Fault Triggers in Open-Source Software: An Experience Report

Domenico Cotroneo*, Michael Grottke†, Roberto Natella*, Roberto Pietrantuono*, Kishor S. Trivedi‡

*DIETI department, Università degli Studi di Napoli Federico II, Via Claudio 21, 80125, Naples, Italy.
†Friedrich-Alexander-Universität Erlangen-Nürnberg, Lange Gasse 20, 90403 Nuremberg, Germany.
‡Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA.
Email: {cotroneo, roberto.natella, roberto.pietrantuono}@unina.it, Michael.Grottke@fau.de, kst@ee.duke.edu

***Abstract*—With software systems becoming increasingly large and complex, many difficulties in coping with software bugs arise for developers. Despite good development practices, thorough testing, and proper maintenance policies, a non-negligible number of bugs remain in the released software. Understanding the type of residual bugs is fundamental for adopting proper countermeasures in current and future software releases. Depending on the *fault triggering conditions* that lead to a failure, developers can introduce fault-tolerance mechanisms and plan verification and validation strategies.**

In this paper, we analyze bugs in four large open-source software systems during their lifecycle, based on the concept of *fault triggers*. We first investigate how the type of system affects the bug type proportions, and their evolution over years. Then, an analysis of bug subtypes is performed, so as to better understand their nature, followed by a comparison with respect to attributes such as their average time to fix and severity.

***Index Terms*—aging-related bug; Bohrbug; Mandelbug; fault classification; fault trigger**

I. Introduction

Failures (i.e., deviations of service delivered from correct service) and *partial failures* (i.e., degradations in functionality or performance) experienced during software usage are preceded by incorrect internal states of the software, known as *errors* [1]. Such errors can be caused by *faults* in the software [1], also referred to as *defects* or *bugs*. By *fault trigger* we mean the set of conditions activating a fault and propagating the resulting error(s) into a failure. Often, fault triggers are complex, involving for instance the timing of events and interactions with other systems, which leads to failures that are very hard to reproduce [2], [3]. These faults are difficult to spot with traditional dynamic testing techniques, since it can be challenging to control and explore complex fault triggers in a testing environment. Therefore, this kind of fault requires specific strategies to be dealt with. Examples are fault tolerance strategies that mask faults, for instance by reinitializing the software state and retrying the failed operation [4], [5], [6], [7], and verification techniques that do not need to actually reproduce the fault trigger during execution, such as code reviews and model checking [8], [9]. Faults that can easily be triggered, instead, require more thorough testing in order to improve reliability. The objective of this study is to analyze the characteristics of bugs reported in large open-source software (OSS) projects.

Past research studies attempted to classify bugs according to this view, by adopting different terminologies with slightly different meanings, such as *hard* vs. *soft faults* [3], or *transient* vs. *non-transient bugs* [10]; other studies focused on *concurrency bugs* [11], [12], [13], [14] as the class of faults causing failures that are difficult to reproduce. To examine fault triggers in a comprehensive way, Grottke and Trivedi [15] developed the following definitions concerning the conditions related to the *fault activation* and the *error propagation*:

- **Bohrbug**: a bug which can easily be isolated and which manifests consistently under a well-defined set of conditions, because its activation and error propagation lack "complexity".
- **Mandelbug**: a bug whose activation and/or error propagation are "complex", where *complexity* can be caused by the possibility of a time lag between the fault activation and the failure occurrence, or by the possible influence of indirect factors, such as the interactions of the software application with its system-internal environment (hardware, operating system, or other applications), the timing of inputs and operations (relative to each other), and the sequencing of inputs and operations.

There is a further subtype of Mandelbugs, that is responsible for a phenomenon increasingly being studied, known as *process aging* [16] or *software aging* [17]. Software aging is a typical problem of long-running software systems in which an increasing failure rate and/or degraded performance is observed. This Mandelbug subtype is defined as follows.

- **Aging-related bug**: a Mandelbug that is capable of causing an increasing failure rate and/or degraded performance, because the rate at which it is activated and/or the rate at which errors caused by it are propagated into (partial) failures increases with the total time the system has been running. Often, such an increasing error propagation rate is caused by the accumulation of internal error states. Since aging-related bugs are a subtype of Mandelbugs, each Mandelbug is either an *aging-related bug* or a Mandelbug that does not cause software aging, called a *non-aging-related Mandelbug*.

Note that these definitions do not focus on the circumstances of one specific manifestation of the bug (e.g., the one that

made the testers notice its presence, or that helped them locate it in the code), but rather on its *potential* manifestation characteristics and its *inherent* features [18]. For example, even if a developer is able to reproduce a failure in a well-controlled environment, the underlying fault is classified as a Mandelbug if its manifestation can result in a transient failure at the user site because one of the criteria of complex fault activation or error propagation (as laid out above) applies. Similarly, even if an aging-related bug is detected before it has had the chance to actually lead to a decreasing performance (e.g., during a code inspection), the fault is still considered aging-related.

Adopting the above classification, we performed an extensive analysis of fault triggers in four large OSS projects: the Linux kernel, the MySQL database management system, the Apache HTTPD server, and the Apache AXIS Web services framework. From publicly-available bug repositories, we inspected problem reports describing the failure occurrences observed, the underlying bugs, and their fix. Based on our classification procedure, we classified each unique fault as Bohrbug (**BOH**), non-aging-related Mandelbug (**NAM**), or aging-related bug (**ARB**), and then analyzed its features. Moreover, we extended the classification to provide additional insights about subclasses of the NAM and ARB categories.

The analysis of these OSS projects provides evidence of Mandelbugs also in OSS systems which are often adopted in mission- and business-critical applications, and are quite different from NASA mission systems on which recent studies have focused [18]. Moreover, analyzing projects of different nature (e.g., in terms of type of system, size, and programming language) allows to relate the type of system with the type of bugs. Furthermore, dealing with OSS projects enables us to publicly release our data to the research community[1], allowing other researchers to adopt the classification more easily, and to carry out further analyses based on our data.

Our analysis reveals the following main findings:

- The proportion of Mandelbugs significantly varies among systems, and can be close to the proportion of Bohrbugs as in the case of the Linux kernel; the proportions can be related to both the size and the nature of the system.
- In every project, the proportion of Mandelbugs seems to converge to a constant value during the lifecycle, although past studies hypothesized that the percentage of Mandelbugs should be predominant in the long term. This may be explained by ineffective quality assurance and testing activities, and by the possibility of introducing new bugs during the project lifecycle.
- The analysis of subtypes indicates that timing-related faults are the largest part of Mandelbugs, although other Mandelbugs, such as those involving interactions with other software and hardware, account for a remarkable share. Similarly, memory-related aging-related bugs predominate, although leaks related to system-dependent data structures are also frequent.
- The time to fix a bug is significantly affected by the bug type, and strategies specifically tailored for Mandelbugs would certainly help.

[1] http://goo.gl/aeKoGR

The paper is organized as follows. In Section II, we discuss relevant studies on bug classification and empirical analyses. In Section III, we describe the approach adopted for classifying bug reports. Section IV provides the results from our analysis of four OSS systems. Section V closes the paper.

II. Related Work

In the last decades software bugs in large and complex systems have extensively been studied for several purposes. Most of the studies were aimed at understanding and characterizing bugs in terms of their location in the code and their features, in order to focus testing on the most fault-prone parts of the code [19], [20], [21], [22]. However, bugs were classified in a way not meant to be generic, but specific to a class of systems.

The study published in [23] proposed a wider classification of software faults that offered a large set of bug types, stressing the notion of fault morphology (i.e., "what is") in terms of code and separating the classification from aspects such as causes ("why it was done"). *Orthogonal Defect Classification* (ODC), by Chillarege *et al.* [24], is a well-known classification scheme for obtaining insights on the development process from the distribution of defect attributes. Attributes include the *defect type*, which reflects the fix made by the programmer. The definition of *defect types* is based on cause-effect relationships between them and the development phase in which defects originate: for example, an excessive number of "function" defects (that is, missing or incorrect functionalities of the system) indicates that high-level design phases should be improved. In a similar way, the *defect trigger* attribute relates the fault to the activities that made it surface during verification and validation (V&V), e.g., code reviews, white- and black-box tests, or in the field, and provides feedback on the V&V process. It is important to note that ODC defect triggers are related to the conditions that *actually* led to the discovery of a specific defect, while the notion of *fault trigger* adopted in this paper refers to the *potential* manifestation characteristics of a fault.

The focus of this work is on defect characteristics that are related to fault triggers and reproducibility, rather than the perspective of the developers that were involved in the mistake. Fault reproducibility was first thoroughly discussed in the seminal work by Gray [3], who distinguished between "solid" (or "hard") faults, for which failure occurrences are easily reproducible, and "elusive" (or "soft") faults, whose activation is not systematically reproducible. Gray named the former type "Bohrbugs", alluding to the physicist Bohr and his rather simple atom model, and the latter one "Heisenbugs", referring to Heisenberg and his uncertainty principle. Grottke and Trivedi [15] revised this nomenclature based on the finding that the term Heisenbug had originally been coined in the 1960s by Lindsay (while working with Gray), referring to "bugs in which clearly the behavior of the system is incorrect, and when you try to look to see why it's incorrect, the problem goes away" [25]; this is a more narrow definition than the one

published later by Gray. Instead of Heisenbug, Grottke and Trivedi identified *Mandelbug* as the complementary antonym of Bohrbug. The "complex" fault activation and/or error propagation of a Mandelbug as defined in Section I endow it with the potential to behave (apparently) non-deterministically or chaotically; this explains the name, alluding to the fractal innovator Mandelbrot. Heisenbugs (in the sense of Lindsay) are a subtype of the more general class of Mandelbugs. Likewise, aging-related bugs are a Mandelbug subtype. For example, for those ARBs causing the accumulation of internal error states there needs to be a time lag between fault activation and the failure occurrence.

Several field data studies analyzed fault reproducibility, even though not adopting the Bohrbug/Mandelbug terminology, and identified various sources of transient behavior of faults that can be related to Mandelbugs, such as: concurrency; timing of external events (e.g., from hardware); wrong memory state; resource leaks [2], [10], [12], [14], [26]. Performing such studies has proven to be a daunting task, since data about non-reproducible faults is hard to collect by their nature.

Notwithstanding the difficulties in finding evidences of Mandelbugs, the cited field studies confirmed that they account for a significant part of bugs in complex software. The analysis of field failures in Tandem computer systems [2], [3] showed that most software failures were caused by bugs whose features are referable to Mandelbugs. More recent analyses found that, although Bohrbugs represent the majority of software faults, Mandelbugs account for a significant share (in the 20-40% range) [10], [27]. Recent work [18] confirmed this result for NASA missions, and also showed that the proportions of Bohrbugs (and, consequently, of Mandelbugs) for different missions seem to stabilize around almost the same value. Moreover, it was found that ARBs represent a non-negligible share of faults even in mission-critical software (4.4%). These results, other than emphasizing the importance of Mandelbugs in such systems, highlight that even in critical and well-tested software a large proportion of Bohrbugs can still remain during the operational phase. In this paper, we carry out the classification based on the criteria adopted by Grottke *et al.* [18], but at a greater detail and focusing on four diverse open-source software applications. We aim at investigating more surgically Mandelbugs than in past work, e.g., by identifying subtypes of NAMs and of ARBs, and by investigating the association between the bug type and other attributes, such as the size and the nature of the project, the manifestation time, the fixing complexity, and the assigned severity.

III. Approach

A. Extended bug type classification

To classify bugs more in detail, we define the following subtypes of a non-aging-related Mandelbug (NAM), based on the different kinds of complexity in fault triggering conditions:

- **LAG**: there can be a time lag between the activation of the fault and the occurrence of a failure;
- **ENV**: the activation and/or error propagation is influenced by the interactions of the software application with its system-internal environment;
- **TIM**: the activation and/or error propagation is influenced by the timing of inputs and operations;
- **SEQ**: the activation and/or error propagation is influenced by the sequencing (i.e., the relative order) of operations.

Of course, these subcategories could also be employed for ARBs. However, it can be expected that the LAG subclass would apply to almost all of them. It is thus more informative to distinguish ARBs according to the various underlying reasons for the software aging phenomenon. Based on our definition of an ARB as well as on the software aging literature [28], [29], we identify the following ARB subtypes:

- **MEM**: ARBs causing the accumulation of errors related to memory management (e.g., memory leaks, buffers not being flushed);
- **STO**: ARBs causing the accumulation of errors that affect storage space (e.g., the bug consumes disk space);
- **LOG**: ARBs causing leaks of "other logical resources", that is, system-dependent data structures (e.g., sockets or inodes that are not freed after usage);
- **NUM**: ARBs causing the accumulation of numerical errors (e.g., round-off errors, integer overflows);
- **TOT**: ARBs in which the increase of the fault activation/error propagation rate with the total system run time is not caused by the accumulation of internal error states (e.g., due to a bug in the Patriot missile defense system [28], [30], the system runtime was incorrectly processed, but the error produced was only propagated into a failure if the system had been running for more than eight hours; error states did not accumulate).

B. Bug sources

We considered four open-source software systems with public and actively-used bug repositories, which provided us with a large number of bugs for the analysis. The chosen software systems are widely adopted in business-critical contexts [31], and they cover different types of software: (i) the Linux kernel, a feature-rich OS used in several domains, from embedded systems to supercomputers; (ii) MySQL, one of the most-used database management systems, accounting for a significant market share among IT organizations; (iii) the Apache HTTPD server, and (iv) the Apache AXIS framework for Web services, adopted by many companies for running their Web applications.

Since these systems are very large and have been around for a long time, tens of thousands of problems have been reported by their users; hence it is unrealistic to analyze all of them.

We therefore selected a subset of these components for each project, and focused the analysis on the problem reports related to them. The selected components/subsystems were: *Network Drivers*, *SCSI Drivers*, *EXT3 Filesystem*, and *Networking/IPV4* for Linux; *InnoDB Storage Engine*, *Replication*, and *Optimizer* for MySQL; *Apache httpd_core*, *Apache httpd_mod_proxy*, *Apache httpd_mod_cgi*, *Apache httpd_mod_ssl* for Apache

TABLE I: Overview of the considered projects.

Project	Language	LoC (selection)	LoC (project)	#reports (selection)	#reports (project)	Time frame
Linux	C	1.31M	9.58M	346	3914	Jul 2003 - May 2011
MySQL	C/C++	453K	1.1M	244	894	Aug 2006 - Feb 2011
HTTPD	C	145K	195K	157	405	Mar 2002 - Oct 2007
AXIS	Java	80K	80K	216	226	Jul 2001 - Nov 2005

HTTPD; for Apache AXIS, we inspected all reports but those related to the *Distribution*, *Documentation*, and *Samples* areas (i.e., only reports about problems that can affect the system during its execution). In the selection, we accounted for the relevance of subsystems/components in terms of usage and number of problem reports, as well as for the coverage of diverse functionalities of the system and of a significant share of the system code. TABLE I provides, for each project, its programming language, the total size in LoC (computed using the `sloccount` utility) of the whole project and of the considered components, the number of problem reports that have been marked as "fixed" and "closed" by the developers (i.e., a fix was found and included in the source code), and the time frame during which these reports were issued.

The bug repositories[2] provide a large amount of information. Although projects may slightly differ with respect to the gathered information (e.g., they use different versioning schemes or fault severity scales), all reports provide:

- the type of the problem report (e.g., if it is a bug report or a request for a new feature);
- the date it was opened, closed, and last modified;
- the severity of the problem (i.e., the perception of its effects by users and developers);
- the version(s) affected by the problem;
- the component or subsystem affected by the problem;
- some textual messages describing the effects of the problem, its diagnosis by developers, and information on whether and how it can be reproduced (e.g., the inputs for triggering the failure behavior at the user's site, or a test case accompanying the bug fix);
- the status (e.g., the problem has been assigned to a developer, it has been solved, etc.).

To work with reliable bug descriptions, we filtered the bug repository, focusing on problem reports that had been solved (i.e., marked as "fixed" and "closed"). Moreover, we restricted the analysis to reports related to the abovementioned components as well as to stable and mature system versions; in particular, the considered versions were Linux 2.6, MySQL 5.1, Apache HTTPD 2, and Apache Axis 1.

C. Classification procedure

Given a problem report, to classify the related bug as a Bohrbug, a non-aging-related Mandelbug, or an aging-related bug, we conducted a manual analysis by examining the textual descriptions and, if available, the test case to reproduce the failure occurrence, the available patches, and additional information attached to the problem report. We defined a classification procedure consisting of the following steps to classify faults in a rigorous way:

1) The problem report was first examined to make sure that it was related to a unique bug; i.e., problems turning out to be operator errors, requests for software enhancements, and duplicates were removed from the analysis. A report was considered a duplicate if either a field in the report or the textual description indicated that the reported problem was caused by the same underlying bug as another report already included in our study.
2) The report was then searched for any information on the activation conditions of the bug (e.g., the set of events and/or inputs required to trigger errors), its error propagation (e.g., how the bug affected the program state and how an erroneous state propagated through the running system), and the failure behavior (e.g., the bug effects perceived by the users).
3) The bug was classified as an ARB if there were indications that the rate with which it is activated and/or the rate with which errors caused by it are propagated into (partial) failures can be an increasing function of the total time the system has been running (e.g., the report refers to leakage and/or gradual corruption of resources, or to the accumulation of numerical errors). Typically, the information in the failure report allowed us to determine the ARB subtype (MEM, STO, LOG, NUM, TOT) as well (e.g., because it was reported that memory expected to be freed had not been freed). Sometimes, it was merely known that the failure rate of a bug tended to increase over time (e.g., because it caused a failure only after a certain function had been called multiple times), but there was not enough information about the exact failure mechanics, like the presence of error accumulation. In such a case, we classified the bug as an ARB of unknown subtype (**ARU**).
4) A bug that was not an ARB was classified as a NAM if we found indications that one of the types of "complexity" of the activation and/or error propagation, embodied in the four subtypes LAG, ENV, TIM, and SEQ, applied to it. Sometimes, we did not have sufficient information about the activation and error propagation conditions of a bug that was reported to sporadically cause failures that could not be reproduced. We then classified this bug as a NAM of unknown subtype (**NAU**).
5) If there was evidence that the bug was neither an ARB nor a NAM, we classified it as a Bohrbug (BOH).

[2]Available at https://bugzilla.kernel.org (Linux 2.6), http://bugs.mysql.com (MySQL 5.1), https://issues.apache.org/bugzilla (Apache HTTPD 2), and https://issues.apache.org/jira/secure/IssueNavigator.jspax (Axis 1).

TABLE II: Examples of NAMs and ARBs.

Project	Bug ID	Type	Description
MySQL	54453	NAM/SEQ	"if you 'alter table .. rename to ..' on a table that has an active transaction open and UNIV_DEBUG is defined, mysqld crashes"
Linux	7207	NAM/LAG	"[The e1000 network driver at suspend/resume does not] explicitly free and allocate irq [...] Restarting the network solved the problem"
HTTPD	8184	NAM/ENV	"The error only occurs intermittently [...] It behaves as if requests are being distributed (via round-robin or the like) and handled sometimes by a worker thread that is not properly initialized"
AXIS	1270	ARB/MEM	"Strings and char[]s are being leaked"
Linux	32832	ARB/LOG	"In 2.6.35 and earlier, shutdown(2) will fully remove a socket. This does not appear to be true any more and is causing software to misbehave."
HTTPD	13511	ARB/STO	"Apache child processes will die trying to write logs which have reached 2GB in size."

6) Sometimes, a report did not contain sufficient details to classify the underlying bug as an ARB, NAM, or BOH. It was then labeled as a bug of unknown type (**UNK**).

During the manual analysis, further information was extracted for the purpose of our analyses: the time at which the report was opened and closed, and the severity stored in the bug repository. To clarify the classification, TABLE II shows examples of NAMs and ARBs, along with statements from the reports that provide information about the fault activation, the error propagation, and the failure occurrence.

IV. Analysis

Our inspection of problem reports, using the procedure previously described, provided a large set of bug data. We first examine the relative frequencies of the bug types in the considered projects, and relate the results with the features of the considered projects. Subsequently, we analyze bug types with respect to some relevant features, including the time to fix the bugs and their severity.

A. Bug type proportions

TABLE III summarizes the absolute numbers and the percentages of each bug type. Among the 963 problem reports, we identified 852 *actual bugs*: these reports include neither operator errors nor duplicates nor problems that do not affect the operational software, such as documentation and compile-time issues. A subset of 816 bugs was classified as BOH, NAM, or ARB. The remaining bugs, which we refer to as UNK, were lacking information for classifying them with certainty. Most of these bugs belong to the Linux project: for this system, some problem reports do not provide a precise diagnosis of the bug, since the related failure disappeared in newer versions of the system (e.g., the bug did not manifest itself anymore after a major rewrite of a module or subsystem).

Comparing the projects with respect to their relative percentages of BOH, NAM, and ARB, it is possible to notice significant differences. In the Linux project, there are more Mandelbugs than Bohrbugs (NAMs and ARBs together account for more than 50% of all bugs), while in the remaining projects the percentage of Bohrbugs is predominant: this percentage ranges between 56.6% and 92.5%. We believe that the first cause for this result is the different nature of the considered projects: Linux and operating systems in general are tightly related to hardware devices; this makes them more prone to incorrect interactions with the hardware and to bugs in event handling, which can lead to transient failures. Another reason for the high percentage of Mandelbugs is the presence of several complex and tightly-interacting subsystems in the Linux kernel. It seems that the proportion of NAMs decreases as we move up in the "software stack"; that is, the proportion is higher for "low-level" code, such as an operating system, and lower for "high-level" code, such as middleware for web applications. This can be expected since in "high-level" code there are fewer interactions with the hardware and less resource management burdens.

We observe that the percentage of ARBs is approximately the same for the Linux, MySQL, and HTTPD projects. Instead, for AXIS the percentage of ARBs is lower than for the other three projects. This can be explained by considering the kind of system and by the fact that AXIS has been developed using the Java language, which provides automated memory management through garbage collection. By contrast, the other three projects adopted the C and C++ languages, in which memory management is handled by developers, and which are therefore more prone to software aging issues. However, it is important to note that Java software is also subject to software aging, even in the presence of garbage collection: This happens in the case of objects that are no longer needed but still referenced, which prevents the garbage collector from reclaiming them [32].

Another perspective on the relative importance of the bug types is given by TABLE IV, which provides an estimated fault density, expressed in faults per kLoC, for each type of bug and each project. To obtain these fault density estimates, we divided the estimated total number of bugs of each type (calculated by multiplying the total number of bug reports, including reports that are still open and UNK reports, with the respective bug type proportion among all *classified* bugs) by the LoC of the considered components shown in TABLE I. This computation necessarily makes the assumption that the bug type proportions among the reports classified are the same as among those bug reports not yet closed or not classified. Of course, the considered projects exhibit different *#bugs/kLoC* ratios, which is a result of the software development process and of quality assurance activities. Nevertheless, we can notice different trends for individual bug types. In fact, the ratio *#BOH/kLoC* decreases fast with an increase in the project size; instead, the decrease in the *#ARB/kLoC* and *#NAM/kLoC*

TABLE III: Total numbers and percentages for each bug type.

Project	#actual bugs	#classified bugs	#BOH	#NAM	#ARB	#UNK	%BOH	%NAM	%ARB	%UNK
Linux	289	267	122	121	24	22	42.2	41.9	8.3	7.6
MySQL	221	209	125	67	17	12	56.6	30.3	7.7	5.4
HTTPD	143	141	116	15	10	2	81.1	10.5	7.0	1.4
AXIS	199	199	184	7	8	0	92.5	3.5	4.0	0.0

TABLE IV: Estimated fault densities for each bug type.

Project	#bugs/kLoC	#BOH/kLoC	#NAM/kLoC	#ARB/kLoC
Linux	0.3434	0.1569	0.1556	0.0309
MySQL	0.4939	0.2954	0.1583	0.0402
HTTPD	3.1054	2.5548	0.3304	0.2202
AXIS	26.1994	24.2245	0.9216	1.0532

ratios is slower. Therefore, when a large software project is considered, the fault densities for Bohrbugs and Mandelbugs may be similar, while for smaller projects the fault density for Bohrbugs tends to be higher. Note that in our sample of projects, there is a high dependency between code size and the kind of system (e.g., Linux is both a large and a "low-level" software); we therefore cannot separate these effects, and both are likely to have an influence on the fault densities.

Fig. 1, Fig. 2, and Fig. 3 show the evolution of the BOH, NAM, and ARB proportions among the classified bugs during project life. For all the projects, the proportions stabilize around a constant value after about two years after project birth. Another interesting result is that the ARB proportions also settle to constant values, which are about the same for three of the four projects (as reported above).

It could have been expected that the proportion of Mandelbugs increases with time: according to Gray's conjecture [3], most of the software faults remaining after thorough testing and years of production are Mandelbugs due to their transient manifestation. However, there are other aspects that have to be taken into account to explain this result. The proportion of Bohrbugs (or Mandelbugs) is not necessarily equal to the percentage of software *failures* caused by Bohrbugs (or Mandelbugs, respectively): the failures actually experienced by the users also depend on the *operational profile*, that is, the kind of system usage that is made by its users. For this reason, our results do not contradict past empirical studies that were concerned with *failures* rather than *faults*, and that reported Mandelbugs as a major cause of software failures [2]. Note that even if Bohrbugs are easy to reproduce and to debug once detected, they are still difficult to detect in large and complex software systems. This may be due to ineffective quality assurance and testing activities, or simply due to the fact that it is impractical to extensively test such large systems. As a result, a significant number of Bohrbugs can still be found after several years. Another possible factor is that OSS in general are continuously evolving during their lifetime, since new features keep being introduced by developers. Therefore, even if Bohrbugs are detected and fixed, more Bohrbugs could be introduced when changing or extending the software.

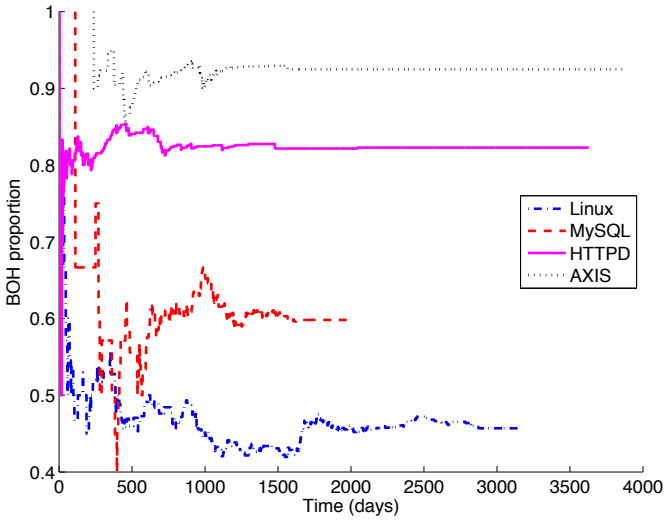

Fig. 1: Proportions of BOH among classified bug reports.

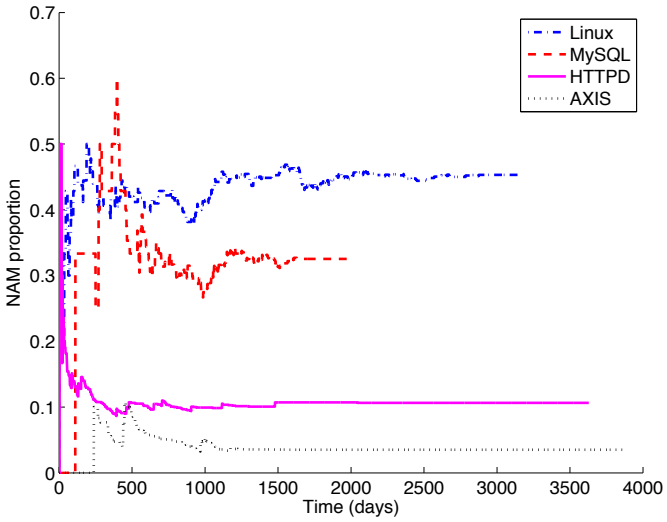

Fig. 2: Proportions of NAM among classified bug reports.

To better understand which factor is most influential on the observed trends in bug type proportions, we analyzed the release dates of minor and major versions of the considered projects. Fig. 4 shows the occurrences of minor and major releases for each project during the same time windows of Fig. 1, Fig. 2, and Fig. 3. It can be seen that for all four projects several minor releases (e.g., Linux 2.6.31, 2.6.32, . . .) occurred during the whole lifecycle. Instead, major releases occur rarely. Considering that after a major release (and even some time before it goes public) most development efforts are devoted

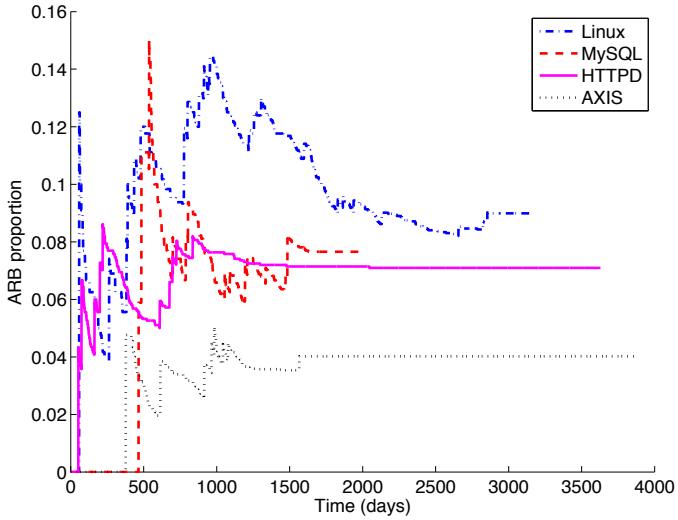

Fig. 3: Proportions of ARB among classified bug reports.

to the new major release (e.g., new important functionalities are introduced in MySQL 5.5 instead of MySQL 5.1), we can assume that the effects of code changes on bug type proportions are less significant from that time. In fact, when a major release occurs, minor releases are mostly focused on bug fixes and minor improvements, and are less likely to introduce new faults; therefore, the bug type proportions after a major release mainly reflect old bugs rather than new ones. In the case of MySQL and HTTPD, major releases occur in the middle of the lifecycle of a previous major version (e.g., the major version 5.5 of MySQL was released while MySQL 5.1 was still being updated and widespread among users), while for AXIS and Linux a major version is released after the end of the lifecycle of the previous major version (i.e., the lifecycles of two major releases do not overlap). For the MySQL and HTTPD projects, at the time of a new major release (at about 1250 days), a stabilization of bug type proportions has already occurred, and there does not seem to be an increasing trend in the proportion of NAMs; therefore, we attribute the stabilization of bug type proportions for these projects to old Bohrbugs that keep being discovered after some time rather than to new Bohrbugs introduced by late releases. For Linux and AXIS, significant changes may have occurred during their lifecycle, but the Bohrbug/Mandelbug proportions among the newly-introduced faults seem to be similar to those among the fixed faults.

In Fig. 5 and Fig. 6, we provide the proportions of NAM and ARB subtypes for all projects (omitting ARB/TOT, which never applied). As for Linux and MySQL, there is a predominance of timing-related bugs, which can be explained by the nature of these systems, where threads concur to access shared resources and have to be properly synchronized. Timing is a less significant problem in Apache HTTPD: although it is a multi-threaded software, there is a high degree of independence between threads, because they seldom have access to shared resources when handling HTTP requests, which renders synchronization problems much less frequent. Environment-

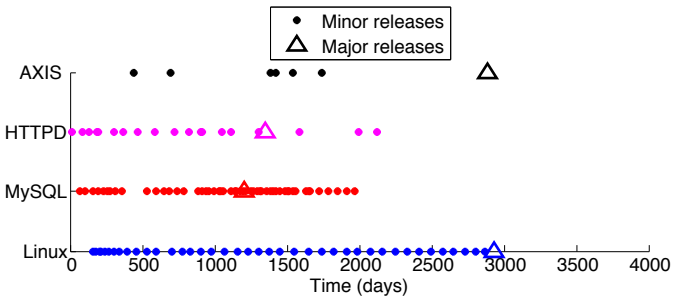

Fig. 4: Minor and major release dates of considered projects.

related faults are also a significant share of NAMs: in Linux, they are often related to hardware management, while in Apache HTTPD and AXIS they are related to the network and the filesystem. Bugs exhibiting a time lag before a failure only affect Linux and MySQL, which have a tendency towards data corruption problems that may cause failures only after these errors have propagated through the system. As for ARBs, there is a strong predominance of memory-related bugs (e.g., memory leaks). Leaks associated with storage and other logical resources were also found. Only few ARBs (a total of two bugs) were related to numerical problems, and in particular to integer overflows. This low number is probably due to the scarcity of floating point arithmetic in the considered projects, which is not used at all in the case of the Linux kernel [33].

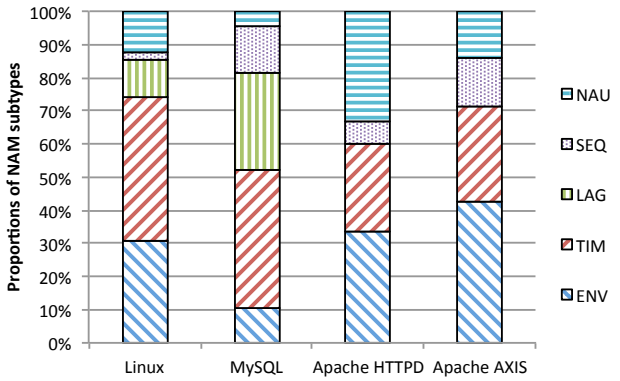

Fig. 5: Proportions of NAM subtypes.

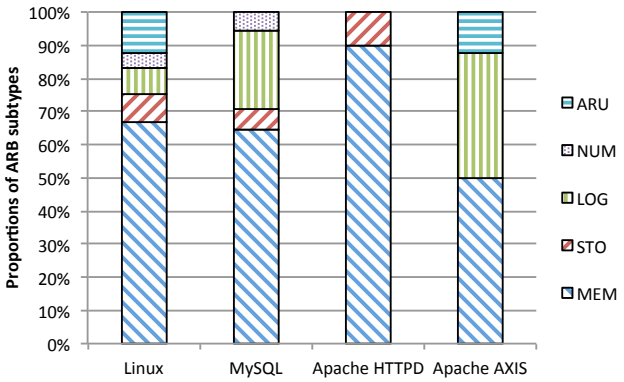

Fig. 6: Proportions of ARB subtypes.

TABLE V: Comparison of time to fix for bug types.

Project	Time to fix: avg. (std. dev.) BOH	NAM+ARB	Test result (adj. p-value) BOH vs. NAM+ARB
Linux	157.34 (226.21)	229.84 (304.24)	**reject** (0.0560)
MySQL	107.92 (176.76)	89.45 (109.60)	do not reject (0.2820)
HTTPD	99.09 (199.44)	116.65 (133.84)	**reject** (0.0973)
AXIS	111.19 (254.99)	186.50 (256.18)	**reject** (0.0973)

B. Time to fix

In order to understand the impact of bug types on the defect management process, we analyzed the time spent by developers on fixing bugs. We hypothesized that the type of a bug has a significant impact on the time to fix, since Mandelbugs tend to be more difficult to reproduce and to diagnose than Bohrbugs. We collected from each bug report the date at which it was issued, as well as the date at which the bug was considered definitively solved by developers, and computed the difference between these two dates. This time period includes the time spent by developers on reproducing the failure reported by the users, diagnosing its root cause, developing a fix for the bug, and validating the effectiveness of the fix through testing and user feedback. It does not include the time required for users to reproduce the failure *on-site* before a bug report is filed; however, we expect that the delay for this activity is no longer for Bohrbugs than for Mandelbugs, given their transient nature, and therefore the time to fix obtained from the bug reports should not bias the comparison in favor of Mandelbugs.

We compared the time to fix for Bohrbugs and Mandelbugs by means of the *Wilcoxon rank-sum test* [34], which assesses whether one of two independent samples tends to attain larger values. TABLE V provides the average and the standard deviation of the time to fix for each class of bugs. It also shows whether the null hypothesis that for both types of bugs the time to fix is sampled from the same distribution can be rejected at a type I error level of 10%, which is the case for a p-value below 0.1. Since multiple comparisons (one per project) are being performed, using unadjusted p-values for making a test decision would lead to a probability of *at least one false rejection* that is larger than the type I error level chosen. We therefore adopted the Benjamini-Hochberg procedure [35], controlling the false rejection probability and retaining a higher power of the tests, to derive adjusted p-values. Despite this correction, which makes the tests more conservative, the null hypothesis can be rejected for three of the four projects (Linux, HTTPD, and AXIS); in each of these cases, the time to fix tends to be greater for Mandelbugs (including both NAM and ARB classes) than for Bohrbugs. One possible cause for this finding is that upon a failure caused by a Mandelbug the developer often requires additional information to understand its nature and to detect the underlying bug in the code. Moreover, Mandelbugs may be located in components that are more difficult to maintain (e.g., a problematic code area that can be dealt with only by few developers in the team). As Mandelbugs seem to be more difficult and time-consuming to cope with than Bohrbugs, strategies specifically tailored for Mandelbugs should be useful to improve the reliability of software systems in a cost-effective way, both by means of fault tolerance mechanisms and by specific testing methods.

TABLE VI: Contingency tables for bug type and severity.

(a) Linux (outcome = do not reject, adj. p-value = 0.8276)

	BOH	NAM+ ARB
Blocking	9	11
High	18	30
Low	7	4
Normal	88	100

(b) MySQL (outcome = do not reject, adj. p-value = 0.8276)

	BOH	NAM+ ARB
Critical	28	17
Serious	41	29
Non-critical	55	36
Performance	1	2

(c) HTTPD (outcome = do not reject, adj. p-value = 0.8276)

	BOH	NAM+ ARB
Blocker	4	0
Critical	16	4
Major	24	7
Minor	8	0
Normal	62	14
Trivial	2	0

(d) AXIS (outcome = do not reject, adj. p-value = 0.4924)

	BOH	NAM+ ARB
Blocker	5	0
Closed	1	0
Critical	6	3
Major	78	6
Minor	15	0
Trivial	1	0

C. Bug severity

Finally, we compared the severity of bugs as perceived by users and developers, who can assign through the bug tracker system an indication of the "importance" of the bug in terms of consequences caused by it. We limited this analysis to severity because it is the only indicator of bug importance available for all the four considered projects. In every project, the severity is expressed using a severity *scale*. Since the scale is different for each project (in terms of the number and names of severity levels), the severities of the bugs of two different projects cannot be compared. We thus focus on analyzing the severity of bugs within the same project. TABLE VI provides the contingency tables for bug type and bug severity, which we analyzed in order to understand whether there is a bug type that is perceived to be more severe than the other one. We adopted *Fisher's exact test of independence* [34], assessing the null hypothesis that two variables are independent. Again, p-values were adjusted using the Benjamini-Hochberg procedure [35]. The null hypothesis cannot be rejected at a reasonable type I error level for any of the projects; the percentage of bugs across severity levels does not seem to be influenced by the bug type. Therefore, we conclude that, although Bohrbugs and Mandelbugs exhibit a different behavior, there is no evidence from the considered projects that their effects in terms of failure severity are perceived to be different. This can be explained by the fact that the distinction between Bohrbugs and Mandelbugs is concerned with fault triggering (e.g., the sequence of inputs or events that make the fault affect the system state), rather than the way in which a bug manifests itself to external users as failures. Therefore, different strategies are needed for dealing with each of them.

D. Limitations

Although the analysis is comforted by the extensiveness of the study, accounting for more than 900 problem reports, care must be taken when interpreting the results and drawing conclusions. While the classification procedure can easily be generalized, results are limited by the chosen applications and components, by the bugs considered, and by the quality of bug reports, like for any empirical study in this field. However, using the outlined criteria for bug selection, we selected a well-defined set of bugs (i.e., the entire set of fixed bugs) from the repositories, avoiding biases related to sampling (e.g., keyword-based sampling, random sampling), which was instead adopted in past empirical studies to deal with the huge number of bug reports [10], [11], [13], [14]. Moreover, we focused our attention on widely-used and diverse projects and components (e.g, we considered both "low-level" code, such as device drivers and a storage engine, and "high-level" code, such as a web framework).

The analysis does not include bugs that have not yet been fixed, since their reports may contain inaccurate or incomplete information. This could bias our estimates, since unfixed bugs may have properties different from the fixed ones; for instance, Mandelbugs may tend to be fixed less frequently than Bohrbugs. However, the previous study of NASA systems [18], for which *all reports* were analyzed due to the availability of detailed failure data, showed trends similar to the ones in this study; it is thus possible that our focusing on fixed bugs has not biased the results. Also, the analysis of fixed bugs found a remarkably large (absolute) number of Bohrbugs, highlighting that Bohrbugs are a serious problem even for mature systems.

V. Discussion

The classification and analysis of bug reports presented in this paper provide insights about how bugs manifest themselves during operation. These kinds of results are useful for (i) understanding the bug characteristics that make failures difficult to reproduce, and (ii) identifying the best countermeasures to cope with bugs during development, testing, and maintenance. The following findings help us on these issues.

The bug type proportions vary with the size and nature of systems. Although Mandelbugs (NAMs and ARBs) account for about 32.9% (25.7% and 7.2%, respectively) of all classified faults, which is in line with previous studies [18], [27], we noticed significant differences among systems. While the size of the software may have some influence, its kind seems to play an important role as well. In fact, non-reproducible behavior of bugs is often related to interactions of the systems with hardware and with low-level resource management. This observation is confirmed by the subsequent analysis of bug subtypes, which distinguishes between the causes of complexity; the *environment* and the *timing* of inputs and events (e.g., concurrency) represent the main subtypes of Mandelbugs, whereas the LAG class is a secondary cause. NAM/LAG bugs exhibit a long chain of events between fault activation and manifestation, which hinders systematic reproduction; they are related to coupling among system components and to code complexity. The predominance of ENV and TIM, along with the greater percentage of NAMs in Linux, suggests that Mandelbugs are more related to low-level interactions and resource management than to software size/complexity. Further analyses are needed for investigating the relationship between NAMs and software metrics. As for ARBs, results confirmed that memory-related problems are the main source of software aging, but the non-negligible percentage of ARBs connected with other system-dependent resources suggests pushing the research on software aging (today mainly focused on memory issues) to investigate other types of ARBs.

Within each project, the bug type proportions stabilize over the years. This finding contradicts the popular opinion that the prevalence of simple bugs (i.e., Bohrbugs) decreases with time, thus leading to an increase in the proportion of Mandelbugs; instead, an approximately constant proportion has been observed. The similarity of this behavior among projects (even in terms of the time to stabilization) suggests that both Mandelbugs and Bohrbugs keep being detected during the overall lifecycle of the product. This appears to be the consequence of ineffective verification activities, leaving many Bohrbugs in the code. It is, in fact, impractical to extensively test large systems. Analyzing the evolution of bug types during the lifecycle of large systems can provide feedback on the effectiveness of quality assurance and on the need for improvements. Another influential factor is the continuous evolution of open-source software, since maintenance actions, such as corrective actions and the introduction of new features, can introduce regression Bohrbugs, as well as Bohrbugs in new functionalities. This observation can be symptomatic of the need to improve maintenance activities: Bohrbugs represent a significant portion of faults and should not be neglected when operating on existing code. This means that more thorough analyses should be made to verify the effects of changes.

Mandelbugs take longer to fix, and require specific strategies to be dealt with. We found a statistically significant difference in the times to fix of Bohrbugs and Mandelbugs, respectively, for three out of four projects. In each of these cases, Mandelbugs tend to have a greater time to fix than Bohrbugs. Adopting strategies and tools for improving the diagnosis of Mandelbugs would improve the fixing time of such bugs. This is the case for the MySQL project, in which there is no statistically significant difference in the times to fix; we attribute this result in parts to the fact that MySQL developers make extensive use of the Valgrind debugging tool for tracking down NAMs and ARBs [36]. Moreover, Mandelbugs are by their nature difficult to detect by testing, and they require more specific techniques to be found during V&V. If the number of Mandelbugs found during operation is high, there are basically two alternatives. The first one is to employ additional V&V techniques for future releases, by introducing model checking, stress testing, code reviews. The second solution is to rely on runtime failure detection [37] and recovery mechanisms [7], [30], to compensate for the longer repair time of these bugs, and avoid system downtime while developers investigate the root cause of problems. Recovery mechanisms include: *restart*

of a component or a service; *reconfiguration* of components (e.g., migration to a diverse environment); *retry* operations. These strategies can be adopted depending on the system and failure type (e.g., a retry can succeed in the case of a timing bug in the software application, while a complete reboot is needed for bugs in the OS). Moreover, software aging issues can be prevented by *software rejuvenation* [16], a technique that proactively restarts a system in order to avoid the occurrence of aging failures.

The severities of bug types are perceived to be similar. The analysis of bug severity highlights that, despite the higher complexity of Mandelbugs that might endow them with more severe consequences, the failure severity assigned by developers shows no significant differences between Bohrbugs and Mandelbugs. These bug types differ in their failure reproducibility, not in their impact on the system. The relative importance of Bohrbugs and Mandelbugs during design, testing, and maintenance activities is therefore determined by their relative proportions. If there is a significant proportion of Mandelbugs, additional testing and recovery strategies are recommended, while more regression and functional testing may be needed if the proportion of Bohrbugs is high.

We believe that in the near future a deeper understanding of bugs from this perspective will be a driving factor in implementing policies for cost-effective software development. Our future work will therefore be devoted to relating these bug features to the software development process.

ACKNOWLEDGMENTS

This work was supported by the Dr. Theo and Friedl Schoeller Research Center for Business and Society, and by the projects TENACE (PRIN n. 20103P34XC) and SVEVIA (PON02 00485 3487758) funded by the MIUR.

REFERENCES

[1] A. Avižienis, J. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. Depend. Sec. Comp.*, vol. 1, no. 1, pp. 11–33, 2004.

[2] I. Lee and R. Iyer, "Software dependability in the Tandem GUARDIAN system," *IEEE Trans. Softw. Eng.*, vol. 21, no. 5, pp. 455–467, 1995.

[3] J. Gray, "Why do computers stop and what can be done about it?" Tandem Computers, Tech. Rep. 85.7, 1985.

[4] Y. Wang, Y. Huang, and C. Kintala, "Progressive retry for software failure recovery in message-passing applications," *IEEE Trans. Comp.*, vol. 46, no. 10, pp. 1137–1141, 1997.

[5] S. Maffeis and D. Schmidt, "Constructing reliable distributed communication systems with CORBA," *Comm. Mag.*, vol. 35, no. 2, pp. 56–60, 1997.

[6] F. Qin, J. Tucek, J. Sundaresan, and Y. Zhou, "Rx: treating bugs as allergies—a safe method to survive software failures," in *ACM SIGOPS Operating Systems Review*, vol. 39, no. 5, 2005, pp. 235–248.

[7] K. S. Trivedi, R. Mansharamani, D. S. Kim, M. Grottke, and M. Nambiar, "Recovery from failures due to Mandelbugs in IT systems," in *Proc. Pacific Rim Intl. Symp. Depend. Comp.*, 2011, pp. 224–233.

[8] V. Basili, S. Green, O. Laitenberger, F. Lanubile, F. Shull, S. Sørumgård, and M. Zelkowitz, "The empirical investigation of perspective-based reading," *Emp. Softw. Eng.*, vol. 1, no. 2, pp. 133–164, 1996.

[9] G. Holzmann, "The model checker SPIN," *IEEE Trans. Softw. Eng.*, vol. 23, no. 5, pp. 279–295, 1997.

[10] S. Chandra and P. M. Chen, "Whither generic recovery from application faults? A fault study using open-source software," in *Proc. 2000 Intl. Conf. Depend. Sys. Netwks.*, 2000, pp. 97–106.

[11] Z. Li, L. Tan, X. Wang, S. Lu, Y. Zhou, and C. Zhai, "Have things changed now? an empirical study of bug characteristics in modern open source software," in *Proc. Wksp. Arch. Sys. Supp. Improv. Softw. Depend.*, 2006, pp. 25–33.

[12] S. Lu, S. Park, E. Seo, and Y. Zhou, "Learning from mistakes: a comprehensive study on real world concurrency bug characteristics." in *Proc. Intl. Conf. Arch. Supp. Prog. Lang. Op. Sys.*, 2008, pp. 329–339.

[13] P. Fonseca, C. Li, V. Singhal, and R. Rodrigues, "A study of the internal and external effects of concurrency bugs," in *Proc. Intl. Conf. Depend. Sys. Netwks.*, 2010, pp. 221–230.

[14] S. K. Sahoo, J. Criswell, and V. Adve, "An empirical study of reported bugs in server software with implications for automated bug diagnosis," in *Proc. Intl. Conf. Softw. Eng.*, 2010, pp. 485–494.

[15] M. Grottke and K. S. Trivedi, "Software faults, software aging and software rejuvenation," *J. Rel. Eng. Ass. Japan*, vol. 27, no. 7, pp. 425–438, 2005.

[16] Y. Huang, C. Kintala, N. Kolettis, and N. Fulton, "Software rejuvenation: Analysis, module and applications," in *Proc. Intl. Symp. Fault-Tolerant Computing*, 1995, pp. 381–390.

[17] A. Pfening, S. Garg, A. Puliafito, M. Telek, and K. S. Trivedi, "Optimal software rejuvenation for tolerating soft failures," *Perf. Eval.*, vol. 27-28, pp. 491–506, 1996.

[18] M. Grottke, A. P. Nikora, and K. S. Trivedi, "An empirical investigation of fault types in space mission system software," in *Proc. Intl. Conf. Depend. Sys. and Netwks.*, 2010, pp. 447–456.

[19] V. Basili and B. Perricone, "Software errors and complexity: an empirical investigation," *Comm. ACM*, vol. 27, no. 1, pp. 42–52, 1984.

[20] N. Fenton and N. Ohlsson, "Quantitative analysis of faults and failures in a complex software system," *IEEE Trans. Softw. Eng.*, vol. 26, no. 8, pp. 797–814, 2000.

[21] T. J. Ostrand, E. J. Weyuker, and R. M. Bell, "Predicting the location and number of faults in large software systems," *IEEE Trans. Softw. Eng.*, vol. 31, no. 4, pp. 340–355, 2005.

[22] A. Chou, J. Yang, B. Chelf, S. Hallem, and D. Engler, "An empirical study of operating systems errors," in *Proc. Symp. Op. Sys. Princ.*, 2001, pp. 73–88.

[23] D. Perry and W. Evangelist, "An empirical study of software interface faults," in *Proc. Conf. Syst. Sci.*, 1985, pp. 113–126.

[24] R. Chillarege, I. S. Bhandari, J. K. Chaar, M. J. Halliday, D. S. Moebus, B. K. Ray, and M.-Y. Wong, "Orthogonal defect classification - a concept for in-process measurements," *IEEE Trans. Softw. Eng.*, vol. 18, no. 11, pp. 943–956, 1992.

[25] S. Bourne, "Interview: A conversation with Bruce Lindsay," *ACM Queue*, vol. 2, no. 8, pp. 22–33, 2004.

[26] M. Sullivan and R. Chillarege, "Software defects and their impact on system availability-a study of field failures in operating systems," in *Proc. Intl. Symp. Fault-Tolerant Computing*, 1991, pp. 2–9.

[27] R. Chillarege, "Understanding Bohr-Mandel bugs through ODC triggers and a case study with empirical estimations of their field proportion," in *Proc. Wksp. Softw. Aging Rejuv.*, 2011.

[28] M. Grottke, R. Matias, and K. Trivedi, "The fundamentals of software aging," in *Proc. Wksp. Softw. Aging Rejuv.*, 2008.

[29] S. Garg, A. Van Moorsel, K. Vaidyanathan, and K. Trivedi, "A methodology for detection and estimation of software aging," in *Proc. Intl. Symp. Softw. Rel. Eng.*, 1998, pp. 283–292.

[30] M. Grottke and K. S. Trivedi, "Fighting bugs: Remove, retry, replicate, and rejuvenate," *IEEE Comp.*, vol. 40, no. 2, pp. 107–109, 2007.

[31] Wikipedia, "LAMP (software bundle)," http://en.wikipedia.org/wiki/LAMP_(software_bundle), 2013.

[32] G. Xu and A. Rountev, "Precise memory leak detection for Java software using container profiling," in *Proc. Intl. Conf. Softw. Eng.*, 2008, pp. 151–160.

[33] R. Love, *Linux Kernel Development*, 3rd ed. Addison-Wesley, 2010.

[34] D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*. CRC Press, 1997.

[35] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *J. Roy. Stat. Soc., Ser. B (Method.)*, pp. 289–300, 1995.

[36] N. Nethercote and J. Seward, "Valgrind: a framework for heavyweight dynamic binary instrumentation," *ACM Sigplan Notices*, vol. 42, no. 6, pp. 89–100, 2007.

[37] M. Cinque, D. Cotroneo, and A. Pecchia, "Event logs for the analysis of software failures: A rule-based approach," *IEEE Trans. Softw. Eng.*, vol. 39, no. 6, pp. 806–821, 2013.