="list-style-type: none"> • Test Coverage • Test Variation • Test Sequencing • Test Interaction 	<ul style="list-style-type: none"> • Workload/Stress • Recovery/Exception • Startup/Restart • Hardware Configuration • Software Configuration • Blocked Test

The ODC documentation does not specify a particular procedure to apply the classification scheme, thus when classifying a defect, we go through the attributes in the order they are discussed in [20]. Obviously, when a defect is found, the first step of the ODC classification is to categorize the occurrence in the open-report attributes. When the defect is corrected, the remaining ODC attributes – the ones in the closed-report section – are classified. Notice that the scale is orthogonal, i.e., there are no overlaps and the order in which attributes are classified, in general, does not impact the results. The only special case is the case of Activity, where an incorrectly classified Activity will lead, in some cases, to an incorrect Trigger. So, in the case of this work, the researchers involved were made aware of this, and were instructed to pay special attention to this characteristic.

The next section summarizes related work, especially works in which the numeric outcome of the application of ODC is clearly reported by the authors. Works that do not clearly mention the outcome of bug analysis, such as the exact distribution of defect types, are excluded from the analysis.

B. Related Work

We conducted a search through the major indexing engines, namely Google Scholar, Scopus, IEEE Xplore, ACM Portal, and Springer Link search engines. Main keywords used for searching included ODC, Orthogonal Defect Classification, Software defect, Software fault. We were able to find 17 papers of various types (e.g., published in peer-reviewed conference proceedings, journals) that use ODC to classify software defects and that, in total, analyze 60 different applications, going through a total of 13171 bugs. Although we found relatively old works, almost one third of the papers found were published in the last 3 years, which signals the importance of carrying out this kind of work for newer types of systems, such as NoSQL databases. The next paragraphs describe the related work found, but, note that we also summarize the results obtained by other authors in Section V, in perspective with our own results (and, therefore, do not present them in depth in this section).

The related work that uses ODC fits in essentially two groups of work. In **Group I** authors carry out *empirical analyses* using ODC to characterize a given set of bugs (e.g., due to the importance of the system or absence of prior data) within a certain system context (e.g., machine learning tools, build systems, safety-critical software). The works belonging to **Group II** also perform empirical analyses, but the work has an additional goal, a second application, which is, most of the times, to use the ODC data for *process improvement* (in *lato sensu*). In this latter group of works, it is frequent to see the authors using the field data to perform root-cause analyses [12] or apply it within the scope of another proposal (e.g., fault injection) [9]. Our work fits directly in the former group.

Regarding the works belonging to **Group I**, which concentrate on performing empirical analyses, Basso et al. [21] use field data of several Java-based systems in production and apply ODC, with the goal of understanding fault representativeness, including security vulnerabilities. The authors put their results in perspective with previous work [9], which targeted systems written in C. Thung et al. [10] perform an empirical study of bugs in machine learning systems. The authors aim to understand how often bugs appear in these systems, which types of bugs surface, how long does it take to fix different types of bugs, and how many files are impacted by the defects. Xia et al. [11] analyze bugs affecting four different build systems (e.g., Ant, Maven), considering that the reliability of these tools could affect the build process reliability. The authors analyze a random sample of 800 bug reports and their corresponding fixes and present the different relative frequency types of bugs. Also discussed is the relationship between types of bugs and their corresponding severities.

Fonseca and Vieira [22] used ODC to manually characterize 655 software bugs related to security using six widely used web applications written in PHP. The results show that, in that context, just a part of the types of software bugs is related to security. The authors also note that web application vulnerabilities result from software bugs that tend to affect a reduced set of statements. The work allows future definition of realistic fault models that originate security vulnerabilities in web applications. Based on the fact that it is rare to find analysis of software defects for safety-critical systems, Silva and Vieira [23] applied ODC to a set of 243 defects obtained from systems that operate in the aerospace and space domains. In addition to the ODC results, the authors have highlighted the challenges of performing the classification, in particular in applying the broad ODC approach to certain specific types of issues. Xuan et al. [24] carry out an empirical evaluation of 300 software bugs in industrial financial systems. These systems tend to be heavily complex, holding many parts with intricate business logic, which makes them different from systems studied in related work. The authors have analyzed bug density, detection time, the distribution of bug categories, and the relationship between categories and bug severity.

Morrisson et al. [25] use ODC to characterize and understand the differences between the discovery and resolution of non-security related defects compared to vulnerabilities. The analysis of the differences allows for security-specific software development process improvements, which are lightly mentioned by the authors, but the work's research questions aim at: i) understanding if detection and discovery of vulnerabilities occur in the same manner as for plain software defects, and ii) are the different types of vulnerabilities discovered and resolved in the same way. The work goes through 1166 bugs in three open-source projects: FireFox, PPHMyAdmin, and Google Chrome and applies ODC+V to better capture vulnerability data. Rahman et al. [26] classify software defects in Puppet scripts belonging to the open software repositories of Mirantis, Mozilla, Openstack, and Wikimedia Commons. The motivation is that, besides the growing importance of this kind of scripts, the nature of the defects in Puppet scripts has not yet been categorized. Using 2 raters per defect (89 raters in total), the authors apply ODC to 3187 defects and observe that configuration-related defects are prevalent.

Regarding the works belonging to **Group II**, Lyu et al. studied the nature, type, detectability, and effect of software defects in program versions built by 34 programming teams for a critical flight application. ODC was used to classify the defects. The authors applied mutation testing techniques to create mutants, using real faults. Among other aspects, the results showed that coverage testing is an effective way for detecting faults, but mutation testing can be a more trustworthy indicator of test quality. Lutz and Morgan [27] used ODC to characterize bugs and discover defect patterns in the spacecraft domain. The results obtained allow the authors to produce a short list of recommendations to avoid undesirable defect patterns and, more generally, for process improvement.

Christmansson and Chillarege [28] use ODC to present a set of errors which emulate software faults. The proposed approach uses field data for generating a set of injectable errors (an error is an effect of a fault [4]), each of which being defined by i) error type, ii) error location, and iii) injection condition. The authors resorted to ODC, as a way of categorizing the required field data, and classified 408 real software defects with just one of the many attributes ODC offers – *Defect type*. The work answers questions related with fault forecasting and has established a solid basis for work in the field of dependable systems.

Durães and Madeira [9] analyzed how software faults can be injected in a source code-independent manner. The authors analyzed 668 real software defects and classified them using ODC. The resulting dataset was analyzed for patterns which were then used in the definition of a set of software fault emulation operators. These operators allow for software faults to be injected in real software, which is particularly important in cases where field data is scarce, regarding the typical faults that affect the software. The emulation operators derived from the field study allowed the authors to propose a fault injection technique named G-SWIFT.

Børretzen and Dyre-Hansen [29] characterized faults in five business-critical industrial projects with the goal of determining possible areas of process improvement. Findings include the fact that a certain fault type (Function/Class/Object) generally dominates the bug reports from all projects and that there is a tendency for some fault types to be marked as more severe as others (e.g., relationship, timing/serialization). The work resulted in recommendations to the organization fault reporting process, in particular the need for enriching the information contained in the fault reports.

The work by Gupta et al. [30] aims at characterize bugs present in reusable software to understand the reasons for the typical lower defect densities in this kind of software. The authors conduct a case study in an industrial setting and apply ODC to 1310 bugs and then perform a qualitative root cause analysis. Results indicate that several factors must be taken into account when analyzing the cause-effect relationship between software reuse and lower defect density. The findings should also encourage practitioners to implement further software reuse policies.

Silva et al. [31] applied ODC in the context of space critical systems, discovering difficulties in classifying about 32% of the bugs analyzed due to context specificities. Thus, the authors propose an adaptation of the taxonomy that is particularly suitable for critical systems, allowing for a more effective classification. Silva et al. [12] carried out a root cause analysis on 1070 software defects belonging to the subsystems of four safety-critical space systems. The authors applied a version of ODC particularly tailored for the context and examined types of defects, impacts, and triggers. The authors propose modifications to the way software is developed and also to the verification and validation activities. The idea is that the analysis of the field data using ODC allows better defect prevention and detection. Finally, Sotiropoulos et al. [32] characterize bugs in navigation software for outdoor robots. The authors use ODC and the analysis of the triggers and effects of the bugs shows that a large part can be disclosed using low-fidelity simulation. Also, the work provides insight into navigation scenarios that could serve as basis for testing.

As we have seen, previous work either fits in purely empirical studies (which is also the case for this work) or in empirical studies accompanied of a second goal (e.g., fault injection). In the related work, defect classification schemes have been used in a very large variety of contexts and tend to produce heterogeneous results (as discussed in Section IV). This strongly suggests that there is the need for understanding the nature of the software defects involved in the particular class of systems studied in the present work. To the best of our knowledge and at the time of writing, we produced the largest ODC dataset among those analyzed in related work, which we make available at [17] for future research. We also use 6 out the 8 ODC attributes (in practice, all attributes for which we had sufficient information – the exception is ‘age’ and ‘source’) in the analysis which is something rare to observe in related work (mostly due to the amount of human effort involved).

III. STUDY DESIGN

This section describes the design of our study, which has the overall goal of providing information regarding software defects present in NoSQL databases. In practice, we go through the following five steps that are detailed in the next paragraphs:

- 1) Selection of NoSQL databases;
- 2) Training a researcher, named *Researcher1*, to learn the classification process, using real defect descriptions (selected and extracted from the databases’ public defect-tracking platforms);
- 3) Selection and extraction of a dataset composed of defect descriptions from the databases’ public defect-tracking platforms;
- 4) Manual classification of each software defect according to the ODC scheme, carried out in five batches for each database;
- 5) Manual verification of the defects classified in *step 4*) (with possible reclassification of some of the defects):
 - a. *Internal Verification* carried out by *Researcher1* to signal errors and improve the overall quality of the dataset. As described in the next paragraphs, in an effort to gradually increase the quality of the classification, this was a task incrementally carried out along with the classification procedure (and not only after the complete classification was finished);

- b. *External Verification* using two new independent classifications, each of them carried out by external researchers (*Researcher2* and *Researcher3*).

We started by **selecting NoSQL databases**, so that we could gather defect data that would allow us to characterize the defects present in these kinds of systems. We aimed for popular open-source NoSQL databases, as we are interested in analyzing systems that are important for users. As these are supported by larger communities, we also had the expectation of finding large quantities of defect descriptions in their public defect-tracking systems, thus providing sufficient data for analysis.

We started by examining the database popularity rankings in *db-engines.com* [33], *stackoverflow.com* [34], and *kdnuggets.com* [35] and found a general agreement in the following ranking (from the most popular to the least): 1) MongoDB; 2) Redis; 3) Cassandra; 4) HBase; and 5) Neo4j. Our intention was to analyze the defects of the top three most popular databases, in a semi-automated manner (e.g., although the analysis is manual, supporting processes, such as defect selection and defect description extraction, are automatic). Of the five, three (MongoDB, Cassandra, and Hbase) use the popular defect-tracking system JIRA, and the remaining two (Redis and Neo4j) use GitHub. The latter were excluded from the analysis in this work, as their defect-tracking platforms severely lacked the necessary information for a proper ODC classification. Additionally, while each JIRA deployment presented a lot more than 1000 defects, Redis and Neo4j's GitHub platforms barely reached half of that. The number of defects at our disposal was an important factor, as the analysis of more defects can allow for more accurate analysis.

The three selected databases (MongoDB, Cassandra, and Hbase) serve the same general purpose but obviously also have technical differences. At the time of writing and, as reported by OpenHub (<https://openhub.net>), MongoDB is mostly written in C++ (42% of the code) and Go (33% of the code). The whole codebase is close to 2 million lines, created by 492 contributors (135 active contributors in the last year). Cassandra's codebase is around 400K lines of code, with 96% of the codebase being in Java that has been created by 372 contributors (80 active contributors in the last year). Finally, HBase's codebase is around 920K lines of code, of which 86% are written in Java that have been committed by 337 contributors (135 active contributors in the last year).

As a way of overcoming the learning curve inherent to applying the ODC scheme and also with the goal of obtaining more reliable results, we chose to use an initial set of 300 defect descriptions (100 defects for each of the three databases) to **train the classification process**. This also allowed for a better understanding of how these databases' JIRA platforms are structured, and the procedures used by the communities and development teams when registering and describing defects. This training effort was done by randomly selecting defect reports from the whole range of defects present in each database's JIRA platform, and classifying them individually. The results obtained from this process were discarded as they were merely used for training and learning.

The **selection and extraction of the set of defects** for applying the ODC scheme was carried out as follows. We started by defining the period of analysis, for which we opted to set no particular limit – we considered all defects present in each of the databases' JIRA platforms as potential candidates for classification. Obviously, for our analysis we use only defects tagged with *closed* and *resolved* (i.e., defects whose existence has been confirmed by the developers and for which a correction already exists), so that our analysis would not involve any false positive (i.e., a report of a defect that actually does not exist). Also, this allows us to perform the full ODC classification, which requires a defect to have passed through both the open-report and closed-report stages.

We first identified all closed and resolved defects, registered until May 1st 2017, which accounts for 7456 for MongoDB, 4576 for Cassandra, and 4905 for Hbase (respectively 8, 8, and 10 year periods of defect reporting). As the classification procedure is manual and we must reduce the human effort involved (also as a way to reduce the risk of erroneous classification), we selected a subset of the defects, which we aimed to be around one fourth of the total closed and resolved defects, for a total of precisely 4096 defects. This resulted in a total of 1618, 1095 and 1383 defects for MongoDB, Cassandra and Hbase.

The next step was to **classify each defect** with the ODC scheme, according to the following steps:

- i. Interpretation of the defect's written description;
- ii. Classification of the open-report ODC attributes;
- iii. Analysis and interpretation of the source code changes;
- iv. Classification of the closed-report ODC attributes.

The information required to perform the abovementioned steps is held in each defect description. This is essentially a textual description that typically includes the conditions that allowed the defect to surface (e.g., a memory leak in a routine), the environment in which it occurred (e.g., the defect surfaced during a resource-intensive procedure), and the corrective measures that were applied in the respective fix (e.g., the defect was fixed by adding the missing functionality to release unused resources), and obviously, a description of the defect itself. Note also that these steps are the direct application of the ODC methodology, without any particular adaptation. The only relevant aspect to mention is that we did not use the age and source attributes because, in general, the defect reports did not include sufficient or clear information on these attributes.

Due to the human intervention involved in the classification process, the final result of classifying a given set of defects may hold errors. This is due to the fact that the application of a given defect classification scheme may involve some subjectivity (i.e., in some cases, two users of the same scheme could classify a given defect in a different way), which is aggravated when the defect is not

described in a complete or clear way. Obviously, there is also human error in the process, even when the description is completely accurate. This is a common issue in works that use defect classification schemes, as mentioned in [36], [37], which is usually mitigated by the use of more than just one rater (i.e., more than one human classifying each bug). Due to the very large size of our dataset and associated huge manual effort, we verify part of the bugs and involve external researchers in the process, as discussed in the next paragraph.

To **verify the classification** of the bugs and, most of all to have some insight regarding the reliability of the annotated dataset, we performed two verification activities against 40% of the bugs (1640 out of 4096 bugs), which we name *Internal* and *External* and that are depicted in Figure 1. During the *Internal Verification* activity, the researcher responsible for originally classifying the bugs (named *Researcher1*) revisited 20% of the bugs (820 of the 4096 bugs) to check for errors and correct any misclassification found. In practice, the whole set of bugs was divided in two, and this internal verification was performed against the first half. This first half (i.e., 2048 bugs) was divided in four batches (512 bugs per batch) that were verified in the following manner. By the end of each batch, the researcher double-checked a total of 40% of the already classified defects in that batch, according to the following distribution: i) the first 20% of the bugs in the batch; ii) a random selection of another 20% from the remaining defects in the batch. The *External Verification* activity involved independent classification by two external researchers (named *Researcher2* and *Researcher3*) of 20% randomly selected, but non-overlapping bugs (i.e., 820 bugs in total, 410 bugs per each of the two researchers) across the whole set of 4096 bugs.




Fig. 1. Diagram of the internal (in orange) and external (in blue) verification of the classification process for each batch.

In Figure 2, we show the error results of the *Internal Verification*, detailing the number of defects that had their classification corrected per batch (i.e., where at least one of the ODC attributes changed its value) and including an error trendline for each of the three databases. In terms of the overall accuracy of the *internal classification*, we obtained 82% of correctness, with the researcher maintaining its classification in 672 bugs of the 820 verified bugs. Notice that this is a pessimistic view of the process in the sense that we mark a bug as incorrectly classified, even if the error refers to just one of the six attributes (with the remaining five being correct). Obviously, there is always some error associated with the application of ODC, which is the result of the human intervention during the process, but also due to the limited or ambiguous information found in some defect reports, which introduces some uncertainty about which value to apply to a given ODC attribute. Despite this, the tendency for lower error values in latter batches, suggests that the researcher improved his skills during the process allowing for a dataset of higher quality. Note also that among the three databases, there are a few differences in the way defects are reported, which explains the higher error values for the first batches in each database.

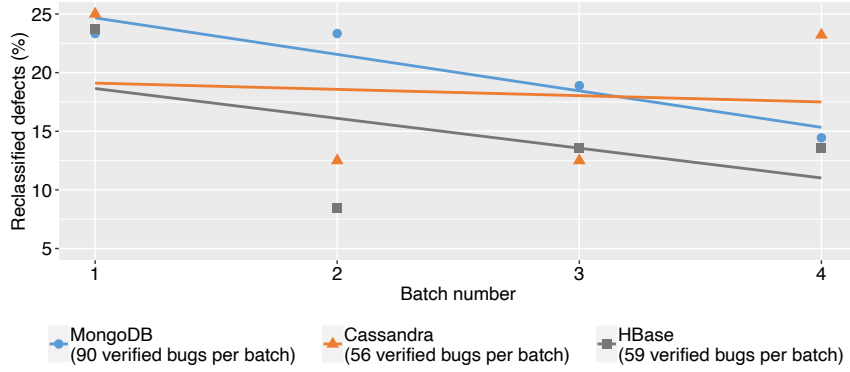


Fig. 2. Reclassified defects in respect to the total number of bugs per batch and for each of the four initial batches, for each database.

Table II shows three confusion matrixes that detail the outcome of the verification procedure, which is composed of three verification tasks, named *Internal* (carried out by *Researcher1*), *External 1* (carried out by *Researcher2*) and *External 2* (carried out by *Researcher3*), regarding the *Defect Type* attribute. We chose to present the detailed results for this case, due to the fact that it is the most widely used ODC attribute in related work [9], [10], [21]. Notice that the count of bugs for *Internal* is 746, which is slightly below the previously mentioned 40% (820 bugs) as we are considering only the bugs where the *Target* attribute had Code as value (thus, qualify for being marked with a defect type) and were verified and kept as Code by *Researcher1*. Thus, 69 bugs were confirmed to not be code defects and the value of the *Target* attribute was changed by *Researcher1* for 5 bugs (i.e., 2 bugs left *Target* Code and 3 bugs became Code problems). Note also that, in each matrix, each cell holds the total number of bugs mapped to a certain attribute value, the values read in the columns represent the outcome of the verification procedure. In light blue, we highlight the true positives (the bugs in which there is agreement between the original classification and the respective verification task).

TABLE II. CONFUSION MATRIXES FOR THE VERIFICATION PROCEDURES

Defect Type	A/I	C	A/M	F/C/O	T/S	I/OOM	R
Internal (746 bugs)							
Assignment/Initialization	34		4			1	
Checking		71	4				
Algorithm/Method	4	6	394	6	1	3	
Function/Class/Object			13	126		1	
Timing/Serialization			1		4		
Interface/O-O Messages			4	1		65	
Relationship				1		1	1
External 1 (410 bugs)							
Assignment/Initialization	20						
Checking		44	2			2	
Algorithm/Method	2	3	238	3	1	1	
Function/Class/Object	1		5	47		3	
Timing/Serialization					7		
Interface/O-O Messages						30	
Relationship							1
External 2 (410 bugs)							
Assignment/Initialization	14		5				
Checking		38	1				
Algorithm/Method	9	7	232	8		2	
Function/Class/Object				53			
Timing/Serialization					2		
Interface/O-O Messages	9		1			21	
Relationship							1

We also analyzed the results obtained by *Researcher2* with Cohen's Kappa, as it is able to measure the agreement between two raters (i.e., *Researcher1* and *Researcher2*) that classify items in mutually exclusive categories [38]. The definition of k is:

$$\kappa = \frac{p_0 - p_c}{1 - p_c} \quad (1)$$

where p_0 is the relative observed agreement between raters (i.e., accuracy) and p_c is the hypothetical probability of chance agreement. If the raters fully agree, then $\kappa = 1$, if there is no agreement beyond what is expected by chance, then $\kappa = 0$, and, finally, a negative value reflects the cases where agreement is worse than random choice. Overall, the following terms apply for the value of κ : *poor* when less than 0, *slight* between 0.01 and 0.2, *fair* between 0.21 and 0.4, *moderate* between 0.41 and 0.6, *substantial* between 0.61 and 0.8, and finally *almost perfect* between 0.81 and 0.99 [39]. The accuracy results of *Researcher2* are presented in Table III,

along with the Kappa value. Due to time restrictions, *Researcher3* classified only the defect type attribute, thus the corresponding results are only presented afterwards, by the end of the next paragraph.

TABLE III. ACCURACY AND COHEN’S KAPPA AGREEMENT FOR RESEARCHER2 EXTERNAL VERIFICATION

ODC Attribute	Accuracy	Kappa	Agreement
Activity	0.71	0.64	substantial
Trigger	0.60	0.58	moderate
Impact	0.83	0.82	almost perfect
Target	0.98	0.98	almost perfect
Defect Type	0.94	0.93	almost perfect
Qualifier	0.93	0.89	almost perfect

As we can see in Table III, we have obtained an *almost perfect* agreement for most of the ODC attributes. The exception is *Activity* with *substantial agreement* and *Trigger* with *moderate agreement*. Regarding *Activity*, the most notable case of disagreement between the researchers is the case where code inspection bugs are marked by *Researcher2* with one of the three possible *Testing* cases (i.e., unit testing, function testing, system testing). In fact, *Researcher2* classified about one third of the 229 bugs marked with code inspection by *Researcher1* as one of the *Testing* cases (15 bugs marked with unit testing, 30 with function testing, and 32 with system testing). This has impact in the *Trigger* classification, as a wrong activity will lead to a wrong *Trigger* (due to the Activity-Trigger mapping presented in Table I, and with exception of bugs that pass from *design review* to *code inspection*, and vice-versa, which does not occur in our study). Therefore, it is expectable that the accuracy values for *Trigger* are lower (due to the accumulated error), which is actually the case. As mentioned, *Researcher3* classified only the defect type attribute, for which it reached an accuracy of 0.91 which corresponds to a Kappa value of 0.9, also achieving an *almost perfect* agreement. These results and analysis reflect the overall quality of the dataset, providing us with the quality assurances that should be present in this type of work. Detailed results are available at [17].

IV. RESULTS

In this section, we describe the results obtained from applying ODC to the 4096 defects collected from the issue-tracking platforms of the three previously mentioned NoSQL databases. We first analyze the distribution of values obtained for each individual ODC attribute, we then analyze the results for pairs of ODC attributes and conclude with a detailed view on the most affected components per database. All detailed results and also supporting code are available at [17].

A. Value Distribution Across ODC Attributes

We begin by overviewing the results obtained for each individual ODC attribute (i.e., activity, trigger, impact, target, defect type, and qualifier). This allows us to understand trends regarding common types of defects or corrective measures, for instance, but is particularly important for dependability researchers to identify representative software defects in the domain of NoSQL databases (e.g., to use in fault injection campaigns). The values obtained for the **activity attribute** for each of the three databases, are shown in Figure 3.




Fig. 3. Activity attribute distribution.

As we can see in Figure 3, the values for the activity attribute are distributed in a very similar fashion among the three databases, with the most common value being “Code Inspection” appearing in around half of the cases. Notice that, in our case, with no special adaptation of ODC, “Code Inspection” refers to human manual inspection but also automatic inspection carried out by tools (e.g., static code analyzers). The fact that this attribute value is associated with around half of the defects essentially shows the importance of this activity in disclosing defects in these database systems. Testing related activities (i.e., “Unit Test”, “Function Test” and “System Test”) show somewhat similar values between each other. It was often difficult to classify defects in terms of this attribute, mostly due to the lack of information, in the defect description, regarding which activity was being carried out at the time the defect

was found. So, whenever the information was not sufficiently clear, *Researcher1* classified the trigger first, which allowed to return and classify the activity, based on the activity-trigger mapping previously mentioned.

Figure 4 shows the distribution of the values obtained for the **trigger attribute**. The different trigger values are grouped by the activity-trigger mapping shown previously. “Logic/Flow” is the most frequent trigger and has its higher count in HBase. “Side Effects” is also a frequent HBase issue (and also in the remaining databases, although with a lower expression), which may suggest a stronger coupling in the way the different parts of the system are written or built. This suggests that the design of the system and the process used for development may require a reflection and improvement. Finally, “Simple Path”, “Test Coverage”, and “Blocked Test” are expectable cases (in the sense that they represent basic testing cases), arising from unit, function, and system testing, respectively. Overall, this emphasizes the *importance of carrying out different testing activities for defect disclosure*. The remaining trigger types show lower occurrence rates, with most of them below 4%. Across databases, there is high variation in the frequency of triggers observed, with no clear pattern. The only visible pattern is the fact that Cassandra occupies the middle position in 16 of the 21 triggers. So, there is a strong variation in the triggers across databases and the different Activities they are associated with, which mostly reflects the way the different communities operate and the way the databases are being developed and verified.




Fig. 4. Trigger attribute distribution.

The **impact attribute** represents the hypothetical effect a defect would have had on a user, if it had been triggered on the field. Figure 5 shows this attribute’s value distribution across the three databases. As we can see in Figure 5, “Capability” and “Reliability” jointly account for about three quarters of the types of impact in each of the three databases. “Capability” is, as defined in [20], the ability of the system to perform its intended and required functions (the customer is not impacted in any of the remaining impact values). Capability is, by definition, a generic attribute value, and it is often used when none of the remaining impact types seem to fit the defect being analyzed. This may justify its higher values in all three databases, but, more importantly, it reflects the fact that a faultload holding representative bugs should impact, in half of the cases, the intended function of the system. In the same manner, in one fourth of the cases, bugs should impact “Reliability”. In this context, reliability represents critical situations in which a system would halt or crash in an unplanned behavior [20].

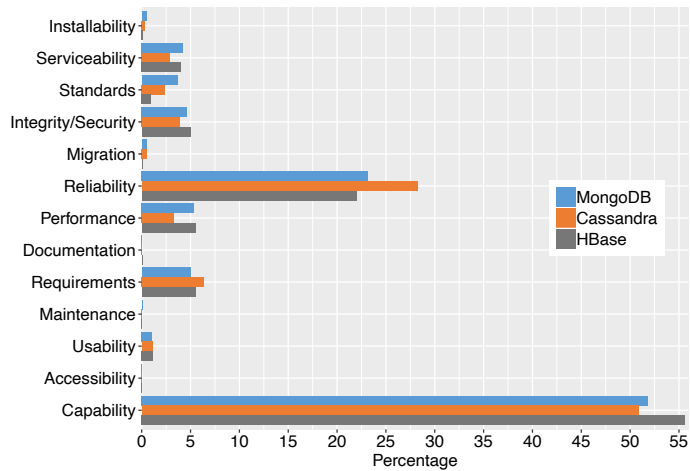


Fig. 5. Impact attribute distribution.

The remaining one fourth of this distribution is covered by all the other impact values. Impacts such as “Installability”, “Maintenance”, “Migration”, “Documentation” or even “Accessibility” were found to be the rarest in the distribution, with some cases showing no occurrences whatsoever. Apart from these, the other impact types appear with some variation, with values ranging between around 5% and 2%. Cassandra is less prone to performance issues, which can be important information for providers that want a long running installation of this database.

Table IV presents the **target attribute**’s distribution (i.e., the entity that was corrected) regarding the three database systems. In what concerns the target attribute, most of the reported defects refer to source code problems (every nine out of ten defects are a code bug, in all three databases). Much less defects were found under “Build/Package” (build or packaging scripts), with the remaining types showing up either very infrequently or never.

TABLE IV. TARGET ATTRIBUTE DISTRIBUTION

Target	MongoDB	Cassandra	HBase
Requirements	0%	0%	0%
Design	0%	0%	0%
Code	93.1%	94.1%	94.7%
Build/Package	6.8%	5.4%	4.4%
Information Development	0.1%	0.5%	0.9%
National Language Support	0%	0%	0%

Figure 6 shows the value distribution for the **defect type attribute**. The distribution of defect type’s values is dominated by “Algorithm/Method”, which is, as defined in [20], a quite generic type of defect that fits many different cases. As an example, a defect that consists of multiple “Assignment/Initialization” corrections, may correspond to an “Algorithm/Method” defect type as opposed to the otherwise assumed “Assignment/Initialization”. Furthermore, cases which contained corrections of both “Assignment/Initialization” and “Checking” types were often classified as “Algorithm/Method”, as there is no option to place multiple tags on a given defect.

The second most frequent type of defect – “Function/Class/Object” refers to large changes in the design of a software system, i.e., a complete change in the way the system performs a certain function, or even the addition or removal of such functions. The need for redesign could either be a consequence of a poor initial design of these databases before the first implementation steps, or simply the need for redesigning large functions in order to bring these databases up-to-date with modern standards, or to fit new requirements, for example. MongoDB seems slightly more prone to interface issues than the remaining, which may be relevant information for developers, as interface issues tend to modify the developer’s perception about a given system’s reliability.

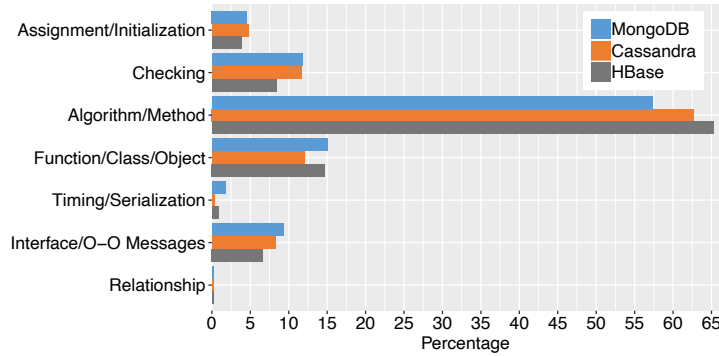


Fig. 6. Defect Type attribute distribution.

Finally, Table V shows how the defects analyzed fit in the **qualifier attribute**. As we can see in Table V, more than half of the defects were found to be of the type “Incorrect”. This means that more than half of the defects were corrected by directly changing (i.e., re-implementing) the affected source code. The second most common qualifier is “Missing”, which occurs when the correction of a defect is done by adding code that was otherwise absent. The least common type of qualifier is “Extraneous” (around 5% in the three databases), where the corrective measure consists of removing unnecessary code.

TABLE V. QUALIFIER ATTRIBUTE DISTRIBUTION

Qualifier	MongoDB	Cassandra	HBase
Missing	33.0%	29.9%	27.0%
Incorrect	63.1%	65.6%	68.9%
Extraneous	3.9%	4.5%	4.1%

Overall, we observed great similarity in values (and thus, in defect trends) across the three databases, although there are a few exceptions. The exact root cause of all exceptional cases is difficult to determine (and is actually out of the scope of this paper), but, in the case of this single attribute view, the relevant part is the overall distribution found. Quite often we found cases where one single value dominated the distribution by occurring more than half of the remaining values together (e.g., in activity, target, defect type and qualifier). Dependability researchers may be able to use this data to select one or a set of representative defects for NoSQL databases (e.g., to use in fault injection campaigns, or to generate code mutations, as carried out respectively in the work by Durães and Madeira [9] and Lyu et al. [40]). The dataset is open (available at [17]), and selection of certain subsets of bugs is obviously also possible, if the goal is, for instance, to explore certain properties of the system (e.g., timing).

B. Value Distribution Across Pairs of Attributes

Analyzing the ODC results according to pairs of attributes is a common practice in this kind of study [6], [9], [16]. These two-way attribute relationships allow us to identify further defect trends. Of the six ODC attributes used in this work, we have excluded activity and target from this phase of the analysis, as activity can be extrapolated from the value of a trigger, and because target is dominated by a single value – “Code” (over 90%).

We begin by analyzing the **impact–trigger** pair. We have previously observed that “Capability” and “Reliability” together cover around three quarters of the impact attribute’s distribution (see Figure 5). Due to this, we limit the analysis to these two values against the different types of triggers, which we present in Figure 7. Note that, to further improve the readability of this figure, we excluded triggers showing up with a frequency lower than 1% for at least one database in the trigger attribute’s distribution.

As we can see in Figure 7, the distribution of the types of triggers tend to be different depending on the type of impact. The inspection-related triggers “Design Conformance” and “Logic/Flow” show a stronger link to defects that impact “Capability”, which is expectable as, by definition, these kinds of triggers can easily impact the ability of the system to perform its intended functions. The testing-related triggers tend to occur more often together with the “Reliability” impact, representing around two thirds of the reliability issues. Indeed, a defect that impacts reliability is something that, in general tends to be detected at runtime, rather than by inspecting source code in which some kinds of issues are hard to catch (e.g., concurrency problems). A particular example can be found with “Blocked Test”, which occurs when a certain system test cannot be concluded or cannot even run due to basic problems, we can see that these three databases struggled, at some point, to run against the developers’ system tests. This does not mean that the system cannot be deployed and run in normal operating conditions, it just means that a certain system test fails to run (e.g., a stress test failing to run due to a misconfiguration of the server). On one hand this may indicate quality of the developers’ system test

suites, on the other hand, may also point out low quality development. In either case, what is relevant is to understand the distribution across impacts so that properly configured verification activities may take place (e.g., selecting the right verification activities to try to trigger certain types of problems).

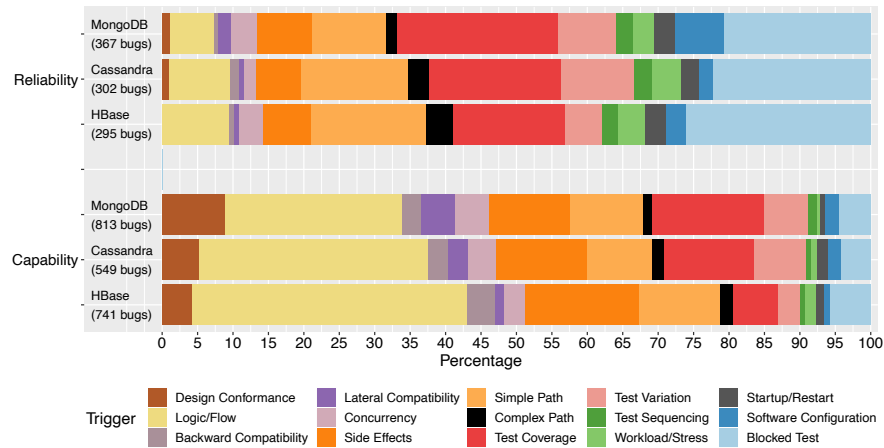


Fig. 7. Impact-Trigger Type pair distribution.

Figure 8 shows the results of the **impact-defect type** pair. As before, we limit impact to the most frequent types “Capability” and “Reliability” and we can see that there are no major differences between the relative distribution of types of defects across both impacts. The only minor differences are related with a relatively higher presence of “Checking” defects in the “Reliability” impact, which is understandable as these types of issues are, many times, associated with robustness problems [41]. Also, “Capability” tends to gather more presence of “Function/Class/Object” defects, which, in agreement with ODC, is in fact a defect that significantly affects “Capability” and this highlights the importance of design in the overall system. So, if the goal is to prioritize capability versus reliability, for instance, tuning the design-oriented activities of the project, or the overall development process, could help diminishing the frequency of “Function/Class/Object” defects in the final product.




Fig. 8. Impact-Defect Type pair distribution.

The last pair involving the impact attribute is the **impact-qualifier** pair. Figure 9 displays the distribution of the “Capability” and “Reliability” impacts for each of the three existing types of qualifier.

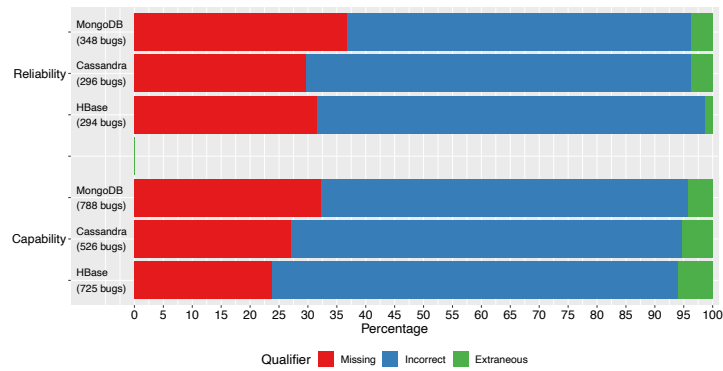


Fig. 9. Impact-Qualifier pair distribution.

Despite there are a few variations between the three databases, in general the distribution of values is quite similar, with “Incorrect”, “Missing” and “Extraneous” being the most, second and least common qualifiers, respectively. Note that this order has also been previously observed in the qualifier’s individual distribution. It is worth mentioning that this order of the qualifiers is also kept in the remaining types of impact (not included in Figure 9).

Figure 10 and Figure 11 show the results involving the **trigger attribute** against **defect type** and **qualifier**, respectively. In both cases, we excluded all types of trigger values that, for all three databases, occurred less than 1/4 of the value for the most frequent trigger type. For instance, given a top trigger value with 20% occurrence rate, the threshold for accepting other triggers would require any of their values to exceed one fourth of this (i.e., 5%). The goal was to focus the analysis on the more relevant pairs between the trigger attribute and other attributes, as the former contains a huge array of values, which would render the figures unreadable.




Fig. 10. Trigger-Defect Type pair distribution.

Figure 10 presents the selected **trigger-defect type** pair and their respective value distribution. We can see that there is some variation of the relative popularity of each defect type, including some differences regarding the overall distribution of defect types.

The most notable cases include HBase/Design Conformance, where the top defect is, by far, “Function/Class/Object”, which strongly suggests either a weaker design of the database or an abnormal need for frequent formal design changes. In either case, developers may be able to reduce the presence of this defect if more time is invested in the design of the system. Another visible case is the fact that “Checking” defects are the second most frequent defect types in the testing-related triggers – Simple Path and Test Coverage (in the overall distribution Function/Class/Object stands at second place). Indeed, this kind of issues is better captured with testing activities than with code inspection or design reviews and it would be abnormal to see defects like Function/Class/Object defects being captured by this kind of triggers more often than Checking defects. Actually, such case happens for HBase under the Test Coverage trigger, which suggests that the related testing activities (i.e., Function Test) may be in need for improvements. Figure 11 shows the distribution of values for the **trigger-qualifier** pair.




Fig. 11. Trigger-Qualifier pair distribution.

In a similar way to the previous figure, here we also refer to the fact that these results generally match the ones presented by these attribute’s individual distributions. In terms of the distribution of the qualifier values, and although there are some variations, we observe the same order of frequency for all trigger types, with “Incorrect” being the most frequent value.

Figure 12 shows the value distribution of the **defect type–qualifier** pair, which is a very common view of this kind of data [6], [9], [16]. It essentially characterizes the correction that was applied – defect type – and the state in which the code was, prior to being corrected – qualifier.




Fig. 12. Defect Type-Qualifier pair distribution.

We can see in Figure 12 that “Checking” is the type of defects that has a higher relative probability of being “Missing” than “Incorrect” or “Extraneous” (e.g., a missing value verification in the code). This also occurs in “Timing/Serialization”, but in this case the bug count is rather low. The case of “Missing” “Checking” defects becomes particularly problematic, in the sense that this type of defects is many times related to instructions that are able to control the flow of a program based on specific conditions – hence “Checking”. Their absence can lead to many kinds of problems, typically robustness problems (if checking is missing at the boundaries of the system) but also security problems (e.g., missing validation for malicious user input) [41].

C. Internal View and Time to Fix

In this section, we provide an internal view of the types of bugs affecting the top 3 components that had the highest number of closed and resolved bugs (in each of the databases). We then close the section showing the average time to fix per type of defect and per database. Table VI overviews each of the most affected components, the descriptions were adapted from each database’s technical documentation.

TABLE VI. DESCRIPTION OF THE THREE DATABASE COMPONENTS WITH HIGHER DEFECT COUNTS

Database	Component	Description
MongoDB	Sharding	Sharding is the process of distributing data across multiple machines (shards) and is handled and balanced automatically by MongoDB, depending on the amount of data and cluster size. Sharding effectively enables easy horizontal scalability and ensures good performance in high throughput environments using commodity hardware.
	Replication	Replication consists in replicating the same data and storing it in multiple servers within a replica set (i.e., a group of processes solely dedicated to this task). It provides redundancy (it is a fault-tolerant mechanism) and increases availability. In some cases, this may even increase read performance, as read requests for the same data may be routed to any of the available machines in a given replica set.

	Querying	Refers to the MongoDB querying API. As opposed to classical SQL used in relational Databases, MongoDB's query model is function-based, as opposed to string-based.
Cassandra	Tools	This component refers to tools that ship with Cassandra such as the Stress Tool used in stress tests, the Node Tool used for managing clusters, the CQL shell (CQLSH) and other tools to handle SSTables.
	CQL	Cassandra Query Language (CQL) is the query language used in the Cassandra database, and it highly resembles the well-known Structured Query Language (SQL) used for querying in relational databases.
	Local Write-Read Paths	Bugs labelled with this component affect one or more of the elements related to writing and reading to and from Cassandra. This may include Memtables (in-memory tables for hot writes and reads), the commit log (a log of commits ensuring redundancy and fault tolerance), SSTables (Sorted Strings Tables, persistent data tables stored in disk), Cache-related logic or even low-level disk I/O issues.
HBase	RegionServer	In HBase, tables are divided into regions, which are themselves managed by Region Servers. These act as nodes that make up the distributed database logic behind HBase.
	Client	Bugs marked with this component affect the HBase Java Client API. This is the interface through which one is able to perform all CRUD operations on HBase tables by using the Data Manipulation Language (DML).
	Master	This is known as the Master Server in HBase. It is responsible for managing all the subsystems such as assigning regions to each server and balancing the load between between Region Servers. In short, it orchestrates an entire HBase cluster.

Figure 13 shows the distribution of defect types found in each of the top 3 most affected components per database. Notice that the overall sum adds up to 3846 (instead of 4096) as some defects are not code defects and, thus, do not qualify for being classified with a ‘defect type’.




Fig. 13. Types of defects found in the three components with higher defect counts in each database.

We computed the Relative Change (RC) [42] values for the different types of defect for each component in respect to the overall values, as follows:

$$RC(x, x_{reference}) = \frac{x - x_{reference}}{x_{reference}} \quad (2)$$

where $x_{reference}$ represents the overall value for a certain type of defect in a certain database and x represents the value for that same type of defect type but observed in a specific database component. Table VII holds the final values, where the cases marked in blue represent RC values greater or equal to $|0.5|$ and where $x_{reference}$ has a value of at least 2%.

TABLE VII. RELATIVE CHANGE VALUES OF THE DEFECT TYPES FOR THE MOST AFFECTED COMPONENTS IN RESPECT TO THE OVERALL VALUES

Defect type	MongoDB			Cassandra			Hbase		
	Sharding	Replication	Querying	Tools	CQL	Local W-RP	regionServer	Client	Master
Assignment/Initialization	-0.3	-0.1	0.0	-0.2	-0.6	-0.2	0.3	0.5	-0.3
Checking	-0.2	-0.1	0.1	-0.3	0.4	0.0	-0.1	-0.6	-0.8
Algorithm/Method	0.0	0.0	0.2	0.1	-0.2	0.2	0.1	0.1	0.0
Function/Class/Object	0.4	-0.1	-0.2	-0.5	0.9	-0.7	-0.3	0.1	0.6
Timing/Serialization	-0.3	1.1	-0.6	-1.0	-1.0	-1.0	0.0	-1.0	0.6
Interface/O-O Messages	0.1	-0.1	-0.8	0.7	-0.1	-0.5	-0.2	-0.6	0.2
Relationship	-1.0	-1.0	-1.0	-1.0	-1.0	-1.0	-1.0	4.2	-1.0

In perspective with the overall values, there are a few tendencies in specific components that are worthwhile mentioning within each database. Regarding **MongoDB**, Timing/Serialization defects tend to have more weight in the *Replication* component, which

emphasizes the component’s sensitivity to timing issues and, at the same time, suggests that stronger verification activities, namely further extensive runtime testing, that particularly takes in consideration timing aspects, are important and a relevant case to consider in the context of this component. Another visible aspect is related with the *Querying* component, where the Interface/O-O Messages defects are relatively lower. The component is very user-centric and possible problems should be captured and quickly resolved during informal testing activities. We have observed that bugs in this highly exposed component actually tend to be fixed very quickly. The average time to fix of Interface/O-O Messages defects in the Querying component of MongoDB is 1.36 days, while the overall time to fix of this kind of defects considering the whole system is 25.8 days. A more detailed analysis of the time to fix per defect type is presented later in this section (please refer to Table VIII).

In the case of **Cassandra**, the *Tools* component presents relatively less Function/Class/Object defects and more Interface/O-O Messages defects. This component includes a number of quite diverse tools, so it is difficult to reason about possible causes, as there is no single nature for the whole component. The only common aspect that is shared by the tools is that they are meant to be used directly by a human user (in very different ways, still). Thus, considering the tools diversity it is more or less expectable that interface issues and messages may play an important role. Regardless of the reason, this data highlights weak spots that could benefit from different verification strategies. In the case of CQL, the component that supports Cassandra’s query language, the highest variations considering the overall values come from Function/Class/Object (with nearly the double of defects of the overall) and also, at a smaller scale, from Assignment defects, which reduce to about half the overall value. The former case reflects the need for a formal design change and the higher count of Function/Class/Object defects just reflects the fact that this component is prone to design changes. In turn, this suggests that a different development process may be needed, so that this kind of changes decrease. The whole overall design of the component may be in need for a larger major design change. The lower presence of assignment defects may be the result of various factors, but it suggests that tests under the presence of limit conditions, which are known to effectively disclose such bugs, are being carried out by developers (which are, in fact, present at the respective testing code repository [43]). In the Local Write-Read Paths component, Algorithm/Method defects have the highest relative weight across components, even considering components from other databases. Also, Function/Class/Object and Interface/O-O Messages show large decreases in this component, but, as the number of bugs is also relatively low, no further conclusions can be made.

Regarding **HBase**, regionServer has little variation in comparison with the overall results. The *Client* component has relatively less presence of Checking and Interface/O-O Messages defects. One possible cause is that the nature of the component makes developer concentrate on this type of issues, potentially leaving a lower number of bugs in the released code. As this component has high exposure, we believe that existing bugs of this type (interface/message problems and checking issues, that are usually performed at the entry point of the API) should be disclosed easily. Finally, in the Master component, Checking bugs also decrease, while Function/Class/Object and Timing/Serialization increase. Again, this is a time sensitive component for which the current verification activities are apparently not sufficient. The higher presence of Function/Class/Object defects suggest that there are deficiencies in the design that could benefit from a higher investment in such important aspect of this component.

Table VIII shows the average time the developers take to fix each type of defect for each of the databases, including a view of the times per different bug priority. The table also includes the standard deviation and respective count of bugs.

TABLE VIII. TIME TO FIX PER TYPE OF DEFECT

Defect Type	Overall time to fix			Priority															
	avg	stdev	count	Blocker - P1			Critical - P2			Major - P3			Minor - P4			Trivial - P5			
				avg	stdev	count	avg	stdev	count	avg	stdev	count	avg	stdev	count	avg	stdev	count	
MongoDB	Assignment/Initialization	29.3	100.0	69	5.9	10.8	55	108.1	179.3	11	170.1	239.4	3	-	-	-	-	-	-
	Checking	48.2	156.2	177	33.3	30.5	2	47.5	159.0	17	43.5	158.9	130	77.5	153.7	24	34.7	56.6	4
	Algorithm/Method	42.7	114.4	864	7.5	14.2	11	28.5	71.2	62	37.0	96.4	696	90.3	192.5	82	146.1	292.3	13
	Function/Class/Object	110.9	277.4	226	9.0	1.8	2	179.1	400.7	23	71.5	156.5	186	544.9	658.5	14	2.7	0.0	1
	Timing/Serialization	21.2	44.8	27	10.5	23.7	23	82.7	76.8	4	-	-	-	-	-	-	-	-	-
	Interface/O-O Messages	25.8	80.9	140	2.1	1.7	3	3.8	2.6	7	23.0	75.0	101	55.8	123.4	21	10.8	18.4	8
	Relationship	0.1	0.1	3	0.1	0.1	3	-	-	-	-	-	-	-	-	-	-	-	-
Cassandra	Assignment/Initialization	28.3	82.7	49	0.9	0.0	1	19.6	49.0	24	21.5	56.8	16	71.6	160.9	8	-	-	-
	Checking	17.8	48.3	120	0.1	0.0	1	9.3	13.1	3	16.5	50.5	53	17.8	43.7	54	30.5	65.7	9
	Algorithm/Method	24.0	69.3	646	11.8	16.1	8	17.4	55.6	23	25.1	79.0	349	25.8	60.1	226	10.4	26.1	40
	Function/Class/Object	63.9	183.9	124	49.6	71.6	4	49.9	122.6	76	92.3	266.2	42	26.4	25.7	2	-	-	-
	Timing/Serialization	1.7	1.4	4	0.1	0.0	1	1.6	0.0	1	1.0	0.0	1	3.9	0.0	1	-	-	-
	Interface/O-O Messages	12.7	31.0	85	0.2	0.0	1	39.7	45.8	4	16.7	38.2	39	7.3	16.9	28	5.1	4.7	13
	Relationship	3.2	1.9	2	5.1	0.0	1	1.3	0.0	1	-	-	-	-	-	-	-	-	-
Hbase	Assignment/Initialization	22.4	67.9	51	13.1	11.9	2	51.4	97.5	7	24.5	77.9	25	11.2	22.5	13	0.2	0.2	4
	Checking	16.3	65.1	110	3.1	6.4	8	5.7	8.3	6	14.3	65.0	62	30.4	85.0	27	3.7	5.4	7
	Algorithm/Method	22.0	89.2	855	16.1	51.9	68	36.9	99.4	92	21.9	98.8	501	19.1	69.6	158	9.0	24.4	36
	Function/Class/Object	35.3	95.4	193	87.3	215.7	22	29.3	45.3	23	28.6	66.3	117	20.9	40.3	29	133.6	128.4	2
	Timing/Serialization	3.9	7.7	11	1.0	0.0	1	0.2	0.0	1	6.4	9.8	6	1.0	0.9	2	0.8	0.0	1
	Interface/O-O Messages	9.8	24.5	87	0.5	0.2	2	12.8	28.7	7	8.9	23.2	44	12.6	28.6	26	5.4	10.0	8
	Relationship	0.4	0.2	3	0.5	0.0	1	0.0	0.0	1	0.6	0.0	1	-	-	-	-	-	-

One of the first visible aspects in Table VIII is that, most of the reported times are associated with quite high standard deviation values, which means that it is hard to predict how much time it will take after a certain type of bug is reported and before a fix is produced. Considering the overall results, Function/Class/Object bugs take the longest to fix. This type of defect should require a formal design change and implies significant capability is affected, so it is interesting to observe that this kind of defects is the one that takes longer to fix, and this applies to all three databases. At the other extreme, we find the rare cases of Relationship bugs, where the few cases observed are corrected either in the same day, or in just two or three days.

Among the fast resolution bugs (and besides Relationship), we usually find Timing/Serialization defects. In Cassandra and HBase these bugs are quickly solved, but they take substantially longer in MongoDB. This suggests that MongoDB is more sensitive to this kind of problems, and they are possibly more difficult to fix. We verified the bug priorities for these cases (marked by the community) and in the case of MongoDB, these bugs were marked with the highest priorities (23 bugs marked with Blocker – the highest priority, and 4 marked with Critical – the second highest), which strongly suggests that they are indeed hard to fix. The remaining databases had Timing/Serialization issues with a clearly better distribution across priorities. Actually, a similar case is observed for Checking, Algorithm/Method, and Interface-O-O Messages. However, the difference is that MongoDB does not have these cases concentrated around the highest or highest priorities. With exception of the abovementioned cases, the time to fix results across databases are quite heterogeneous (especially considering the different priorities involved), which is possibly the result of the intervention of different communities, along with the specificities of each system involved.

V. DISCUSSION

In this section, we highlight and further discuss the main results presented in the previous section and put them in perspective with the results analyzed in related work.

The results obtained in this work, in particular for the **single attribute** analysis presented earlier, show that, regardless of the attribute being considered, *defects extracted from different systems tend to concentrate around just a few attribute values* (e.g., among the thirteen existing impact values, “Capability” and “Reliability” are by far the most common ones). In some cases (e.g., impact, defect type), the distribution is generally around attribute values that are, by definition, somewhat generic and fit a larger number of scenarios. Also, we observed cases where a single value dominated the distribution, by occurring more than half the times than the remaining attributes together. However, in certain attributes, such as the trigger attribute, the results tend to be more scattered, possibly due to the absence of generic trigger types. In certain attributes, we observed several cases of values presenting results ranging from non-existent, to a small fraction of the most common value in the distribution. Additionally, we have observed that in the vast majority of cases, the relative values remain quite similar across the three NoSQL databases.

The results of the analyzed **attribute pairs** in many cases show the same value order observed for the individual distributions. However, there are notable cases worth highlighting. *Testing activities are more than two times more frequent in reliability defects than in capability defects*, which signals the importance of these verification activities to disclose defects associated with this kind of impact. So, a system where reliability is a priority may benefit of a higher investment in testing activities. A similar case was observed with *checking defects, which are more frequent when impact is reliability than when it is capability* which again provides important information for the selection of verification activities.

Contrary to the remaining cases, *in HBase/Design Conformance, the top defect type is, by far, “Function/Class/Object”*, which strongly suggests either a weaker design of the database or an abnormal need for frequent formal design changes. Also regarding HBase, the fact that Function/Class/Object defects appear more frequent than checking defects under the Test Coverage trigger, suggests that *Function Testing strategies in MongoDB may be in need for improvement*. Finally, we observed no major differences when crossing qualifier with other ODC attributes. The only notable case is the fact that *Checking is a defect type that has a higher probability of being “Missing” than “Incorrect” or “Extraneous”*, as opposed to all other types of defects (at least those with sufficiently high bug counts).

Despite the fact that these databases were built in different programming languages (i.e., MongoDB mostly in C++ and Cassandra and HBase in Java), by different development teams, potentially following different methodologies, results show that the defect distribution, share many similarities. This suggests that *the nature of the systems may be linked to the types of defects* affecting such systems. In fact, when we drill down to the component level, and especially when we go through the most affected components per database, we see that *the nature of the different database components appears to be connected to the distribution of defects*, which is quite visible in certain components. As an example, the Replication component in MongoDB, a time-sensitive component, shows a stronger presence of timing/serialization defects. Obviously, there is no rule of thumb and these observations are mere observed facts, that may not apply to other NoSQL databases with similar characteristics.

We now analyze **our results in perspective with previous work** on defect classification using ODC. In this final analysis, we focus on the defect type attribute, as in the abovementioned previous works this is the only attribute for which we found comparable data (rarely, some works also use trigger and impact, but with just a small portion of the wide array of values we use in this paper). As a summary, Christmansson and Chillarege [28] presented a set of errors which emulate software faults and Lyu et al. [40] aimed at evaluating the effectiveness of software testing and software fault tolerance. Lutz and Morgan [27] characterized bugs with to discover defect patterns and Durães and Madeira [9] analyzed how software faults can be injected in a source code-independent

manner. Børretzen and Dyre-Hansen [29] characterized faults in industrial projects for process improvement. Fonseca and Vieira [22] used ODC to characterize security patches of widely used web applications and Basso et al. [21] characterize Java faults to understand their representativeness, including security vulnerabilities. Gupta et al. [30] characterize defects in reusable software to understand the causes for lower defect densities in this kind of software. Thung et al. [10] characterized bugs in machine learning tools and Xia et al. [11] characterize bugs in software build systems. Silva and Vieira [23] assessed ODC's adaptability to critical software environments and Xuan et al. [24] study bugs in industrial financial systems. Silva et al. [31] use ODC as basis to classify critical systems engineering issues and propose an adaptation of the taxonomy. Silva et al. [12] performed a root cause analysis on defects from safety-critical software systems. Morrissou et al. [25] use ODC to characterize and understand the differences between the discovery and resolution of defects compared to vulnerabilities and Sotiropoulos et al. [32] characterize bugs in robot middleware to understand which can be detected through simulation. Rahman et al. [26] purely target defect classification, highlighting the fact that the nature of defects in the nowadays very popular Puppet scripts had not been yet categorized.

Table IX summarizes the main results obtained in these previous works. The columns in the table (from the left to the right) identify the author and year of the work being analysed, the reference for the work used in this paper, the names of the applications analyzed, the relative percentage found per each value of defect type, the total number of comparable defects (i.e., code or design defects classified with one of the 7 defect types in the ODC specification [20]) and their corresponding percentage regarding the total number of defects analysed per application, and also the global number of defects per work. Additionally, we show the number of raters used (those marked with an asterisk correspond to information not present, or not clear, in the paper and that was provided by the paper authors), the programming language, and the type of application being analyzed. The following should be noted: i) defect types not used in a specific work are represented with dashes (-); ii) works which used custom ODC values (e.g., [12], [23]) had any additional defect types excluded from the comparison; and iii) as a visual aid, we present the number of order of each defect type for each work, from 1 (the most popular defect type found for the particular application being analyzed) to 7 (the least frequent defect type). These numbers also correspond, in order, to the colors red, orange, yellow, green, blue, light gray, dark gray.

TABLE IX. ORTHOGONAL DEFECT CLASSIFICATION RESULTS FOUND IN RELATED WORK

Work	Software	Order	Assign / Init	Order	Check	Order	Alg / Met	Order	Func / Class	Order	T/S	Order	Inter / OOM	Order	R	Comparable defects (C)	C / T	Total Defects (T)	Total defects per work	Raters per defect	Language	Domain	
This work	MongoDB	5	4.3	3	10.9	1	53.4	2	14.0	6	1.7	4	8.7	7	0.2	1506	93%	1618	4096	1+1	C++	NoSQL database	
	Cassandra	5	4.5	3	11.0	1	59.0	2	11.3	6	0.4	4	7.8	7	0.2	1030	94%	1095	4096	1+1	Java	NoSQL database	
	Hbase	5	3.7	3	8.0	1	61.8	2	14.0	6	0.8	4	6.3	7	0.2	1310	95%	1383	4096	1+1	Java	NoSQL database	
Christmannson and Chillarege, 1996	Tandem Guardian90 and IBM OS	2	19.1	3	15.2	1	37.8	5	7.6	4	13.2	6	7.1	1	-	408	100%	408	408	-	-	Spacecraft navigation system	
Lyu et al., 2003	RSDIMU	2	31.9	1	33.8	3	19	4	14.1	1	-	5	1.2	-	426	100%	426	426	-	C	Spacecraft navigation system		
Lutz and Morgan, 2004	Deep space antenna controller	3	2.7	-	-	-	1	78.2	-	-	-	2	4.8	-	-	126	75%	167	167	-	-	Space	
Durães and Madeira 2004	Linux	3	22.6	2	25.8	1	33.3	4	12.9	-	-	5	5.4	-	-	93	100%	93	668	1*	C	Operating System	
	Vim	2	21.3	3	18.2	4	9.1	5	0	-	-	1	54.5	-	-	249	100%	249	668	1*	C	Text Editor	
	FCIV	4	11.3	3	13.2	1	52.8	2	15.1	-	-	5	7.5	-	-	53	100%	53	668	1*	C	Game	
	Joe	2	25.6	1	44.9	3	15.4	5	0	-	-	4	14.1	-	-	78	100%	78	668	1*	C	Text Editor	
	MingW	3	10	2	38.3	1	46.6	5	0	-	-	4	5	-	-	60	100%	60	668	1*	C	Programming toolset	
	ScummVM	2	24.3	3	8.1	1	56.8	4	6.8	-	-	5	4.1	-	-	74	100%	74	668	1*	C++	Game emulator	
	CDEX	2	18.2	2	18.2	3	9.1	0	0	-	-	1	54.5	-	-	11	100%	11	668	1*	C++	CD Audio Extractor	
	Pat2html	1	55	3	5	2	40	0	0	-	-	0	0	-	-	20	100%	20	668	1*	C++	PDF converter	
	Gaim	4	4.3	1	52.2	2	26.1	3	13	0	-	-	4	4.3	-	-	23	100%	23	668	1*	C	Instant messaging client
	ZSNES	1	66.7	0	2	33.3	0	0	0	-	-	0	0	-	-	3	100%	3	668	1*	Assembly, C, C++	Game emulator	
	Bash	1	100	0	0	0	0	0	0	-	-	0	0	-	-	2	100%	2	668	1*	C	Command interpreter	
	Firebird	1	50	1	50	0	0	0	0	-	-	0	0	-	-	2	100%	2	668	1*	C++	Browser	
	Borretzen and Dyre-Hansen, 2007	Anonymized P1	2	9.5	3	6.3	7	1.1	1	25.3	4	1.4	5	0.3	5	0.3	217	44%	490	901	-	Java	Business critical - Data registration
		Anonymized P2	4	7.4	2	15.4	3	12	1	24	5	2.3	7	1.1	6	1.7	135	64%	212	901	-	Java	Business critical - Administration tool
Anonymized P3		2	14.6	4	2.4	3	4.9	1	53.7	4	2.4	0	0	0	0	33	78%	42	901	-	Java	Business critical - Merging of applications	
Anonymized P4		2	26.7	0	4	6.7	1	36.7	0	3	10	0	3.3	28	83%	34	901	-	Java	Business critical - Transaction tool			
Anonymized P5		2	14	4	7.5	3	8.6	1	24.7	7	1.1	5	5.4	6	4.3	81	66%	123	901	-	Java	Business critical - Transaction tool	
Fonseca and Vieira, 2008	PHP-Nuke, -Fusion, -myAdmin, -BB, Drupal, WordPress (v)	2	6	4	2.4	1	86	-	-	-	-	3	5.6	-	-	679	100%	679	679	1*	PHP	Web apps	
Tânia et al., 2009	Vuze BitTorrent client	3	15.2	4	11.2	1	36.8	2	28.8	-	-	5	8.0	-	-	125	100%	125	574	3*	Java	P2P BitTorrent client	
	Freemind	3	11.1	4	2.2	1	46.7	2	28.9	-	-	5	11.1	-	-	90	100%	90	574	3*	Java	Diagram application	
	Jedit	3	14.1	5	11.3	1	36.6	2	25.4	-	-	4	12.7	-	-	71	100%	71	574	3*	Java	Text Editor	
	Phex	2	90.0	0	5.0	1	60.0	3	10.0	-	-	3	10.0	-	-	20	100%	20	574	3*	Java	P2P Bnutella client	
	Struts	3	18.2	4	10.0	1	48.5	5	4.0	-	-	2	20.2	-	-	99	100%	99	574	3*	Java	Library	
	Tomcat	4	12.4	2	23.2	1	58.0	5	2.4	-	-	3	13.6	-	-	169	100%	169	574	3*	Java	Web Server	
Gupta et al., 2009	Java Enterprise Framework	3	16.7	4	3.1	4	3.1	1	21.2	-	-	0.0	2	19.9	-	143	64%	223	1310	-	Java	Reusable Java Framework	
	Digital Cargo Files	2	12.4	3	10.4	5	3.3	1	31.3	6	1.2	4	9.9	-	-	300	69%	438	1310	-	Java	Document storage	
	Shipment and Allocation	5	6.7	3	10.9	2	22.8	1	29.8	6	1.6	4	9.3	-	-	526	81%	649	1310	-	Java	Business process management	
Thung et al., 2012	Mahout, Lucene, OpenNLP	3	13	4	11.4	1	22.6	-	-	-	-	4.8	2	15.2	-	335	67%	500	500	-	Java	Machine learning	
Xia et al., 2013	Ant, Maven, Cmake, Qmake	2	8.36	3	7.62	4	4.49	-	-	-	-	5	1.62	1	30.59	-	428	54%	800	800	-	Java, C++	Build Systems
Silva and Vieira, 2014	Satellite on-board SW start-up	6	0.9	2	6.5	3	2.8	1	14.0	4	0.9	4	0.9	-	-	28	26%	107	243	2*	C, Ada, Assembly	Satellite on-board software start-up	
	Satellite on-board SW start-up	0.0	3	7.0	4	1.2	1	20.9	2	8.1	0.0	0.0	0.0	-	-	32	37%	86	243	2*	C, Ada, Assembly	Satellite on-board software start-up	
	Satellite software	0.0	2	15.0	0.0	0.0	1	55.0	0.0	0.0	0.0	0.0	0.0	-	-	14	70%	20	243	2*	C, Ada, Assembly	Satellite software	
	Airborne critical component	0.0	0.0	10.0	3	6.7	1	10.0	5	3.3	3	6.7	-	-	-	11	37%	30	243	2*	C, Ada, Assembly	Airborne critical component	
Xuan et al., 2015	PMS, β - Analyzer and OrderPro	3	12.33	4	8.33	2	17.67	-	-	-	-	5	2.33	1	27.33	-	169	56%	300	300	-	Java, Python, C	Industrial financial systems
Silva and Vieira, 2016	Satellite subsystems	-	3	6.6	2	11.4	1	21.3	-	-	-	3	6.6	-	-	606	82%	739	739	2*	C, Ada, Assembly	Space critical Embedded Systems	
Silva et al., 2017	Satellite subsystems	5	4.3	3	6.4	2	9	1	19.1	6	3.1	4	5.2	-	-	503	47%	1070	1070	2*	C, Ada, Assembly	Satellite subsystems	
Morrison et al., 2017	Chrome	3	13.7	4	10.9	2	29.1	1	36.0	6.0	2.3	7	1.7	5	6.3	175	100%	175	1166	2+1	C++	Browser	
	Chrome (v)	3	12.1	1	39.3	2	30.6	4	8.7	2	1.2	6	2.3	5	5.8	173	100%	173	1166	2+1	C++	Browser	
	Firefox	3	25.2	2	25.6	1	35.3	4	1.1	1	1	0	5	2.2	6	0.6	317	100%	317	1166	2+1	C++	Browser
	Firefox (v)	3	18.1	1	40.3	2	32.7	4	7	7	0	5	1.6	6	0.3	315	100%	315	1166	2+1	C++	Browser	
	PHPMYAdmin	2	33.3	1	38.7	3	19.4	4	5.4	5	1.1	5	1.1	5	1.1	93	100%	93	1166	2+1	PHP	Web app	
	PHPMYAdmin (v)	2	32.3	1	48.4	3	17.2	4	2.2	5	0	5	0	5	0	93	100%	93	1166	2+1	PHP	Web app	
Sotiropoulos et al., 2017	OpenRobots (pom, p3d, dtm)	2	21.1	4	6.1	1	21.2	5	3	3	3	3	15.2	-	-	23	70%	33	33	4*	C	Robot platform	
	Mirantis	1	49.3	6	1.9	3	6.5	4	6.4	5	2.6	2	12.5	-	-	344	100%	344	3187	2	Puppet	Infrastructure as Code	
Rahman et al., 2018	Mozilla	1	36.5	2	17	5	7.7	3	10	4	8.4	6	1.9	-	-	558	100%	558	3187	2	Puppet	Infrastructure as Code	
	OpenStack	1	57.5	2	6.7	4	5.9	6	2.4	3	6.5	5	2.9	-	-	1987	100%	1987	3187	2	Puppet	Infrastructure as Code	
	Wikimedia	1	62.7	5	12	2	3.3	3	4.3	4	5.1	6	2.6	-	-	298	100%	298	3187	2	Puppet	Infrastructure as Code	

As we can see in Table IX, the diversity of applications and contexts studied is quite large, including browsers, operating systems, text editors, satellite software, game emulators, machine learning tools, or build tools. In fact, ODC has been applied to a large number of domains, with the exclusion of storage intensive systems – which have now been handled in this work. The closest works (only from a domain perspective) may be the work by Børretzen and Dyre-Hansen, 2007 and Gupta et al., 2009 which mention transaction tools and document storage, respectively, but do not analyze NoSQL databases, or even databases in general. Alongside with the highly diverse domains, we also observe applications built in different programming languages, with Java being often present.

The number of defects studied in related work is usually around a few hundred per work. This is related with the huge effort required to perform the ODC classification. Note that not all of the defects used in each work can be used for this analysis, as in certain cases the authors customized ODC to their specific context and in the remaining cases some defects were not code or design defects, thus, are not classifiable with a ‘defect type’. We observed it is a common practice to use one or two raters (i.e., a human that takes care of all the classification process), however the vast majority of the works does not even specify how many raters were used. We were able to determine the number of human raters used for a few works, after contacting the authors. This count is marked with an asterisk under the column raters per defect. We also observed that when multiple raters are used, it is very rare that the authors discuss the inter-rater agreement.

Regarding the results obtained in previous works there are just a few observable trends, which is very likely the result of the high diversity of applications and contexts analyzed, aggravated by the general non-uniformity of each research work (e.g., dataset size, procedure applied, number of raters used). The results observed in related work strongly suggest that *there are no a priori expectable results for a certain system*, as the large number variables involved seem to play an important role and, in some cases, are difficult to measure, namely in what concerns the whole development process used (unless, in such cases, the system is being developed in a very confined or controlled manner, such as in some mission-critical systems). This calls in for the need of having new studies that specifically target new systems, such as NoSQL databases, that are being built in relatively uncontrolled environments, as it is usual in open source projects.

Despite the huge variability observed, there is a clear separation between “Timing/Serialization” and “Relationship” defect types and the remaining, as these to occupy the bottom places. “Interface/O-O Messages” and “Checking” occupy middle positions more often. “Assignment” defects are found more often occupying the second position and “Algorithm/Method” seems to be most frequently at the first position, although “Function/Class/Object” is also top defect. Other than this, *the specificity of the systems and the way they are built seems to have an overwhelming weight in the overall distribution of defect data*. This means that, *for activities where representative software faults need to be selected, it is necessary to previously analyze defect data for the particular context*, so that activities such as fault injection campaigns can be effectively carried out.

The data obtained here could be useful for software developers working in this kind of systems, allowing them to direct more the inspection and testing efforts on certain areas of the software systems being developed. This could not only improve development process and product quality for these databases, but also allow time to be saved in otherwise extensive defect finding/fixing activities.

VI. THREATS TO VALIDITY

In this section, we present threats to the validity of this work and discuss mitigation strategies. We start by mentioning the fact that, in this work, we have analyzed 4096 bug reports, which is *a subset of all reported bugs* affecting the studied NoSQL databases (about one fourth of the total number of bug reports). Thus, our results do not perfectly represent the whole population of bugs affecting these systems. This option for a subset of bugs was due to the huge amount of human effort involved, but still we must mention that, to the best of our knowledge, it is the largest dataset found among related work.

Each *bug report was, at a first stage, classified manually* by a single person, which may raise questions regarding the reliability of the dataset. This may introduce some error in the results either due to the lengthy human intervention required, or due to some subjectivity present in the process which may be due to the intrinsic characteristics of the ODC method itself, to the lack/erroneous information present in the bug report, and finally due to the technical expertise of the human that is applying the ODC method. To mitigate such issues, we took the following measures: i) the researcher involved in the classification process was trained using a set of 300 bugs that were discarded from the results; ii) we added a verification procedure carried out by the same person (also to have some understanding about the number of errors potentially present and as a means to improve the researcher classification skills); iii) we added an additional verification procedure carried out by two external researchers. Due to the huge amount of effort involved *both verification steps used a subset of the 4096 bugs* (a total of 40% of the bugs were reanalyzed), in which we found relatively low counts of errors and, in most of the cases, almost perfect inter-rater agreements.

In this work, we have *studied three NoSQL databases*, and thus the results cannot be generalized to all NoSQL database systems. Still, we have selected among the most popular ones, so that our results are meaningful to a potentially larger community of practitioners and researchers.

VII. CONCLUSION

In this paper, we analyzed software defects reported for three of the most popular NoSQL databases. The large number of defects analyzed using ODC, especially considering previous work (e.g., [9], [23], [44]), leads us to consider that these are representative of

the whole set of defects for these popular databases. However, due to the human intervention in the classification process, results may still include some error, which we tried to minimize, by carrying out verification activities.

We have discussed the *main findings* of this work, of which we highlight: i) First of all, the large variation in the distribution of defect types found in related work, which in practice means that we cannot assume, *a priori*, the existence of a certain defect distribution, which implies that this kind of study is carried out whenever it is important to know which are representative defects; ii) The results for NoSQL databases revealed, in terms of defect type prevalence, unique distributions for the types of defects (not visible in any of the related works analyzed); iii) We observed similarity in the single-attribute view of all three databases, also with, sometimes, a single value dominating the distribution. This, in conjunction with the previous finding and with the fact that all three databases are built and verified in a similar way (i.e., in a relatively uncontrolled open-source environment and by open and dynamic communities), means that the overall nature of the system appears to play an important role in the distribution of defects; iv) Testing activities are above than two times more frequent in reliability defects than in capability defects; checking defects are more frequent when the impact is reliability and are more likely to be associated with the “missing” qualifier; v) There are clear disparities in the distribution of defect types in different system components, which, in several cases, relate to the nature of the component (e.g., a replication component holding Timing/Serialization defects); vi) Certain types of defects are consistently associated with longer times to fix across all three databases (e.g., Function/Class/Object).

There were also a few *lessons learned* throughout the classification process and analysis of the results. It became obvious that effective training of the researcher that is responsible for performing the classification is crucial to have a good quality outcome in this kind of process. The contribution of additional raters is crucial, as it allow us to, on one hand, improve the final quality of the dataset, and, on the other hand, provides us with information regarding the possible presence of errors in the dataset. Also important to mention is that, there is a connection between two ODC attributes (Activity and Trigger) where certain triggers map to certain activities. This means that an incorrect activity may lead to an incorrectly classified trigger. The user applying ODC must take special attention when marking the values for these attributes, which we found to be a source of error. A difficult aspect is certainly related with understanding the exact causes of the results. Although we point out some possible reasons for the main issues raised by the results, following up with structured techniques like root cause analysis is out of the scope of this particular work, and is left for future work.

Overall, the resulting data represents vital information for NoSQL database developers, who could mostly benefit from it in defect prevention and removal efforts, thus contributing for more reliable systems. Additionally, researchers can use this data as a starting point in other works (e.g., fault injection experiments on systems of similar nature). As future work, we intend to explore the possibility of automating the classification procedure using machine learning algorithms.

ACKNOWLEDGMENTS

This work has been partially supported by the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 823788 (project ADVANCE); by the project METRICS, funded by the Portuguese Foundation for Science and Technology (FCT)—agreement no POCI-01-0145-FEDER- 032504 ; and by the project MobiWise: From mobile sensing to mobility advising (P2020 SAICTPAC/0011/2015), co-financed by COM- PETE 2020, Portugal 2020 - Operational Program for Competitiveness and Internationalization (POCI), European Union’s ERDF (European Regional Development Fund), and the Portuguese Foundation for Science and Technology (FCT). We would like to thank Henrique Marques and Gonçalo Carvalho, for their voluntary and fundamental help in using ODC to classify 820 defects. It would have not been possible to conclude this work without their valuable collaboration.

REFERENCES

- [1] N. Leavitt, “Will NoSQL databases live up to their promise?,” *Computer*, vol. 43, no. 2, 2010.
- [2] H. Marks, “Code Spaces: A Lesson In Cloud Backup,” *Network Computing*, 03-Jul-2014. [Online]. Available: <https://www.networkcomputing.com/cloud-infrastructure/code-spaces-lesson-cloud-backup/314805651>. [Accessed: 08-Nov-2017].
- [3] G. V. Hulme, “Amazon Web Services DDoS Attack And The Cloud,” *Dark Reading*. [Online]. Available: <https://www.information-week.com/news/229204417>. [Accessed: 08-Nov-2017].
- [4] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, “Basic concepts and taxonomy of dependable and secure computing,” *IEEE Trans. Dependable Secure Comput.*, vol. 1, no. 1, pp. 11–33, 2004.
- [5] P. Mellor, *CAD: Computer aided disasters*, vol. 1. 1994.
- [6] R. Chillarege, “Orthogonal Defect Classification,” in *Handbook of Software Reliability Engineering*, Chapter 9, 1996, pp. 359–400.
- [7] “IEEE Standard Classification for Software Anomalies,” *IEEE Std 1044-2009 Revis. IEEE Std 1044-1993*, 2010.
- [8] R. B. Grady, *Practical Software Metrics for Project Management and Process Improvement*. Prentice-Hall, 1992.
- [9] J. Durães and H. Madeira, “Emulation of Software Faults: A Field Data Study and a Practical Approach,” *IEEE Trans. Softw. Eng.*, vol. 32, no. 11, pp. 849–867, Nov. 2006.
- [10] F. Thung, S. Wang, D. Lo, and L. Jiang, “An Empirical Study of Bugs in Machine Learning Systems,” in *2012 IEEE 23rd International Symposium on Software Reliability Engineering*, 2012, pp. 271–280.
- [11] X. Xia, X. Zhou, D. Lo, and X. Zhao, “An Empirical Study of Bugs in Software Build Systems,” in *2013 13th International Conference on Quality Software*, 2013, pp. 200–203.
- [12] N. Silva, J. C. Cunha, and M. Vieira, “A field study on root cause analysis of defects in space software,” *Reliab. Eng. Syst. Saf.*, vol. 158, pp. 213–229, Feb. 2017.
- [13] MongoDB Inc., “MongoDB,” *MongoDB*. [Online]. Available: <https://www.mongodb.com/index>.
- [14] Apache Software Foundation, “Apache Cassandra,” *Apache Cassandra*. [Online]. Available: <http://cassandra.apache.org/>.

- [15] Apache Software Foundation, "Apache HBase," *Apache HBase*. [Online]. Available: <https://hbase.apache.org/>.
- [16] L. Zhi-bo, H. Xue-mei, Y. Lei, D. Zhu-ping, and X. Bing, "Analysis of software process effectiveness based on orthogonal defect classification," *Procedia Environ. Sci.*, vol. 10, pp. 765–770, 2011.
- [17] J. Agnelo, N. Laranjeiro, and J. Bernardino, "NoSQL ODC Dataset, Results, and Support Code," Mar-2019. [Online]. Available: <https://eden.dei.uc.pt/~cni/papers/2019-jss.zip>.
- [18] A. Wood, "Software reliability growth models," *Tandem Tech. Rep.*, vol. 96, no. 130056, 1996.
- [19] P. Wilson, L. Dell, and G. Anderson, *Root Cause Analysis: A Tool for Total Quality Management*. Milwaukee, Wisconsin: ASQ Quality Press, 1993.
- [20] IBM, "Orthogonal Defect Classification v 5.2 for Software Design and Code." 12-Sep-2013.
- [21] T. Basso, R. Moraes, B. P. Sanches, and M. Jino, "An investigation of java faults operators derived from a field data study on java software faults," in *Workshop de Testes e Tolerância a Falhas*, 2009, pp. 1–13.
- [22] J. Fonseca and M. Vieira, "Mapping software faults with web security vulnerabilities," in *2008 IEEE International Conference on Dependable Systems and Networks With FTCS and DCC (DSN)*, 2008, pp. 257–266.
- [23] N. Silva and M. Vieira, "Experience Report: Orthogonal Classification of Safety Critical Issues," in *2014 IEEE 25th International Symposium on Software Reliability Engineering*, 2014, pp. 156–166.
- [24] X. Xuan, X. Zhao, Y. Wang, and S. LI, "An Empirical Study of Bugs in Industrial Financial Systems," *IEICE Trans. Inf. Syst.*, vol. E98.D, pp. 2322–2327, Dec. 2015.
- [25] P. J. Morrison, R. Pandita, X. Xiao, R. Chillarege, and L. Williams, "Are vulnerabilities discovered and resolved like other defects?," *Empir. Softw. Eng.*, vol. 23, no. 3, pp. 1383–1421, Jun. 2018.
- [26] A. Rahman, S. Elder, F. H. Shezan, V. Frost, J. Stallings, and L. Williams, "Categorizing Defects in Infrastructure as Code," *ArXiv180907937 Cs*, Sep. 2018.
- [27] R. Lutz *et al.*, "Research infusion collaboration : finding defect patterns in reused code," Technical Report, Dec. 2004.
- [28] J. Christmansson and R. Chillarege, "Generation of an error set that emulates software faults based on field data," in *Fault Tolerant Computing, 1996. Proceedings of Annual Symposium on*, 1996, pp. 304–313.
- [29] J. A. Børretzen and J. Dyre-Hansen, "Investigating the Software Fault Profile of Industrial Projects to Determine Process Improvement Areas: An Empirical Study," in *Software Process Improvement*, 2007, pp. 212–223.
- [30] A. Gupta, J. Li, R. Conradi, H. Rønneberg, and E. Landre, "A Case Study Comparing Defect Profiles of a Reused Framework and of Applications Reusing It," *Empir. Softw. Engg.*, vol. 14, no. 2, pp. 227–255, Apr. 2009.
- [31] N. Silva and M. Vieira, "Software for Embedded Systems: A Quality Assessment Based on Improved ODC Taxonomy," in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, New York, NY, USA, 2016, pp. 1780–1783.
- [32] T. Sotiropoulos, H. Waeselyneck, J. Guiochet, and F. Ingrand, "Can Robot Navigation Bugs Be Found in Simulation? An Exploratory Study," in *2017 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, 2017, pp. 150–159.
- [33] "DB-Engines Ranking - popularity ranking of database management systems." [Online]. Available: <https://db-engines.com/en/ranking>. [Accessed: 01-May-2017].
- [34] "SQL & NoSQL Ranking - StackOverkill." [Online]. Available: <http://www.stackoverkill.com/ranking/sql-nosql>. [Accessed: 01-May-2017].
- [35] "Top NoSQL Database Engines." [Online]. Available: <http://www.kdnuggets.com/2016/06/top-nosql-database-engines.html>. [Accessed: 01-May-2017].
- [36] K. Henningsson and C. Wohlin, "Assuring fault classification agreement-an empirical evaluation," in *Empirical Software Engineering, 2004. ISESE'04. Proceedings. 2004 International Symposium on*, 2004, pp. 95–104.
- [37] K. El Emam and I. Wiczorek, "The repeatability of code defect classifications," in *Software Reliability Engineering, 1998. Proceedings. The Ninth International Symposium on*, 1998, pp. 322–333.
- [38] J. Cohen, "A Coefficient of Agreement for Nominal Scales," *Educ. Psychol. Meas.*, vol. 20, no. 1, pp. 37–46, Apr. 1960.
- [39] J. R. Landis and G. G. Koch, "The Measurement of Observer Agreement for Categorical Data," *Biometrics*, vol. 33, no. 1, pp. 159–174, 1977.
- [40] M. R. Lyu, Z. Huang, S. K. S. Sze, and X. Cai, "An empirical study on testing and fault tolerance for software reliability engineering," in *14th International Symposium on Software Reliability Engineering, 2003. ISSRE 2003.*, 2003, pp. 119–130.
- [41] N. Laranjeiro, M. Vieira, and H. Madeira, "A Robustness Testing Approach for SOAP Web Services," *J. Internet Serv. Appl. JISA*, vol. 3, no. 2, pp. 215–232, Sep. 2012.
- [42] "Using and Understanding Mathematics: A Quantitative Reasoning Approach." [Online]. Available: <https://www.pearson.com/us/higher-education/product/Bennett-Using-and-Understanding-Mathematics-A-Quantitative-Reasoning-Approach-5th-Edition/9780321652799.html>. [Accessed: 14-Feb-2019].
- [43] Apache Software Foundation, *Apache Cassandra Unit Testing Code*. The Apache Software Foundation, 2019.
- [44] N. Silva, F. Moreira, J. Carlos Cunha, and M. Vieira, "A Process for Finding and Tackling the Main Root Causes that Affect Critical Systems Quality," in *Certifications of Critical Systems - The CECRIS Experience*, Chapter 4: River Publishers, 2017, pp. 81–102.