

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

1-2020

Memory and resource leak defects and their repairs in Java projects

Mohammadreza GHANAVATI
Heidelberg University

Diego COSTA
Heidelberg University

Janos SEBOEK
Heidelberg University

David LO
Singapore Management University, davidlo@smu.edu.sg

Artur ANDRZEJAK
Heidelberg University

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Programming Languages and Compilers Commons](#), and the [Software Engineering Commons](#)

Citation

GHANAVATI, Mohammadreza; COSTA, Diego; SEBOEK, Janos; LO, David; and ANDRZEJAK, Artur. Memory and resource leak defects and their repairs in Java projects. (2020). *Empirical Software Engineering*. 25, (1), 678-718.

Available at: https://ink.library.smu.edu.sg/sis_research/4501

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/328015993>

Memory and Resource Leak Defects and their Repairs in Java Projects

Preprint · September 2018

CITATIONS

0

READS

134

5 authors, including:



Mohammad Ghanavati

Universität Heidelberg

5 PUBLICATIONS 59 CITATIONS

SEE PROFILE



Diego Elias Costa

Concordia University Montreal

20 PUBLICATIONS 66 CITATIONS

SEE PROFILE



David Lo

Singapore Management University

381 PUBLICATIONS 7,503 CITATIONS

SEE PROFILE



Artur Andrzejak

Universität Heidelberg

114 PUBLICATIONS 2,090 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Efficient software testing [View project](#)



Distributed Storage Systems [View project](#)

Memory and Resource Leak Defects and their Repairs in Java Projects

Mohammadreza Ghanavati ·
Diego Costa · Janos Seboek ·
David Lo · Artur Andrzejak

Received: date / Accepted: date

Abstract Despite huge software engineering efforts and programming language support, resource and memory leaks are still a troublesome issue, even in memory-managed languages such as Java. Understanding the properties of leak-inducing defects, how the leaks manifest, and how they are repaired is an essential prerequisite for designing better approaches for avoidance, diagnosis, and repair of leak-related bugs.

We conduct a detailed empirical study on 452 issues from 10 large open-source Java projects. The study proposes taxonomies for the leak types, for the defects causing them, and for the repair actions. We investigate, under several aspects, the distributions within each taxonomy and the relationships between them. We find that manual code inspection and manual runtime detection are still the main methods for leak detection. We find that most of the errors manifest on error-free execution paths, and developers repair the leak defects in a shorter time than non-leak defects. We also identify 13 recurring code transformations in the repair patches. Based on our findings, we draw a variety of implications on how developers can avoid, detect, isolate and repair leak-related bugs.

Keywords empirical study · memory leak · resource leak · leak detection · root-cause analysis · repair patch

Mohammadreza Ghanavati
E-mail: ghanavati@uni-heidelberg.de

Diego Costa
E-mail: diego.costa@informatik.uni-heidelberg.de

Janos Seboek
E-mail: janos.seboek@stud.uni-heidelberg.de

David Lo
E-mail: davidlo@smu.edu.sg

Artur Andrzejak
E-mail: artur.andrzejak@informatik.uni-heidelberg.de

1 Introduction

Leaks are unreleased system resources or memory objects which are no longer used by an application. In memory-managed languages such as Java, C#, or Go, a garbage collector handles memory management. Garbage collector uses object reachability to estimate object liveness. It disposes of any heap objects which are no longer reachable by a chain of references from the root objects. However, if an unused object is still reachable from other live objects, garbage collector cannot reclaim the space. Aside from memory, finite *system resources* such as file handles, threads, or database connections require explicit management specified in the code. It is the responsibility of the programmer to dispose of the acquired resource after using it, otherwise, a resource leak is likely.

Leak-related bugs has a high severity (Tan et al. 2014) and can finally result in performance degradation and program crash. Hence, they should be resolved in an early stage of development. However, due to their non-functional characteristics, leaks are likely to escape traditional testing processes and become first visible in a production environment. The root cause of a memory leak can differ from the allocation which exhausts the memory (Jump and McKinley 2007). Some leaks can only be triggered if an abnormal behavior occurs such as an exception or a race condition. These factors make leak diagnosis hard and error-prone.

Defects induced by memory and resource leaks are among the important problems for both researchers and practitioners. Microsoft engineers consider leak detection and localization as one of the top ten most significant challenges for software developers (Lo et al. 2015). This problem is addressed by various researchers, tools, and programming languages. Many previous work targeted memory and resource leak diagnosis by leveraging static and dynamic analysis. Static leak detection techniques include value-flow reachability analysis (Cherem et al. 2007), data-flow analysis (Orlovich and Rugina 2006), object ownership analysis (Rayside and Mendel 2007), loop invariant analysis (Yan et al. 2014), and automated resource management (Dillig et al. 2008; Nguyen et al. 2015; Torlak and Chandra 2010; Weimer and Necula 2004a). The main challenge of static analysis is lack of scalability and high rate of false positives. To mitigate this issue, researchers apply dynamic analysis techniques for leak diagnosis. The major lines of approaches include staleness detection (Hauswirth and Chilimbi 2004; Bond and McKinley 2006; Novark et al. 2009; Bond and McKinley 2009; Jung et al. 2014), growth analysis (Jump and McKinley 2007; Ghanavati and Andrzejak 2015; Sor et al. 2013), analysis of captured state (Mitchell and Sevitsky 2003; Clause and Orso 2010; Xu et al. 2011), and hybrid approaches (Xu and Rountev 2013).

Programming languages provide support for programmers to prevent occurrences of leak-inducing defects. For instance, Java 7 introduces a new

language construct, called **try-with-resources**¹ to dispose of the objects that implement the *autoclosable* interface. Various open-source or proprietary tools (e.g., FindBugs², Infer³) also aim to help programmers to find the potential leaks in the software codebase. For example, FindBugs provides some rules⁴ to warn programmers about potential file descriptor leaks.

Despite the above-mentioned academic work, language enhancements, and tool supports, a number of challenges are still open. The impact of these efforts depends on whether they target prevalent or rare issue types, whether they can handle difficult cases, and whether their assumptions are realistic enough to be applicable in practice. Programming language enhancement such as `try-with-resources` or tool support such as FindBugs help to find only the resource leaks and not memory leaks. Many of the academic work are motivated by anecdotal evidence or by empirical data collected only from small sets of defects. For example, Xu and Rountev (2008) propose a method for detecting memory leaks caused by obsolete references from within object containers but provide only a limited evidence that this is a frequent cause of leak-related bugs in real-world applications. As another example, Leakbot (Mitchell and Sevitsky 2003) introduces multiple sophisticated object filtering methods based on observations derived from only five large Java commercial applications.

A systematic empirical study of a large sample of leak-related defects from real-world applications can help both researchers and practitioners to have a better understanding of the current challenges on leak diagnosis. We believe such a study can be beneficial in following directions:

Benefit 1. A representative study can characterize the current approaches for leak diagnosis used in practice. This can guide researchers to find limitations of leak detection approaches and motivate further improvements. The results would provide a comprehensive basis for design and evaluation of new solutions.

Benefit 2. It helps programmers to avoid mistakes made by the other programmers and shows some of the best practices for leak diagnosis.

Benefit 3. It can be a verification for the assumptions used in previous work. For example, it is interesting to verify empirically whether there is a large amount of leaks caused by collection mismanagement in real-world applications. The positive answer to this could confirm the assumption of Xu and Rountev (2008) on memory leak detection.

To the best of our knowledge, the research body of empirical studies on resource and memory leak-related defects is relatively thin in comparison with the large body of studies about other bug types (e.g., semantic or performance bugs). The existing studies (Machida et al. 2012; Tan et al. 2014) provide only little information about characteristics of detection types, root causes, and

¹ <https://docs.oracle.com/javase/tutorial/essential/exceptions/tryResourceClose.html>

² <http://findbugs.sourceforge.net>

³ <http://www.fbinfer.com>

⁴ <http://findbugs.sourceforge.net/bugDescriptions.html>

repair actions of leak defects. To fill this gap, we conduct a detailed empirical study on 452 real-world memory and resource leak defects gathered from 10 large, open-source Java applications.

We manually study the collected issues and their properties: leak types, detection types, common root causes, repair actions, and complexity of fix patches. Based on our findings, we draw several implications on how to improve avoidance, detection, localization, and repair of leak defects. In particular, this study tries to answer the following research questions:

- . **RQ1.** What is distribution of leak types in studied projects?
- . **RQ2.** How are leak-related defects detected?
- . **RQ3.** To what extent are the leak-inducing defects localized?
- . **RQ4.** What are the most common root causes?
- . **RQ5.** What are the characteristics of the repair patches?
- . **RQ6.** How complex are repairs of the leak-inducing defects?

The preliminary idea of this work is presented in a two-pages short paper in ICSE 2018 (Ghanavati et al. 2018). This work provides the following contributions:

Characterization study. We conduct an empirical study on 452 bugs from 10 mature, large Java applications. To the best of our knowledge, this is the first work which studies characteristics of leak-related bugs from real-world applications in a comprehensive way while using a large set of issues from diverse open-source applications.

Taxonomies. We propose taxonomies for leak types (Section 4.1), detection types and methods (Section 4.2), root causes (Section 4.4), and repair actions (Section 4.5).

Analysis. We investigate the distributions of leaks across the categories within each taxonomy and the relation between the taxonomies. Our findings show that source code analysis and resource monitoring are the main techniques to detect leaks. Our analysis using a state-of-the-art resource leak detection tool (i.e., Infer) highlights that the static analysis tools require further improvement to detect different leak types in practice. We find that 75% of the leaks are triggered during the error-free execution paths. We identify 13 recurring code transformations in the repair patches. We also show that developers resolved the studied issues in about 6 days on median.

Implications. We use our findings to draw a variety of implications on the leak prevention and diagnosis for both researchers and practitioners (Section 5).

Replicability. To make our study replicable and reusable for the community, we make the dataset and the results available online⁵.

This paper is organized as follows. Section 2 provides a short background about leak definition and issues in the bug tracking systems. Section 3 describes the design of our empirical study. In Section 4, we present the answers to the research questions. In Section 5, we present the implications drawn from our

⁵ https://github.com/heiqs/leak_study

observations and findings. Section 6 discusses potential threats to the validity of our study. Section 7 surveys related work. Finally, Section 8 concludes the paper.

2 Background

2.1 Leak definition

Leaks occur due to mismanagement of memory or finite systems resources. In this section, we briefly explain these two types.

Memory leak. Contrary to the unmanaged languages such as C or C++ in which programmer is responsible for freeing the memory, in memory-managed languages such as Java or C#, a garbage collector reclaims the space. A programmer can rely on the garbage collector to release references due to dangling pointers or lost pointers. However, if the references to the unused objects are present in the running process, they cannot be garbage-collected. As a sequence, a memory leak might be triggered. In other words, a *memory leak* in Java occurs when process maintains unnecessary references to some unused objects.

Resource leak. In Java, finite system resources like connections, threads, or file handles are wrapped in special *handle objects*. Programmer accesses such a resource by normal object allocation. However, in contrast to memory management, the developer should dispose of a system resource by making an explicit call to the disposal method of the handle object (or by ensuring that a thread has stopped). Besides this, all unnecessary references to such objects should be removed to prevent the potential memory leak. Hence, a *resource leak* occurs when the programmer forgets to call the respective close method for a finished handle object. Similar to memory leak, a resource leak gradually depletes system resources which degrades performance and can lead to a failure.

In this paper, we use the term *leak* for both memory and resource leaks. We also occasionally use the term *disposing of an object* for either closing a resource or releasing (deallocating) memory (in Java, by removing all references to an object).

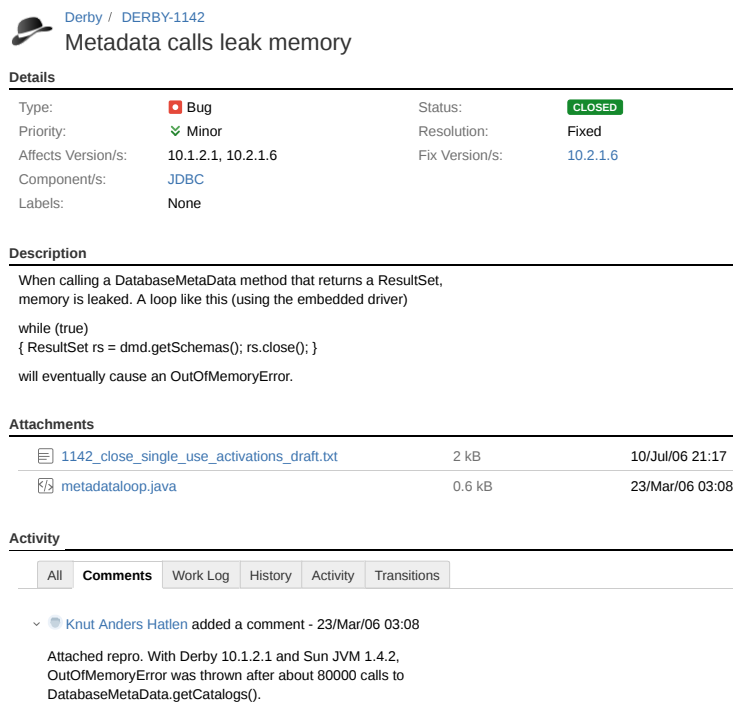
2.2 Issue Report

Modern projects often use an Issue Tracking System (ITS) to collect the issues reported by users, developers, or software quality teams. An issue typically corresponds to a bug report or a feature request. Bugzilla⁶, Jira⁷, and GitHub issue tracker⁸ are examples of ITS systems. Jira is one of ITS used by Apache

⁶ <https://www.bugzilla.org/>

⁷ <https://issues.apache.org/jira/projects/>

⁸ <https://github.com/>



Derby / DERBY-1142
Metadata calls leak memory

Details

Type:	❌ Bug	Status:	CLOSED
Priority:	✔️ Minor	Resolution:	Fixed
Affects Version/s:	10.1.2.1, 10.2.1.6	Fix Version/s:	10.2.1.6
Component/s:	JDBC		
Labels:	None		

Description

When calling a DatabaseMetaData method that returns a ResultSet, memory is leaked. A loop like this (using the embedded driver)

```
while (true)
{ ResultSet rs = dmd.getSchemas(); rs.close(); }
```

will eventually cause an OutOfMemoryError.

Attachments

1142_close_single_use_activations_draft.txt	2 kB	10/Jul/06 21:17
metadataloop.java	0.6 kB	23/Mar/06 03:08

Activity

All **Comments** Work Log History Activity Transitions

▼ Knut Anders Hatlen added a comment - 23/Mar/06 03:08

Attached repro. With Derby 10.1.2.1 and Sun JVM 1.4.2, OutOfMemoryError was thrown after about 80000 calls to DatabaseMetaData.getCatalogs().

Fig. 1 An issue report from JIRA.

repository. As we study the projects hosted in Apache repository⁹, we build our dataset based on the issue reports filed in JIRA. Each issue report in JIRA is identified with a unique identifier. An issue report contains a variety of information such as title, description, comments, and related fix patches. It also contains metadata information such as type, status, priority, resolution, and associated timestamps (e.g., created or resolved timestamps). Figure 1 shows a snippet of an issue report. All the information provided in issue reports makes the issue tracker a rich environment to get more insights on bugs and their corresponding repairs.

3 Empirical Study Design

In this section, we describe the design of our empirical study. Figure 2 gives an overview of our methodology. In the remainder of this section, we illustrate the research questions, studied applications, and data collection process.

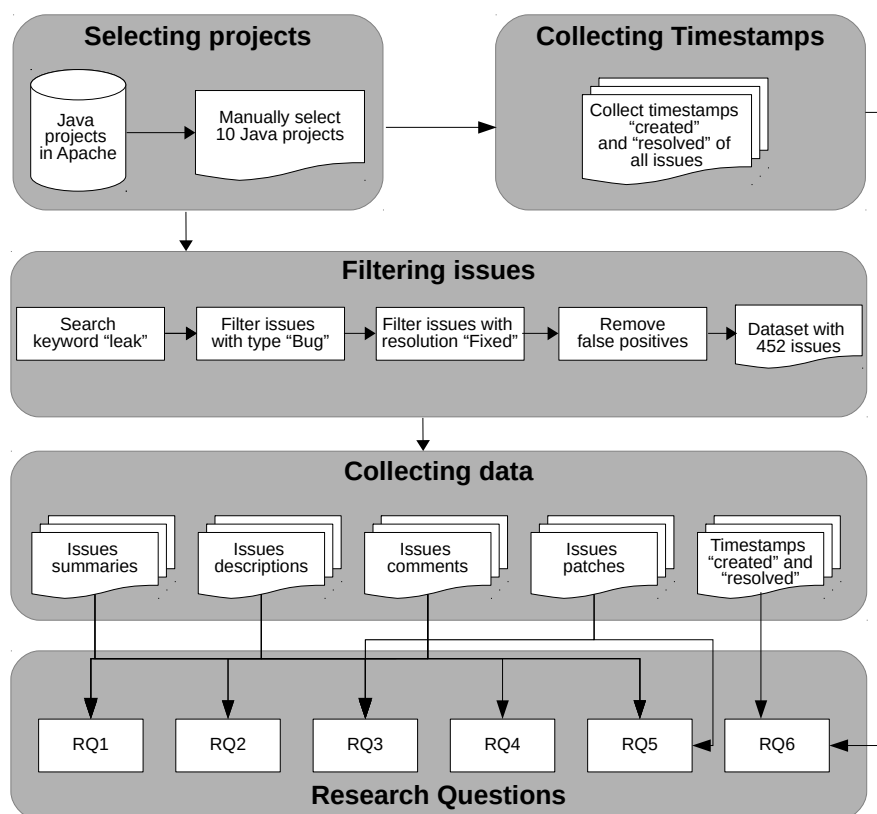


Fig. 2 Overview of the empirical study design.

3.1 Studied Projects

We perform a study on ten open-source Java projects. We investigate the leak-related issues from a wide variety of software categories to ensure the diversity of the studied projects.

Table 1 lists the studied projects. AMQ¹⁰ is an open-source message broker with the support of cross language clients and protocols. CASSANDRA¹¹ is a distributed database targeting high scalability and availability. CXF¹² is an open source framework for developing services using front-end programming APIs. DERBY¹³ is an open-source relational

⁹ Note that all the projects in our study have a mirror project in GitHub

¹⁰ <http://activemq.apache.org>

¹¹ <http://cassandra.apache.org>

¹² <http://cxf.apache.org>

¹³ <http://db.apache.org/derby>

Table 1 Overview of studied projects. The Java LOC for each project is obtained from Open Hub.

Project	Category	First Commit	#Committers	#kLOC
AMQ	Distributed messaging	2004	58	1158
CASSANDRA	Distributed database	2009	45	313
CXF	Web service	2007	38	674
DERBY	Relational database	2004	44	689
HADOOP	Distributed computing	2006	163	1260
HBASE	Distributed database	2007	57	1115
HIVE	Data warehouse	2009	63	1074
HTTPCOMP.	Network client/server	2004	18	115
LUCENE	Search framework	2004	67	557
SOLR	Search framework	2008	67	416

database. HADOOP¹⁴ is a distributed computing platform including four main components: HADOOP Common, HDFS, MapReduce, and YARN. HBASE¹⁵ is a distributed, scalable and big data store. HIVE¹⁶ is an SQL-enabled data warehouse for large datasets. HTTPCOMPONENT¹⁷ with its two components core and client is a tool set for working with the HTTP protocol. LUCENE¹⁸ is a high performance, cross-platform text search engine. SOLR¹⁹ is an open-source full-text enterprise search server based on LUCENE.

We study these projects for two reasons. First, they are large-scale and open-source projects with a mature codebase with years of development. We believe that by using such a well-established and well-developed applications, we can get results representative for mature Java projects. Column "#kLOC" in Table 1 shows the size of the Java source code of the studied projects ranging between 115 to over 1200 kLOC. Second, their issues are reported and tracked in a bug tracking system, called JIRA. Similar to other bug trackers (e.g., Bugzilla), reports in JIRA are well-described and provide sufficient information to answer the research questions investigated in this study.

3.2 Research Questions

The following research questions guide our study:

RQ1. What is distribution of leak types in studied projects?

In Section 4.1, we analyze the dominant leak types in each project. We use this analysis in next research questions to distinguish the properties of different leak types.

¹⁴ <http://hadoop.apache.org>

¹⁵ <http://hbase.apache.org>

¹⁶ <http://hive.apache.org>

¹⁷ <http://hc.apache.org>

¹⁸ <http://lucene.apache.org/core>

¹⁹ <http://lucene.apache.org/solr>

RQ2. How are leak-related defects detected? Understanding different detection types can help leak detection approaches to improve the detection accuracy. In Section 4.2, we investigate how developers or users report the leak-inducing defects and how the leaks manifest at runtime. We analyze different detection and manifestation types and study their relation to the leak types.

RQ3. To what extent are the leak-inducing defects localized? Bug localization is the first step in bug diagnosis. The extent of the bug can highly affect the number of files that need to be fixed to repair the bug. In this question, we analyze the locality of leak-inducing defects (Section 4.3).

RQ4. What are the most common root causes? Section 4.4 describes the common root causes of leak defects. We investigate the prevalence of each root cause and their relation to the leak types.

RQ5. What are the characteristics of the repair patches? In Section 4.5, we identify the repair actions applied by the developers to repair the leak-related defects and investigate the frequency of each considering different leak types. We also search to find recurring code transformations in the repair patches. We identify 13 common repair patterns from our dataset. In this question, we investigate whether the automated program repair techniques (i.e., the process of providing the repair patches automatically) such as template-driven patch generation are applicable for fixing the leak-related defects.

RQ6. How complex are repairs of the leak-inducing defects? In Section 4.6, we measure the code churn, change entropy, and diagnosis time to assess the complexity of the changes needed to repair the leak-inducing defects. This analysis provides insights about the difficulty of repairing the leak-related defects and shows which type of leaks can be repaired with less effort in terms of time and amount of code changes.

3.3 Data Extraction

We collected the leak-related issues from the issue tracker in June 2016. The issues were reported between January 2004 and June 2016.

To build a suitable dataset for our study, we apply a four-step filtering methodology: (1) keyword search, (2) issue type filtering, (3) resolution filtering, and (4) manual investigation. This four-step filtering method yields a dataset with 452 leak-related issues, each representing a unique leak bug report (i.e., none are duplicates of another). We make the dataset available online²⁰.

Keyword search. We use a simple heuristic and select issues that contain the keyword "leak" in the issue title or issue description. The keyword search is a well-known method used by previous empirical studies (Jin et al. 2012a; Zhong and Su 2015; Nistor et al. 2013) to filter the issues of interest from the others. Note that other related keywords might lead to many false positives

²⁰ https://github.com/heiqs/leak_study

Table 2 Studied projects with statistics on number of issues (explained in Section 3.3). Columns “#MLeak”, “#RLeak” and “Total” show the numbers of memory and resource leak issues per application, and their totals, respectively.

Project	#Issues	#Bugs	#Fixed	#MLeak	#RLeak	Total
AMQ	123	116	88	54	26	80
CASSANDRA	77	65	45	19	16	35
CXF	62	61	44	29	8	37
DERBY	50	36	23	12	4	16
HADOOP	236	201	132	43	76	119
HBASE	92	65	44	11	29	40
HIVE	78	69	47	19	25	44
HTTPCOMP.	31	28	24	8	12	20
LUCENE	77	65	42	13	21	34
SOLR	74	60	33	11	16	27
Total	900	766	522	219	233	452

causing high manual efforts to prune non-relevant issues. Despite the simplicity of keyword search, this heuristic proved to be highly precise due to the high quality of issue reports and related data in the studied projects. Wu et al. (2011) highlight that even simple heuristics can yield the same precision and recall as more sophisticated search techniques when applied to a well-maintained bug tracker. Using the keyword search, we identify 900 leak-related issues. Column “#Issues” in Table 2 shows the number of filtered issues for each project.

Issue type filtering. Each issue in the bug tracker can be classified as “Bug”, “Task”, “Test”, and so on. As we are only interested in leak-related bugs, we first filter issues with type “Bug”. Among the 900 issues filtered by keyword search, there are 766 issues labeled as a bug (column “#Bugs” in Table 2).

Issue resolution filtering. To analyze how developers repair a leak defect we need to restrict our analysis to fixed bugs. For this, we filter issues with the resolution label “Fixed”. This reduces the dataset to 522 issues (column “#Fixed” in Table 2).

Manual investigation. In the final step, we remove the false positives from our dataset. We manually filter out the following issues:

- Non-leak-related bugs retrieved by our keyword search heuristic. For instance, in issue CXF-3390²¹, the term *leak* is used in “information leak” which is not related to this study.
- Wrongly reported leaks. These issues should be marked as “Invalid”, but are closed as “Resolved” in the bug tracker.

3.4 Tagging Leak-Related Defects

To analyze the properties of the leak-related defects, we need to classify the issues for each dimension of interest (i.e., leak type, detection type, detection

²¹ <https://issues.apache.org/jira/browse/CXF-3390>

Table 3 Cohen’s kappa measurement.

Dimension	Cohen’s Kappa
Leak Type (RQ1)	0.86
Detection Type (RQ2)	0.83
Detection Method (RQ3)	0.70
Defect Type (RQ4)	0.68
Repair Type (RQ5)	0.56

method, defect type, and repair type). However, we only have qualitative information such as issue description, developers discussions, and repair patches. There is no label provided in the bug tracker for classification of the attributes that we are interested for reported leaks. To derive properties for the bugs in our dataset, we need to quantify the qualitative information. For this purpose, we perform an iterative process similar to *Open Coding* (Seaman 1999; Seaman et al. 2008). In our study, the input of the coding process for each issue are issue summary, issue description, developers discussions, and repair patches. The first author of the paper (a Ph.D. student), classified a sample set of the issues to determine the possible categories for each dimension. After identifying the initial types for each category, the second and the third authors (a Ph.D. student and a undergraduate student) join the first author to discuss about the categories and label the remaining issues. We held many meetings, spent many hours, and performed multiple iterations to achieve a cohesive labeling.

The tagging process is iterative. Each time a new type is identified, the coders (first three authors) verify it in a decision-making meeting. If a majority of the coders agree on the new type, they go through all the previously tagged issues and check if the issues should be tagged with the new type. This also minimize the threat of human error during labeling process. To further reduce the error probability and in case of difficulty in classifying of the issues, all the coders check and discuss the complex issues to find the appropriate categories. The conflicts were resolved by discussing and coming to an agreement.

To validate the manual labeling process, we apply the following procedure. The first and second author perform a classification of a statistically representative sample of the dataset with a confidence level of 95% and a confidence interval of 10%. This results in a sample set of 79 out of 452 issues. We calculate the inter-rate agreement with Cohen’s kappa metric (Cohen 1960; Artstein and Poesio 2008). Table 3 shows the result of our analysis. The lowest Cohen’s kappa value is for the repair type, although it shows a moderate agreement between the two coders. The reason for disagreements is that the categories in this attribute are not mutually exclusive. Therefore, there is a probability that each coder has a different interpretation for the same issue. After rating, the two raters discussed their disagreements to reach consensus.




Fig. 3 Frequency of the leak types per project.

3.5 Uniqueness of categories

During tagging task, we encounter the issues with the possibility of assigning them to multiple categories. For example, in Hadoop 6833²², a leak is reported due to the forgotten call to the *remove* method of a collection. The developers repaired the bug by adding the *remove* call in the exception path:

```

--- src/java/org/apache/hadoop/ipc/Client.java
+++ src/java/org/apache/hadoop/ipc/Client.java
@@ -697,6 +697,7 @@ public class Client {
    } else if (state == Status.ERROR.state) {
        call .setException(new RemoteException(WritableUtils.readString(in),
                                             WritableUtils.readString(in)));
+   calls .remove(id);
    } else if (state == Status.FATAL.state) {
        // Close the connection
        markClosed(new RemoteException(WritableUtils.readString(in),

```

One could label this issue as *collection mismanagement*. However, if the exception is not thrown no leak is triggered. Therefore, we use the underlying cause as the main root-cause category (here *bad exception handling*). For the repair action, we assign a bug to the category used by the developer to repair the defect. In this example, we label the repair action as *remove element*.

4 Empirical Study Results

In this section, we answer the research questions. For each research question, we describe the motivation behind the question, the approach used in answering the research question, and the findings derived from the analysis.

4.1 RQ1: What is distribution of leak types in studied projects?

Motivation. In this research question, we want to find the primary leaked resource for each issue. The leak type classification will be used in further research questions to determine the existence of different patterns for different

²² <https://issues.apache.org/jira/browse/HADOOP-6833>

leak types. We also use this investigation to assess the leak diversity on the studied projects.

Approach. For most of the studied issues, the reporters or developers explicitly mentioned the leak type. For such cases, we assign the leak type as reported. In case of no explicit mention of the leak type, we manually analyze the title, description, and developers discussions to assign the correct leak type.

Taxonomy of leak types. Our analysis yields a taxonomy of leak types with following four categories:

Memory. We group in this category all issues reported due to unreleased references to Java objects, such as mismanagement of collections or circular references.

File handle. We group in this category leaks related to file descriptors. These issues are related to the mismanagement of Java file handlers such as I/O streams.

Connection. We group in this category leaks triggered by non-closed network or database connections.

Thread. We group in this category leaks caused by unclosed, yet unused threads. A thread leak occurs when a no-longer-needed thread is unnecessarily kept alive. This thread then leaks its internal resources, which cannot be released by the JVM.

Results. Figure 3 shows the distribution of the leak types for each project. We use this data to find the dominant leak types in the projects and in the project categories.

Finding 1. The three leak types corresponding to the resource leaks (i.e., file handle, connection, and thread) are the most common leak types in six out of the ten projects. Resource leaks (with 233 issues) are slightly more reported than memory leaks (with 219 issues).

Finding 2. Each project shows a distinct distribution of the leak types. LUCENE and HADOOP have a higher frequency of the *file handle* leak type with this leak type corresponding to 55.9% and 42.9% of the issues, respectively. Projects AMQ (67.5%), CASSANDRA (54.3%), CXF (78.4%), and DERBY (75.5%) contain more memory leak issues. Connection leaks are more frequently reported in HBASE (37.5%), HTTPCOMP (30%), and Hive (27.3%). This analysis shows the diversity of the leak types in the studied projects. Even projects within the same category show different distributions of the leak types.

Summary. Resource leaks (233 out of 452 issues) are slightly more often reported than memory leaks (219 issues). Leak type distribution is different across the projects.

4.2 RQ2: How are leak-related defects detected?

Motivation. Each issue report provides information about leak symptoms, environmental setup, and methods used to detect a leak. Understanding how leaks are detected can provide valuable insights on leak diagnosis. It also shows in which direction the researchers and tool builders should help programmers in leak detection. In this question, we want to find whether the leaks are detected during runtime and whether the static analysis is used for leak detection.

Approach. To find detection type for each issue, we use three data sources: issue title, issue description, and developers discussions. Using this data, we analyze the frequency of the detection types, distribution of detection methods, and their relation to different leak types.

Taxonomy of leak detection. Leak-inducing defects can be discovered with and without runtime failures or performance degradation. They can be detected via manual analysis of the source code, an unexpected runtime failure (in particular, an out-of-memory error), or abnormal usage of resources. We classify detection types into two major categories: source code-based detection and runtime detection. In the following, we explain these two detection types in more detail.

Source code-based detection. In this category, we classify issues such that the leak detection is performed before execution of the program and there is no reported runtime information in the issue report, nor reports on leak-related failures. We observe that issue reporters describe these issues with phrases such as "can potentially cause a leak" or "can lead to a leak". The main techniques to detect leaks prior to the runtime are *manual code inspection* and *static analysis tools*.

Manual inspection of the source code (or code review) is a process in which developers inspect a set of program elements (e.g., methods, classes) in order to improve the quality of software (Huizinga and Kolawa 2007; McConnell 1993; Sommerville 2010). It is one of the most common static detection methods used by developers in practice. This detection type requires the knowledge of how a leak can be introduced as well as understanding of the application's behavior. For instance, in AMQ-5745²³, manual inspection revealed cases where bad exception handling could yield resource leaks on the AMQ codebase.

Static analysis tools can be used to identify potential leak defects during the development process. There are many free and proprietary static analyzers which are able to detect specific leak types (e.g., FindBugs, Infer).

Runtime analysis. Some leak-related failures are observed and reported when a user/developer encounters a performance degradation in a production environment, an out-of-memory error is raised, or a test is failed. Issue reporters often use phrases such as "consistently observing memory growth" or "meet memory leak in a production environment". In these issues, the bug reporter often provides additional material such as heap profile, memory dump, a log file, or a stack trace. This supplementary data can help

²³ <https://issues.apache.org/jira/browse/AMQ-5745>




Fig. 4 Frequency of the detection types per leak type.

developers on localizing the root cause of the leak defect more efficiently. Leaks usually manifest in the runtime with a symptom. From our observation, we identify three symptoms reported in the issue reports: failing tests, out-of-memory errors, and warning messages.

The output of a failing test case may expose a leak. The test can be a system test, a unit test or any other application-provided test. For example, in LUCENE-3251²⁴, a failing unit test exposed a file-handle leak. The user provided the stack trace of the failing test in the issue report:

```

Testsuite : org.apache.lucene.index.TestAddIndexes
Testcase : testAddIndexesWithRollback(org.apache.lucene.index.TestAddIndexes):
Caused an ERROR
MockDirectoryWrapper: cannot close: there are still open files : {_co.cfs=1}
java.lang.RuntimeException: MockDirectoryWrapper: cannot close: there are still open files : {_co.cfs=1}
    at org.apache.lucene.store.MockDirectoryWrapper.close(MockDirectoryWrapper.java:483)

```

In some cases, the growth of memory usage leads to an out-of-memory (OOM) error during runtime. This is a severe symptom as the underlying system often crashes when an OOM error occurs. For example, DERBY-5730²⁵ reported a severe memory leak which might lead to a system crash due to an out-of-memory error. In DERBY-5730, the reporter mentioned that after removing the suspicious call, the test program was successfully executed with a much lower heap size.

Developers also implement algorithms for detection of specific leak defects. They usually warn the user about the potential presence of a leak with a message during program's execution. For example, in CXF-5707²⁶, a message warned the user for a potential leak during the performance test of the *netty-http-server* module:




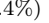








”SEVERE: LEAK: ByteBuf.release() was not called before it's garbage-collected. Enable advanced leak reporting to find out where the leak occurred.”

²⁴ <https://issues.apache.org/jira/browse/LUCENE-3251>

²⁵ <https://issues.apache.org/jira/browse/DERBY-5730>

²⁶ <https://issues.apache.org/jira/browse/CXF-5707>

Table 4 Distribution of detection methods for memory and resource leaks.

Detection Type	Detection Method	Memory Leaks	Resource Leaks
Source code-based detection	Manual code inspection	 64 (29.2%)	 106 (45.5%)
	Static analyzer	 0 (0.0%)	 1 (0.4%)
Runtime detection	Failed test	 16 (7.3%)	 38 (16.3%)
	Out-of-memory error	 38 (17.4%)	 12 (5.2%)
	Warning message	 7 (3.2%)	 7 (3.0%)
	Runtime (exclude above)	 94 (42.9%)	 69 (29.6%)

Results. Figure 4 shows the distribution of detection types in relation to the leak types. Table 4 illustrates the relationship between detection types, detection methods, and leak types.

Finding 1. More resource leaks (106 issues) are detected via source code-based detection than memory leaks (64 issues). Runtime detection is the dominant detection type for detecting memory leaks with 155 out of 219 issues (70.8% of the issues). The reason why more resource leaks are detected with source code-based detection techniques comes from the difference in memory and resource management. In Java, a programmer should explicitly dispose of the resources after usage. Due to explicit management, potential resource leaks can more often be captured through the code review or using static analyzers. Contrary to this, the JAVA Virtual Machine (JVM) manages the heap space and releases the unused objects when they become unreachable. Detecting unused references with code inspection is a hard task, as the programmer needs to have a profound understanding of the program’s workflow.

Finding 2. 281 (about 62.2%) issues are detected or manifest during the runtime. In these issues, users often use a third-party memory analyzer (e.g., *jmap*, *MAT*²⁷, *yourkit*²⁸), or OS-specific commands (e.g., *lsdf*) to verify the presence of the leaks. The information collected from these tools and commands can considerably help the developers to reproduce and diagnose the leak defects.

Finding 3. Users detected leaks in 14 issues (3.1%) via warning messages. From our dataset, we observe that in three applications, developers implemented leak detection mechanisms. This result shows that it is a good practice for developers to provide integrated leak detection mechanisms to accelerate diagnosis of the leak-related defects.

Finding 4. Out-of-memory errors are observed more than three times in memory leaks-related issues. OOM error is one of the most severe leak symptom and should be particularly prevented in a production environment.

Finding 5. In 54 (about 12%) issues, users detected the leaks via a test case. We also observe that for some issues, developers added a test case to

²⁷ <https://www.eclipse.org/mat/>

²⁸ <https://www.yourkit.com/>

the repair patches for future leak detection. This result shows the possibility of the software tests as a lightweight tool for leak detection. Previous work (Fang et al. 2015; Ghanavati and Andrzejak 2015) confirm the effectiveness of software tests for leak detection. The utility of a test case is twofold. First, it can be used as an oracle for automated leak detection and bug isolation. Second, it can be an oracle for automated leak repair techniques as they need test cases to verify the correctness of their proposed fix patches. As leaks are highly environment - and input - sensitive, the automated test input generation should provide inputs which can trigger the leaks in different execution paths.

Finding 6. Only in one issue (CASSANDRA-7709²⁹), the leak is detected and reported by a static analysis tool. As we only consider the reported issues, we cannot generalize that the static analysis tools are not used. However, it is important to know why these tools are not used for other reported issues with the similar characteristics as the detected issue. One reason might be that there are still obstacles in fully usage of such debugging tools. Such obstacles can be high false positives, complex usage procedure, or lack of knowledge about these tools. Researchers or tool builders should improve current debugging tools to detect not-yet covered bugs, simplify the tool usage, and spread them widely in the community.

Summary. Source code-based detection is more common in resource leak detection (45.9%). Runtime detection is the dominant detection type for memory leaks (70.8%). Out-of-memory errors are observed about three times more frequently in conjunction with the memory leaks than with the resource leaks.

4.3 RQ3: To what extent are the leak-inducing defects localized?

Motivation. Fault localization is the first step of bug diagnosis. Locality of a fault can be defined in different granularity such as statement, method, and file. In case of leak-related defects, they can affect a region (e.g., multiple modules, classes, etc.) in the codebase of an application (Mitchell and Sevitsky 2003). Accurate defect localization is vital in providing the correct repair patch. Otherwise, the patch will not fix the bug completely and even introduces a new bug (Yin et al. 2011). In this research question, we investigate the locality of leak-inducing defects. In particular, we want to find out: (1) how many source code files are changed to repair a defect, and (2) which types of files are changed in each repair patch.

Approach. To assess the locality of leak defects, we analyze the distribution of modified source code files. For each issue, we collect the files changed in the repair patches. We also investigate the file type of modified source files by collecting the file extensions. We ignore test files in the repair patches if

²⁹ <https://issues.apache.org/jira/browse/CASSANDRA-7709>




Fig. 5 Heatmap of the number of modified Java source code files per project.

the tests added or modified for future leak detection and not for repairing purposes.

Results. Figure 5 shows the heatmap representation of the amount of modified Java source code files for each project.

Finding 1. In 57% of the issues, developers changed only one source code file to repair the defect. In about 81% of the issues, three source code files are modified at most. This result implies the high locality of leak-inducing defects considering file level granularity.

Finding 2. In 14 issues, developers repaired the defect via adding at least one Java source code file without deleting any code file. Albeit occurring on rare situations, such cases require more sophisticated repair strategies. Most of the current automated program repair approaches (Weimer et al. 2009; Kim et al. 2013; Le et al. 2016) can only provide correct repair patches with only one source code line. Hence, it is still not feasible for existing automated program repair techniques to provide complex repair patches such as adding a complete method or class.

Finding 3. In 12 issues, no Java source code file is changed. In six issues, source code files written in C are modified. In three cases, developers modified the XML files to use a non-leaky third-party library as a dependency for that project. Three issues are repaired by changing source code files written in Scala and Ruby. The reason for changing different file types is that in some studied projects, specific modules are implemented in different programming languages than Java. For example, bzip2 (a compression method) implementation in HADOOP is written in C. We also observe that in 15 issues only test files are changed. It is because these leaks are introduced in the test suite and not in the source code of the applications. We also observe that in 24 issues, no repair patch is provided as the bug was already fixed in previous versions of the applications.

Although only few defects are repaired via modifying files written in other programming languages, there is still a need of having skills and knowledge on different languages to repair all are required to repair the leak defects in some specific scenarios.

Summary. 57% of the leak defects are repaired via changing only one source code file. Only in one-fifth of the reported leaks, more than three source files were modified. In 6% of the issues, files from other languages, such as C, Scala, and Ruby are modified to fix the leak-related defect.

4.4 RQ4: What are the most common root causes?

Motivation. In this research question, we want to find out which root causes are dominant, and whether there are significant differences in the common root causes for different leak types.

Approach. To find the root cause, we use three data sources for each issue: issue title, issue description, and developers discussions. The categories for root cause are not mutually exclusive. For issues with the possibility of assignment to multiple categories, we select the most specific category as explained in Section 3.5.

Taxonomy of defect types. Table 5 lists the taxonomy of the defect types. We describe the most common root causes here.

Non-closed resource. The programmer should close any system resources such as file handles, connections, and threads when they are no longer needed. Otherwise, a resource leak is likely. For example, in HBASE-12837³⁰, zookeeper connections created in the constructor of `ReplicationAdmin` left unclosed.

Bad exception handling. According to Java documentation³¹, an exception is an event which disrupts the normal flow of the program's instructions. When an exception is thrown, any resources accessed during the normal execution of the program remain open. If a programmer does not properly handle the exceptions, a leak is prone to happen, as shown in the following quote from an issue from Lucene:

“Programmer should handle the exception properly instead of swallowing it.”

For instance, in LUCENE-3144³², `FreqProxTermsWriter` leaks open file handle if an exception is thrown during `flush()`.

Collection mismanagement. The mismanagement of elements in a collection can result in memory leak. Such leak occurs when a programmer assumes that garbage collector collects all unused objects, even if they are still referenced. Leaks due to collection mismanagement can lead to severe memory waste, in particular when the collection is used as a static member. The reason is

³⁰ <https://issues.apache.org/jira/browse/HBASE-12837>

³¹ <https://docs.oracle.com/javase/tutorial/essential/exceptions/definition.html>

³² <https://issues.apache.org/jira/browse/LUCENE-3144>

Table 5 Taxonomy of root causes. Column "#Issues" states the total number of issues per root cause.

Description (Short)	#Issues
Non-closed resource at error-free execution (nonClosedRes)	135 (29.87%)
Object not disposed of if exception is thrown (exception)	93 (20.58%)
Dead objects referenced by a collection (collection)	89 (19.69%)
Unreleased reference at error-free execution (unreleasedRef)	54 (11.95%)
A race condition defect (concurrency)	17 (3.76%)
Wrong call schedule of disposal method (callSchedule)	15 (3.32%)
Over-sized cache or buffer (cache)	12 (2.65%)
Incorrect API usage (wrongAPI)	10 (2.21%)
Unreleased reference due to thread-local variable (threadLocal)	9 (1.99%)
ClassLoader keeps a bi-directional reference to a class (classloader)	8 (1.77%)
Leaks related to Java native interface (jni)	8 (1.77%)
Leak inside a third-party library (leakyLib)	2 (0.44%)

Table 6 Detailed frequency of root causes in relationship with leak type. Column "Conn." states numbers for connection leaks.

Defect	Memory	Resources			
		Total	File	Conn.	Thread
<i>nonClosedRes</i>	-	57.9%	53.2%	54%	71.2%
<i>exception</i>	7.3%	33%	40.3%	36%	15.3%
<i>collection</i>	40.6%	-	-	-	-
<i>unrelRef</i>	24.7%	-	-	-	-
<i>concurrency</i>	5.5%	2.1%	0.8%	6%	1.7%
<i>schedule</i>	3.2%	3.4%	3.2%	2%	5.1%
<i>cache</i>	5.5%	-	-	-	-
<i>api</i>	3.7%	0.9%	0.8%	-	1.7%
total	219	233	124	50	59

that the static fields are never garbage-collected. Issue YARN-5353³³ reports a severe memory leak due to keeping the tokens in the `appToken` map of the `ResourceManager` even after task completion.

Concurrency defect. A leak can be caused by a race condition preventing the disposal of a resource or releasing references to unused objects. Issue LUCENE-6499³⁴ reports a file handle leak if files are concurrently opened and deleted.

Results. We investigate the frequency of the root causes across the leak types. Table 5 and Table 6 summarize the results. Table 5 lists the common root causes and their corresponding number of issues. Table 6 shows a more detailed information on the relative frequencies of the most prevalent root causes grouped by the leak types.

³³ <https://issues.apache.org/jira/browse/YARN-5353>

³⁴ <https://issues.apache.org/jira/browse/LUCENE-6499>

Finding 1. The majority of the defects (about 75% of the cases) manifest when a normal execution path is exercised. The most common root cause is also the non-closed resource in a regular (error-free) execution path (*nonClosedRes*) with about 30% of the cases. This finding is interesting. The error-free execution paths are more often executed and checked. Therefore, it should be less likely for defects to manifest in normal execution paths (Weimer and Necula 2005). However, our analysis shows that this is not the case for leak-related defects. Our analysis shows the value of software tools and tests which check whether resources are disposed of at the end of the normal execution paths.

Finding 2. Bad exception handling (*exception*) is the second-most frequent root cause with 20.58% of the issues (93 issues). This even increases to 33% of the issues if we only consider resource leaks. We also observe that this root cause is about 5 times more common for resource leaks than for memory leaks. One reason for such observation is that - by definition - exception paths happen in exceptional situations, being less frequently tested than normal execution paths. Even correctly-behaved programs in normal execution paths, can manifest error in exceptional paths (Weimer and Necula 2005, 2004a). This observation implies that the proper exception handling plays an important role for preventing leaks especially resource leaks.

Finding 3. Collection mismanagement (*collection*) is the most common root cause for memory leaks (40.6% of the cases). This finding verifies the applicability of existing automated approaches for detecting memory leaks caused by collection mismanagement (e.g., Xu and Rountev (2008)).

Summary. Most leaks are caused by four root causes. Collection mismanagement (40.6% of the issues) and non-closed resources (58% of the issues) are the dominant root causes for memory and resource leaks, respectively. The majority of the leaks (75% of the cases) manifest in the regular execution paths.

4.5 RQ5: What are the characteristics of the repair patches?

Motivation. One approach for automated program repair is to search for common repairs from previous fix patches and provide repair candidates to fix bugs (Kim et al. 2013; Le et al. 2016; Liu et al. 2013; Selakovic and Pradel 2016; Song and Lu 2014). Align with this direction, we investigate the repair actions and code transformations in the repair patches to check whether there are common patterns for fixing the leak-inducing defects.

Approach. For each issue in our dataset, we manually check the issue summary, the issue description, the developer discussions, and the repair patches to understand and find the repair action for each defect. When analyzing the patches, we apply the following considerations. First, we are only interested in the defects within the codebase of the application. Hence,

Table 7 Taxonomy of repair actions. Column "#Issues" states the total number of issues per repair action.

Description(Short)	#Issues
R1: Dispose of resource in regular execution paths (disposeReg)	104 (23.01%)
R2: Dispose of objects in exceptional path (disposeExcep)	92 (20.35%)
R3: Remove the elements from a collection (removeElm)	100 (22.12%)
R4: Release the reference (releaseRef)	67 (14.82%)
R5: Shutdown thread after finishing the task (threadDown)	39 (8.63%)
R6: Improve thread safety by avoiding race condition (threadSafe)	22 (4.87%)
R7: Use an efficient API to improve memory usage (correctAPI)	10 (2.21%)
R8: Modify strong reference to a weak reference (weakRef)	8 (1.77%)
R9: Use a non-leaky Library (nonLeakyLib)	2 (0.44%)
R10: Bugs not belonging to the above categories (others)	8 (1.77%)

we ignore modifications of the test files in the repair patches. Second, in 24 issues, the defects are already repaired by developers in other issues or another version of the application, but were not tagged as "duplicate" in the bug tracker. We decide to ignore these issues. Every label is attributed to a specific repair action whenever possible. We categorize the fix patch to a generic category only when no specific repair action would fit the repair description.

Taxonomy of repair actions. Table 7 lists the repair actions. We describe the prevalent actions here.

Dispose of resource in a regular path (disposeReg). Resource leak defects introduced in *regular* execution paths can be resolved via simply calling the *dispose* method after the resource usage. In Java, this is achieved by calling the *close* *dispose* method. For example, in HADOOP-7090³⁵, the developer refers to closing the I/O streams in a *finally* block as a *good practice*. Following is a partial patch for this issue:

```

--- org/apache/hadoop/io/BloomMapFile.java
+++ org/apache/hadoop/io/BloomMapFile.java
@@ -186,10 +186,17 @@
    @Override
    public synchronized void close() throws IOException {
        super.close();
        DataOutputStream out = fs.create(new Path(dir, BLOOM_FILE_NAME), true);
        bloomFilter.write(out);
        out.flush();
        out.close();
+       DataOutputStream out = null;
+       try {
+           out = fs.create(new Path(dir, BLOOM_FILE_NAME), true);
+           bloomFilter.write(out);
+           out.flush();
+           out.close();
+           out = null;
+       } finally {
+           IOUtils.closeStream(out);
+       }
    }

```

³⁵ <https://issues.apache.org/jira/browse/HADOOP-7090>

Release reference. Any unused object in Java should be reclaimed by GC. If this object is still reachable by a live object, GC will not release its memory footprint. In such cases, the responsibility lies on the programmer to release the references preventing the object for being garbage collected (e.g., by nullifying the references to the unused objects). For example, HBASE-5141³⁶ reports a memory leak due to keeping references, even the corresponding task is finished. The fix patch nullifies the no-longer-needed objects. Following is the partial patch:

```

--- org/apache/hadoop/hbase/monitoring/MonitoredRPCHandlerImpl.java
+++ org/apache/hadoop/hbase/monitoring/MonitoredRPCHandlerImpl.java
@@ -217,6 +217,13 @@
...
+ @Override
+ public void markComplete(String status) {
+     super.markComplete(status);
+     this.params = null;
+     this.packet = null;
+ }
+

```

Proper exception handling (disposeExcp). Programmer should dispose of the objects or resources in all *exceptional* execution paths. Otherwise, a leak is likely to happen when an exception is thrown. Issue AMQ-3052³⁷ reports a memory leak in `securityContexts`. It occurs when the `addConnection()` fails after a successful authentication check. The developer fixed the bug via adding a `try-catch` block and calling disposal method in the `catch` block:

```

--- org/apache/activemq/security/SimpleAuthenticationBroker.java
+++ org/apache/activemq/security/SimpleAuthenticationBroker.java
@@ -92,7 +92,13 @@
...
- super.addConnection(context, info);
+ try {
+     super.addConnection(context, info);
+ } catch (Exception e) {
+     securityContexts.remove(s);
+     context.setSecurityContext(null);
+     throw e;
+ }

```

Remove element from collection (removeElem). No longer needed members of a collection should be removed by the programmer, allowing the garbage collector to reclaim the memory. A common repair action is the call of `remove()` method of a collection to clear useless elements from being referenced by the collection object. For example, in issue YARN-3472³⁸, already expired and removed tokens are not removed from `allTokens` map resulting in a potential memory leak. Developer fixed the issue by adding a call to `remove` method which removed the expired token from the map.

```

--- org/apache/hadoop/yarn/server/resourcemanager/security/DelegationTokenRenewer.java
+++ org/apache/hadoop/yarn/server/resourcemanager/security/DelegationTokenRenewer.java
@@ -577,6 +577,7 @@ private void requestNewHdfsDelegationTokenIfNeeded(
...
+ if (t.token.getKind().equals(new Text("HDFS_DELEGATION_TOKEN"))) {
+     iter.remove();
+     allTokens.remove(t.token);
+ }
...

```

³⁶ <https://issues.apache.org/jira/browse/HBASE-5141>

³⁷ <https://issues.apache.org/jira/browse/AMQ-3052>

³⁸ <https://issues.apache.org/jira/browse/YARN-3472>




Fig. 6 Heatmap of relationship between root causes and repair actions.

Shutdown finished thread (threadDown). A live thread of the application should be destroyed by the programmer when the thread task is completely finished. Adding a call to the shutdown method or adding a disposal method are the common repair actions for fixing the leaks caused by threads. HDFS-9003³⁹ reports a thread leak when a standby NameNode initializes the quota. Here, the thread pool is not shut down. To fix this bug, the developers added a call to the shutdown method.

```

--- org/apache/hadoop/hdfs/server/namenode/FSLmage.java
+++ org/apache/hadoop/hdfs/server/namenode/FSLmage.java
@@ -880,6 +880,7 @@ static void updateCountForQuota(BlockStoragePolicySuite bsps,
    root, counts);
    p.execute(task);
    task.join();
+   p.shutdown();

```

Results. In following, we show the results of our analysis on the repair patches in three sub-questions. First, we study the frequency of the repair actions. Second, we analyze the mapping between the root causes and the repair actions to find the relationship between these two taxonomies. Finally, we report the common code transformations found in the fix patches.

Finding 1. Table 7 lists the common repair actions along with the number of issues corresponding to them. 77% of the resource leaks are repaired with three major actions: *disposeReg*, *disposeExcep*, and *threadDown*, while 76% of the memory leaks are fixed with two repair actions *releaseRef* and *removeElm*.

Finding 2. Figure 6 reveals an almost one-to-one mapping between some root causes and repair actions (e.g., *exception* \rightarrow *disposeExcep*, *collection* \rightarrow *removeElm*, *concurrency* \rightarrow *threadSafe*, *concurrency* \rightarrow *threadSafe*). Leak defects with the root cause type *nonClosedRes* are repaired with repair actions *threadDown* and *disposeReg*. Leak defects with the root cause type *unreleaseRef* are repaired with repair actions *releaseRef* and *weakRef*.

³⁹ <https://issues.apache.org/jira/browse/HDFS-9003>

Table 8 Recurring code transformations and examples of code before and after the transformations.

- **Code transformation 1:** Conditional disposal of resource.
Example: `dispose(obj) → If (obj != null) obj.dispose()`
- **Code transformation 2:** Add disposal method call.
Example: `None → obj.dispose()`
- **Code transformation 3:** Add disposal method.
Example: `None → void dispose()`
- **Code transformation 4:** Set obsolete reference to null.
Example: `None → obj=null`
- **Code transformation 5:** Add catch/try-catch block.
Example: `Type obj = new Type() → try {Type obj = new Type()} exception {dispose(obj)}`
- **Code transformation 6:** Add finally/try-finally block
Example: `Type obj = new Type() → try {Type obj = new Type()} finally {dispose(obj)}`
- **Code transformation 7:** Add try-with-resources statement.
Example: `Type obj = new Type() → try {Type obj = new Type()}`
- **Code transformation 8:** Change condition expression.
Example: `If (cond1) obj.dispose() → If (cond1 and cond2) obj.dispose()`
- **Code transformation 9:** Change method call parameters.
Example: `obj.method(x, y) → obj.method(x, z)`
- **Code transformation 10:** Change static object to a non-static.
Example: `static Type obj = new Type() → Type obj = new Type()`
- **Code transformation 11:** Change to weak reference.
Example: `new HashMap<key, value>() → new HashMap<key, WeakReference(value)>()`
- **Code transformation 12:** Replace method call.
Example: `obj.method1() → obj.method2()`
- **Code transformation 13:** Change collection.
Example: `obj = new <collection1>() → obj = new <collection2>()`

Finding 3. We find 13 recurring patterns in the repair patches. Table 8 lists the recurring code transformations and the code examples before and after the transformations. We observe that 57% of the issues in the patch analysis dataset are partially or fully repaired with one or a combination of the recurring code transformations. This result shows that template-driven patch generation techniques (e.g., work Kim et al. (2013); Meng et al. (2013); Pan et al. (2009)) might be applicable for repairing the leak-related defects.

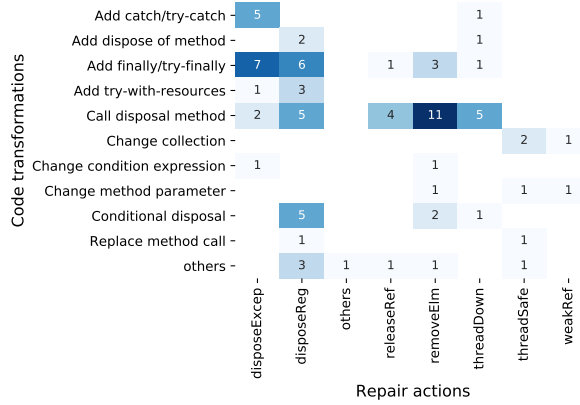


Fig. 7 Heatmap of recurring code transformations and single repair actions.

Finding 4. 57% of the issues in the patch analysis dataset are partially or fully repaired with one or a combination of recurring code transformations. This result shows that template-driven patch generation techniques (e.g., work Kim et al. (2013); Meng et al. (2013); Pan et al. (2009)) might be helpful for repairing the leak-related defects.

Finding 5. We find out that 83 (out of 452) issues are repaired with a single code transformation. We analyze the quantitative relationship between the repair actions and the most common code transformations. Figure 7 shows the heatmap of the quantitative analysis. For this heatmap, we only consider repair patches with a single code transformation. Code transformation *Add finally/try-finally* is often used in the repair actions *disposeReg* or *disposeExcep*. Code transformation *Add catch/try-catch* is the most used code transformation for repair action *disposeExcep*. We also observe a direct relationship between the repair action *RemoveElm* and the code transformation *Call disposal method*.

Summary. Overall, four repair actions are used by developers to repair over 80% of the issues in our dataset. About 57% of the fixed issues are repaired with one or a combination of 13 recurring code transformations.

4.6 RQ6: How complex are repairs of the leak-inducing defects?

Motivation. This research question addresses the complexity of changes applied to repair the leak-inducing defects. Besides this, we analyze the diagnosis time for different repair actions. We also compare the diagnosis time between leak-related and non-leak-related defects. In this question, we want to find how complex are the repair patches. The results can provide

more insights on how complex are the repair patches and how long does it take to repair a leak-inducing defect.

Approach. To assess the complexity of changes, we compute the code churn and change entropy (Hassan 2009).

Code churn is the sum of added and deleted lines in a repair patch. We only consider changes in the code statements and ignore comments or empty lines when calculating the code churn metric.

We use change entropy to find scatteredness of the changes. Derived from Shannon entropy in information theory, the change entropy measures the complexity of the changes. To measure the change entropy, we use the normalized Shannon entropy (Hassan 2009; Chen et al. 2016). It is defined as:

$$H = \frac{-\sum_{i=1}^n p(file_i) * \log_e p(file_i)}{\log_e(n)},$$

where n is the total number of files in a repair patch and $p(file_i)$ is defined as the number of lines changed in $file_i$ over the total number of lines changed in every file of that repair patch. Change entropy achieves its maximum value when all the files in a repair patch have the same number of modified lines. In contrast, we can achieve minimum entropy when only one file has the total number of modified lines. Using the entropy, we can find how complex are the repair patches. The higher the entropy, the more complex the repair patch.

To assess the diagnosis time, we collect two timestamps (i.e., created, and resolved) from each issue report. The created timestamp is the time a bug report is filed for the first time in the bug tracker. The resolved timestamp is the time of the latest patch applied to repair the bug. The diagnosis time is the difference between created and resolved timestamps. Strictly speaking, this time period can be further broken down to bug assignment, root-cause locating, patch design, and so on. Unfortunately, we cannot get such fine-grained information from the bug tracker. Besides, in some issues, the bug assignment timestamp took place after the repair patches were proposed by developers. In some other issues, no bug assignment is applied. Previous works also use similar timestamps to measure the diagnosis time of a bug (Song and Lu 2014).

Results. In following, we show the results of our analysis on the complexity of the repair patches.

Distribution of code churn. Figure 8 shows the box plot of code churn for different repair actions. The line within each box points to the median value of the code churn for that repair action.

Finding 1. In about all repair actions, the median of code churn is less than 20 lines of code while the repair action *disposeExcp* shows the highest median value.

Finding 2. Figure 9 shows the distribution of added and removed lines over studied projects. In all the projects, the median of added lines (29.5 lines) shows a larger value than the removed lines (16.5 lines). Hence, the fault repairing changes often increase the codebase of the applications.




Fig. 8 Distribution of code churn per repair action.




Fig. 9 Distribution of number of added and removed lines over studied project.




Fig. 10 Distribution of change complexity over the repair patches.

Finding 3. Figure 10 shows the distribution of change complexity over the repair patches. The distribution appears to be bimodal with the main peak around zero and a lower one around one. The change complexity analysis shows that the changes applied for repairing leak-inducing defects are more concentrated in fewer files and are less scattered. This result can be a useful finding for automated fault localization as it shows the high localization in leak-inducing defects.




Fig. 11 Distribution of diagnosis time per repair action.




Fig. 12 Diagnosis time comparison of leak-related and other bugs in studied projects.

Diagnosis time. Figure 11 shows distribution of diagnosis time across repair actions. Figure 12 shows the distribution of the diagnosis time for the leak-related and other defects in the studied projects. To calculate the diagnosis time of other defects, we collect the created and resolved timestamps of all bugs with the resolution "FIXED" (except the leak-related defects) from the studied projects in the same time frame that we collected the leak defects.

Finding 4. On median, about 6 days is needed for developers to fix a leak-inducing defect. This time is slightly lower than the diagnosis time for repairing non-leak defects (about 6.6 days). One reason could be that the leak-related defects are important for users and developers. The data in our dataset also confirms this. The issue priority in about 84% of the issues in our dataset are labelled as *Blocker*, *Critical*, or *Major* (which are the highest priority levels in the bug tracker). This corroborates with the assumption that leak-inducing defects impose a high negative impact on the performance of the applications, and are highly prioritized by development teams.

Table 9 The evaluation of Infer static analyzer on a sample of resource leaks from our dataset. Column "Detected?" reports whether Infer could detect the defect reported in the respective issue. "Code-based detection" refers to source code-based detection. "Defect" type and "Repair" type are explained in Section 4.4 and Section 4.5, respectively.

Issue	Detected?	Detection	Defect	Repair
AMQ-5745	✓	Code-based	nonClosedRes	disposeReg
AMQ-6051	No	Runtime	exception	disposeExcep
CASSANDRA-7709	No	Runtime	exception	disposeExcep
CASSANDRA-9134	No	Runtime	nonClosedRes	disposeReg
DERBY-5480	✓	Runtime	nonClosedRes	disposeReg
HADOOP-10203	No	Runtime	nonClosedRes	disposeReg
HADOOP-10490	✓	Runtime	nonClosedRes	disposeReg
HADOOP-11014	No	Code-based	exception	disposeExcep
HADOOP-11056	No	Code-based	exception	disposeExcep
HADOOP-11349	No	Code-based	exception	disposeExcep
HADOOP-11368	No	Runtime	nonClosedRes	threadDown
HADOOP-11414	No	Code-based	exception	disposeExcep
HADOOP-9681	✓	Runtime	nonClosedRes	disposeReg
HBASE-10461	No	Code-based	exception	disposeExcep
HBASE-10995	✓	Code-based	nonClosedRes	disposeReg
HBASE-13601	No	Runtime	exception	disposeExcep
HBASE-13797	✓	Code-based	nonClosedRes	disposeReg
HDFS-1118	No	Code-based	exception	disposeExcep
HDFS-1753	No	Code-based	exception	disposeExcep
HDFS-5099	No	Runtime	nonClosedRes	disposeReg
HDFS-5671	No	Runtime	exception	disposeExcep
HDFS-6208	No	Code-based	nonClosedRes	disposeReg
HDFS-6238	✓	Runtime	nonClosedRes	disposeReg
HIVE-12250	No	Runtime	nonClosedRes	disposeReg
HIVE-12790	No	Runtime	nonClosedRes	disposeReg
HIVE-13405	No	Code-based	exception	disposeExcep
MAPREDUCE-6528	No	Runtime	exception	disposeExcep
YARN-2484	No	Code-based	exception	disposeExcep
YARN-2988	✓	Code-based	nonClosedRes	disposeReg
YARN-4581	No	Runtime	exception	disposeExcep

Summary. The change entropy shows that the changes are more concentrated in fewer files and therefore less scattered. The median diagnosis time of the leak-inducing defects is about 6 days.

4.7 Other Results

In this section, we provide other findings found by our study.

Efficiency of static analysis tools. In RQ2 (Section 4.2), we showed that only in one issue (i.e., CASSANDRA-7709), the resource leak was reported using a static analyzer. There are many static analysis tools which support resource leak detection. Note that these tools mostly cannot detect memory leaks due to presence of garbage collector and lack of runtime information.

Table 10 Comparison of common code transformations found in our study with previous work. 27Repairs refers to Pan et al. (2009).

Our study	PAR	R2FIX	27Repairs
Conditional disposal of resource	✓	✓	✓
Add disposal method call	×	✓	×
Add disposal method	×	×	✓
Set obsolete reference to null	×	×	×
Add catch/try-catch block	×	×	✓
Add finally/try-finally block	×	×	×
Add try-with-resources statement	×	×	×
Change condition expression	✓	×	×
Change method call parameters	✓	×	✓
Change static object to a no-static	×	×	×
Change to weak reference	×	×	×
Replace method call	✓	×	✓
Change collection	×	×	×

However, one could ask why these tools are not mentioned in the studied bug issues. One reason might be that the developers already used such tools in development phase to reduce the potential leaks. We already showed that more than half of the studied leaks are resource leaks. It is interesting to know how many of these issues could be detected by static analysis tools.

For this purpose, we perform an evaluation on our dataset. We randomly select 30 issues reporting resource leaks from our dataset. As static analysis tool, we choose Infer which is used by large software organizations⁴⁰. We selected Infer because it is an open source tool and can detect resource leaks in Java. The applicability of Infer for resource leak detection is also confirmed in work van Tonder and Goues (2018).

Table 9 shows the result of our evaluation. From 30 issues, Infer was able to detect the leak defects reported in eight issues. To apply Infer, we first have to build the buggy version of the application in question which contains the leak. After a successful build, Infer produces a file reporting all potential resource leaks. Finally, we investigate whether the file contains the reported leak. We further investigate the eight issues detected by Infer to find the shared characteristics among those issues. In all cases, the leaks occurred in normal execution paths. The analysis shows that Infer was not able to detect leaks triggered in exceptional paths in the sample set. We also observe that developers repaired the leak defects by disposing of the unclosed resources in a `try-finally` block. This result can encourage the researcher and tool developers to improve current static analysis tools for leak detection.

Comparison of common code transformations. In RQ5 (Section 4.5), we showed 13 common code transformation found in the studied fix patches. Previous work also reported common repair patterns (Kim et al. 2013; Liu et al. 2013; Pan et al. 2009). PAR (Kim et al. 2013) found 10 manual repair patterns for Java. (Liu et al. 2013) used 8 patterns (2 of them for repairing

⁴⁰ <http://www.fbinfer.com>

memory leaks) to provide patches for bugs in C. Pan et al. (2009) introduced 27 automatically extractable repair patterns.

We compare our 13 patterns with previous work to find which patterns are not reported before. Table 10 shows the result of our evaluation. The result shows that six code transformations were not reported before. We can also observe that "conditional disposal of resource" was also used in all studied previous work.

5 Implications of the Study

Based on the findings of our empirical results, we discuss the implications of our study for both researchers and practitioners.

Prevalence of leak types. Understanding which types of leaks are prevalent in a project can help to avoid and detect leak defects efficiently. The results of Section 4.1 show that every studied projects has a particular dominant leak type. This knowledge can be exploited by prioritizing the most effective detection methods for the dominant leak types. As shown in Section 4.2, memory leaks and resource leaks have distinct best practice detection methods which can be used in a software development process.

Manual code inspection is the dominant detection method for resource leaks. Projects with a large number of resource leaks can benefit from this detection method. One can further improve this by using techniques like code self-review based on the Personal Software Process (PSP) (Humphrey 2000) with checklist items adapted for detection of resource leaks. For memory leaks, about 70.8% of the issues are detected or observed using the runtime information. Projects with a large number of memory leaks should consider the regular usage of the profiling tools. Profiling measures different metrics such as memory or time complexity of a program during its runtime. With this data, programmers can continuously check the resource usage of the program and react faster to the abnormal behavior.

In practical terms, the knowledge of the dominant leak types can be gained via (1) mining distribution of the leak types (or at least the dominant ones) from the bug trackers and repositories, and (2) improving the bug trackers with a labeled classification of the leaked resource.

Good practices. Good practices can considerably reduce the probability of introducing a leak defect. Such practices can be obtained for example from Java documentation or from existing research work. Here we describe two good practices.

Use try-with-resources on AutoCloseable resources. Introduced in Java 7, `try-with-resources` statement is an efficient method for better management of the closeable resources. It ensures that each resource is closed at the end of the try statement. Our empirical study shows that 33% of the resource leaks are caused by bad exception handling. The `try-with-resources` statement can help to avoid such defects as many current Java applications support Java 7 or higher.

Prevent having a strong reference from the value object to its key in a `HashMap`. As opposed to regular references, weak references do not protect the objects from being disposed of by the garbage collector. This property makes them suitable for implementing cache mechanisms through `WeakHashMap`, where the entry will be disposed of as soon as the key becomes unreachable. If the value objects of a `HashMap` refers to its own key, the programmer should wrap the value as `WeakReference` before putting the value into the map as recommended by the Java documentation⁴¹. Otherwise the key cannot be discarded.

Software testing for leak detection. Software tests can be used as a lightweight leak detection tool. They are beneficial for decreasing the cost of leak defects by triggering the leaks before the production phase. Our study shows that over 12% of the leak defects are detected as the result of a failing test (Table 4). Works like (Fang et al. 2015; Ghanavati and Andrzejak 2015) corroborate with our results by showing the effectiveness of leak detection via testing.

Fault localization. Fault localization is the first step in automated program repair. Defects with high locality can be repaired with low code churn. Our results showed that leak defects are highly localized. First, in 57% of the issues, only one file was modified. This value increases to 71% for repairs with changing two files at most. Second, in 96% of the issues, only Java files are changed. These findings can encourage researchers to improve and develop techniques for the automated repair of leak defects.

Template-driven patch generation. Previous works proposed patch-generation techniques based on the templates derived from existing human-written patches (Kim et al. 2013; Le et al. 2016). Work (Selakovic and Pradel 2016) showed the existence of common patterns for performance problems in JavaScript. We evaluated the potential of providing template-driven repairs for leak defects through studying repair patches. We found 13 common code transformations used by developers. About 57% of the issues from patch analysis dataset are repaired by a combination of one or more of these code transformations. These results show the potential of template-driven patch generation techniques for repairing leak-inducing defects.

6 Threats to Validity

In this section, we discuss the threats to the validity of our study.

Construct validity. The quality of dataset used in our study is a threat to construct validity. We used JIRA as the bug tracker to collect leak defects. This set of defects are not necessarily include all leak defects in the studied applications. These sets might be different since most likely not all leak defects from set B are reported in the issue tracker. Conversely, some investigated

⁴¹ <https://docs.oracle.com/javase/7/docs/api/java/util/WeakHashMap.html>

defects might never manifest at runtime. This might be especially the case for issues found by source code-based detection (see Table 4). However, since we investigate a large number of defects and focus on distributions and their relations, we expect that our findings describe the characteristics of the whole defect population in general.

We also used a simple keyword search to find leak-related bugs. Issues that do not contain the keyword "leak" could skip our data collection process. We searched for other leak-related keywords, but our query yield many false positives. To minimize such threats related to insufficient or skewed sampling of the leak defect population, we used a large set of leak-related bugs (452 issues) from 10 large-scale projects from a variety of application categories.

Internal validity. The experimenter bias and errors are threats to internal validity. In this study, we heavily use manual analysis. When generating taxonomies defined in our study, we manually extract the contents of the issues and use our knowledge to assign a bug to a category. To mitigate this problem, the authors involved in the labeling process discussed any conflicts to reach a consensus. We have spent many hours studying all data related to each defect such as issue title and description, developer discussions, and repair patches. We applied Cohen's kappa metric to measure the inter-rater agreement. The kappa values ranged from 0.56 to 0.86 which shows a substantial agreement in most of the dimensions. We also make our dataset available online to improve the replicability of our study.

External validity. Threats to external validity relate to the generalizability of our findings and implications. We collect our dataset from different categories of open source projects. The projects may not be representative for closed source projects. Our results are derived from 10 Java projects and some findings may not apply to projects written in other languages.

7 Related Work

There is a large body of work in detection, localization, and repairing functional and non-functional bugs. Here we cover work related to our study, grouped in three research directions.

Empirical study. There are many work studying characteristics of bugs (Zhong and Su 2015; Jin et al. 2012b; Song and Lu 2014; Selakovic and Pradel 2016). Zhong and Su (2015) extracted and analyzed the characteristics of the real bug fixes from six Java projects. Close to our study are work of Machida et al. (2012) and Tan et al. (2014) which investigated memory-related bugs. Machida et al. (2012) investigated five open source Java projects related to cloud computing and found 55 leak-related defects. They showed that in all studied projects leak-related bugs exist with the ratio ranged from 0.4% to 1.4% of the total bugs. The majority of 55 leaks were file descriptor leaks comprising of 30% of the cases. Tan et al. (2014) studied the characteristics of three open source projects Written in C. They

showed that memory-related bugs is one of three main causes of bugs (in addition to concurrency and semantic bugs). They found that 16.7 to 40.0% of the memory bugs across the studied projects are caused by memory leaks. They also showed that the severity of memory leaks is high as most of them result in a crash.

Our study differs from the above mentioned studies. We studied both resource and memory leak-related defects from 10 open source Java applications. We performed an in-depth analysis on leak defects and their repairs providing taxonomies for leak type, leak detection, fault localization, root-causes, and repair actions. We found that there are common repair patterns for fixing the leak defects. We also evaluated the complexity of the the repair patches. Finally, we drawn actionable implications based on our observations and findings. Hence, we believe that our study considerably differs from the previous work in both size of the studied dataset and depth of analysis.

Automated diagnosis of memory and resource leaks. Various techniques are proposed by researchers to diagnose memory and resource leaks.

Memory leaks. Static analysis is used to detect memory leaks (Xie and Aiken 2005; Heine and Lam 2003; Cherem et al. 2007; Orlovich and Rugina 2006). These approaches like other static approaches in fault localization suffer from the lack of scalability and precision. LeakChecker (Yan et al. 2014) tries to decrease the inaccuracy via investigating loops provided by the developer as an oracle for memory leak detection.

To mitigate the problems of static analysis, other researchers leverages dynamic analysis to detect memory leaks. The major directions of dynamic leak detection are: staleness detection (Hauswirth and Chilimbi 2004; Bond and McKinley 2006; Novark et al. 2009; Jung et al. 2014), growth analysis (Jump and McKinley 2007; Sor et al. 2013; Matias et al. 2014; Fang et al. 2015; Ghanavati and Andrzejak 2015), analysis of captured state (Mitchell and Sevitsky 2003; Clause and Orso 2010; Xu et al. 2011), and hybrid approaches (Xu and Rountev 2008; Rayside and Mendel 2007).

(Xu and Rountev 2008) focuses on the memory leak defects related to collections and try to rank the suspicious statements by assigning a leak confidence value based on staleness and memory usage. In our study, we quantitatively show that collections are one of the major root causes of the memory leaks defects.

Some studies introduced approaches to tolerate the memory leaks by keeping the program in a running state (Bond and McKinley 2008; Rayside and Mendel 2007; Bond and McKinley 2009). They achieve this by reducing the performance degradation (e.g., with predicting and reclaiming the leaky objects at runtime).

Resource leaks. Many approaches have purposed to detect resource leaks in Java and C (Dillig et al. 2008; Torlak and Chandra 2010; Shaham et al. 2003; Cherem and Rugina 2004; Weimer and Nacula 2004b,a). They usually use static analysis techniques to find the unclosed resources in different execution

paths. There are also researches which try to detect resource leaks in Android (Guo et al. 2013; Banerjee et al. 2018).

Automated leak repair. Recently, automated program repair attracted the attention of researchers. Pioneering work GenProg (Weimer et al. 2009) introduced a patch generation technique based on a genetic search algorithm. Kim et al. (2013) proposed an automated program repair technique based on patterns learned from real patches. It generates correct patches for 27 out of 119 bugs. All the provided fix patterns are simple and one-line statements. Prophet (Long and Rinard 2016) learns properties of successful patches to guide finding the appropriate candidates. HDRRepair (Le et al. 2016) leverages information derived from history of the previous patches of hundreds Java projects to select the correct patch candidates. All the mentioned techniques differ in defining the search space and the approach to find the accurate patch.

Semantic-based techniques have also been explored (Nguyen et al. 2013; Mechtaev et al. 2016). Angelix (Mechtaev et al. 2016) is a good example of this category which extracts semantic constraints from the application codebase and generates fixes using program synthesis.

Automated leak repair is still embryonic, and only few works exist in literature (van Tonder and Goues 2018; Gao et al. 2015; Yan et al. 2016; Liu et al. 2013). van Tonder and Goues (2018) purposed "Footpatch" on top of Infer. Footpatch could generate fixes for resource leaks in C and Java as well as fixes for memory leaks in C. (Gao et al. 2015; Yan et al. 2016) leveraged static and dynamic analysis to fix memory leaks in C. They analyze the execution paths of each allocation/deallocation and insert a `free` when no release is encountered. Liu et al. (2013) used two repair patterns (*AddFree* and *MvFree*) and provide correct patches for 16 out of 41 memory leaks in C.

8 Conclusions and Future Work

Diagnosis of leak-inducing defects are one of the main challenges for both researchers and practitioners in software development and maintenance. Understanding the characteristics of resource and memory leaks can provide useful information to further improve leak diagnosis techniques. For this purpose, we conducted a detailed empirical study on a large dataset (452 issues from 10 mature projects) to understand how leaks are detected, which defects create them, and which types of repairs exist. Our findings and implications showed that even by simple changes in the quality assurance processes (e.g., code review, testing), the avoidance and diagnosis of leaks could be significantly improved.

In our future work, we will investigate why automated leak detection tools are rarely used for leak detection in practice. We will also evaluate approaches for automated repair of the leak-inducing defects with the focus on template-driven patch generation techniques. We plan to implement a fault injector which simulates the distribution of the leak types and the defect types in real

applications. It can serve as a realistic benchmarking tool for evaluation of methods and tools for leak diagnosis.

References

- Artstein R, Poesio M (2008) Inter-coder agreement for computational linguistics. *Comput Linguist* 34(4):555–596, DOI 10.1162/coli.07-034-R2, URL <http://dx.doi.org/10.1162/coli.07-034-R2>
- Banerjee A, Chong LK, Ballabriga C, Roychoudhury A (2018) Energypatch: Repairing resource leaks to improve energy-efficiency of android apps. *IEEE Transactions on Software Engineering* 44(5):470–490, DOI 10.1109/TSE.2017.2689012, URL doi.ieeecomputersociety.org/10.1109/TSE.2017.2689012
- Bond MD, McKinley KS (2006) Bell: Bit-encoding online memory leak detection. In: *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XII*, pp 61–72, DOI 10.1145/1168857.1168866, URL <http://doi.acm.org/10.1145/1168857.1168866>
- Bond MD, McKinley KS (2008) Tolerating memory leaks. In: *Proceedings of the 23rd ACM SIGPLAN Conference on Object-oriented Programming Systems Languages and Applications, OOPSLA '08*, pp 109–126, DOI 10.1145/1449764.1449774, URL <http://doi.acm.org/10.1145/1449764.1449774>
- Bond MD, McKinley KS (2009) Leak pruning. In: *ACM Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), ASPLOS XIV*, pp 277–288, DOI 10.1145/1508244.1508277, URL <http://doi.acm.org/10.1145/1508244.1508277>
- Chen TH, Shang W, Yang J, Hassan AE, Godfrey MW, Nasser M, Flora P (2016) An empirical study on the practice of maintaining object-relational mapping code in java systems. In: *Proceedings of the 13th International Conference on Mining Software Repositories, ACM, MSR '16*, pp 165–176, DOI 10.1145/2901739.2901758, URL <http://doi.acm.org/10.1145/2901739.2901758>
- Cherem S, Rugina R (2004) Region analysis and transformation for java programs. In: *Proceedings of the 4th International Symposium on Memory Management, ISMM '04*, pp 85–96, DOI 10.1145/1029873.1029884, URL <http://doi.acm.org/10.1145/1029873.1029884>
- Cherem S, Princehouse L, Rugina R (2007) Practical memory leak detection using guarded value-flow analysis. In: *Proceedings of the 2007 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '07*, pp 480–491, DOI 10.1145/1250734.1250789, URL <http://doi.acm.org/10.1145/1250734.1250789>
- Clause J, Orso A (2010) Leakpoint: Pinpointing the causes of memory leaks. In: *Proceedings of the 32Nd ACM/IEEE International Conference on Software Engineering - Volume 1, ICSE '10*, pp 515–524, DOI 10.1145/1806799.1806874, URL <http://doi.acm.org/10.1145/1806799.1806874>
- Cohen J (1960) A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20(1):37–46, DOI 10.1177/001316446002000104, URL <https://doi.org/10.1177/001316446002000104>, <https://doi.org/10.1177/001316446002000104>
- Dillig I, Dillig T, Yahav E, Chandra S (2008) The closer: Automating resource management in java. In: *ACM International Symposium on Memory Management (ISMM)*, pp 1–10, DOI 10.1145/1375634.1375636, URL <http://doi.acm.org/10.1145/1375634.1375636>
- Fang L, Dou L, Xu GH (2015) Perflblower: Quickly detecting memory-related performance problems via amplification. In: *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic*, pp 296–320, DOI 10.4230/LIPIcs.ECOOP.2015.296, URL <http://dx.doi.org/10.4230/LIPIcs.ECOOP.2015.296>
- Gao Q, Xiong Y, Mi Y, Zhang L, Yang W, Zhou Z, Xie B, Mei H (2015) Safe memory-leak fixing for c programs. In: *Proceedings of the 37th International Conference on Software Engineering - Volume 1, ICSE '15*, pp 459–470, URL <http://dl.acm.org/citation.cfm?id=2818754.2818812>
- Ghanavati M, Andrzejak A (2015) Automated memory leak diagnosis by regression testing. In: *2015 IEEE 15th International Working Conference on Source Code Analysis and*

- Manipulation (SCAM), pp 191–200, DOI 10.1109/SCAM.2015.7335415
- Ghanavati M, Costa D, Andrzejak A, Seboek J (2018) Memory and resource leak defects in java projects: An empirical study. In: Proceedings of the 38th International Conference on Software Engineering Companion, ICSE '18 Companion, pp 410–411, DOI 10.1145/3183440.3195032, URL <https://doi.org/10.1145/3183440.3195032>
- Guo C, Zhang J, Yan J, Zhang Z, Zhang Y (2013) Characterizing and detecting resource leaks in android applications. In: Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering, IEEE Press, Piscataway, NJ, USA, ASE'13, pp 389–398, DOI 10.1109/ASE.2013.6693097, URL <https://doi.org/10.1109/ASE.2013.6693097>
- Hassan AE (2009) Predicting faults using the complexity of code changes. In: Proceedings of the 31st International Conference on Software Engineering, IEEE Computer Society, ICSE '09, pp 78–88, DOI 10.1109/ICSE.2009.5070510, URL <http://dx.doi.org/10.1109/ICSE.2009.5070510>
- Hauswirth M, Chilimbi TM (2004) Low-overhead memory leak detection using adaptive statistical profiling. In: Proceedings of the 11th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XI, pp 156–164, DOI 10.1145/1024393.1024412, URL <http://doi.acm.org/10.1145/1024393.1024412>
- Heine DL, Lam MS (2003) A practical flow-sensitive and context-sensitive c and c++ memory leak detector. In: Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation, PLDI '03, pp 168–181, DOI 10.1145/781131.781150, URL <http://doi.acm.org/10.1145/781131.781150>
- Huizinga D, Kolawa A (2007) Automated defect prevention: best practices in software management. John Wiley & Sons, USA
- Humphrey W (2000) The personal software process (psp). Tech. Rep. CMU/SEI-2000-TR-022, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, URL <http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=5283>
- Jin G, Song L, Shi X, Scherpelz J, Lu S (2012a) Understanding and detecting real-world performance bugs. In: Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, pp 77–88, DOI 10.1145/2254064.2254075, URL <http://doi.acm.org/10.1145/2254064.2254075>
- Jin G, Song L, Shi X, Scherpelz J, Lu S (2012b) Understanding and detecting real-world performance bugs. In: Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, pp 77–88, DOI 10.1145/2254064.2254075, URL <http://doi.acm.org/10.1145/2254064.2254075>
- Jump M, McKinley KS (2007) Cork: Dynamic memory leak detection for garbage-collected languages. In: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp 31–38, DOI 10.1145/1190216.1190224, URL <http://doi.acm.org/10.1145/1190216.1190224>
- Jung C, Lee S, Raman E, Pande S (2014) Automated memory leak detection for production use. In: Proceedings of the 36th International Conference on Software Engineering, ICSE 2014, pp 825–836, DOI 10.1145/2568225.2568311, URL <http://doi.acm.org/10.1145/2568225.2568311>
- Kim D, Nam J, Song J, Kim S (2013) Automatic patch generation learned from human-written patches. In: Proceedings of the 2013 International Conference on Software Engineering, ICSE '13, pp 802–811, URL <http://dl.acm.org/citation.cfm?id=2486788.2486893>
- Le XBD, Lo D, Le Goues C (2016) History driven program repair. In: IEEE 23rd International Conference on Software Analysis, Evolution and Reengineering, IEEE, SANER '16
- Liu C, Yang J, Tan L, Hafiz M (2013) R2fix: Automatically generating bug fixes from bug reports. In: Proceedings of the 2013 IEEE Sixth International Conference on Software Testing, Verification and Validation, ICST '13, pp 282–291, DOI 10.1109/ICST.2013.24, URL <https://doi.org/10.1109/ICST.2013.24>
- Lo D, Nagappan N, Zimmermann T (2015) How practitioners perceive the relevance of software engineering research. In: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ACM, New York, NY, USA, ESEC/FSE 2015,

- pp 415–425, DOI 10.1145/2786805.2786809, URL <http://doi.acm.org/10.1145/2786805.2786809>
- Long F, Rinard M (2016) Automatic patch generation by learning correct code. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '16, pp 298–312, DOI 10.1145/2837614.2837617, URL <http://doi.acm.org/10.1145/2837614.2837617>
- Machida F, Xiang J, Tadano K, Maeno Y (2012) Aging-related bugs in cloud computing software. In: Proceedings of the 2012 IEEE 23rd International Symposium on Software Reliability Engineering Workshops, ISSREW '12, pp 287–292, DOI 10.1109/ISSREW.2012.97, URL <http://dx.doi.org/10.1109/ISSREW.2012.97>
- Matias R, Andrzejak A, Machida F, Elias D, Trivedi K (2014) A systematic differential analysis for fast and robust detection of software aging. In: 2014 IEEE 33rd International Symposium on Reliable Distributed Systems, pp 311–320, DOI 10.1109/SRDS.2014.38
- McConnell S (1993) Code complete: a practical handbook of software construction. Microsoft Press, USA
- Mechtaev S, Yi J, Roychoudhury A (2016) Angelix: Scalable multiline program patch synthesis via symbolic analysis. In: Proceedings of the 38th International Conference on Software Engineering, ICSE '16, pp 691–701, DOI 10.1145/2884781.2884807, URL <http://doi.acm.org/10.1145/2884781.2884807>
- Meng N, Kim M, McKinley KS (2013) Lase: Locating and applying systematic edits by learning from examples. In: Proceedings of the 2013 International Conference on Software Engineering, IEEE Press, Piscataway, NJ, USA, ICSE '13, pp 502–511, URL <http://dl.acm.org/citation.cfm?id=2486788.2486855>
- Mitchell N, Sevitsky G (2003) LeakBot: An automated and lightweight tool for diagnosing memory leaks in large Java applications. Lecture Notes in Computer Science 2743:351–377
- Nguyen HDT, Qi D, Roychoudhury A, Chandra S (2013) Semfix: Program repair via semantic analysis. In: Proceedings of the 2013 International Conference on Software Engineering, ICSE '13, pp 772–781, URL <http://dl.acm.org/citation.cfm?id=2486788.2486890>
- Nguyen K, Wang K, Bu Y, Fang L, Hu J, Xu G (2015) Facade: A compiler and runtime for (almost) object-bounded big data applications. In: Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15, pp 675–690
- Nistor A, Jiang T, Tan L (2013) Discovering, reporting, and fixing performance bugs. In: Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13, pp 237–246, URL <http://dl.acm.org/citation.cfm?id=2487085.2487134>
- Novark G, Berger ED, Zorn BG (2009) Efficiently and precisely locating memory leaks and bloat. In: Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '09, pp 397–407, DOI 10.1145/1542476.1542521, URL <http://doi.acm.org/10.1145/1542476.1542521>
- Orlovich M, Rugina R (2006) Memory leak analysis by contradiction. In: Proceedings of the 13th International Conference on Static Analysis, SAS'06, pp 405–424, DOI 10.1007/11823230_26, URL http://dx.doi.org/10.1007/11823230_26
- Pan K, Kim S, Whitehead EJ Jr (2009) Toward an understanding of bug fix patterns. Empirical Softw Engg 14(3):286–315, DOI 10.1007/s10664-008-9077-5, URL <http://dx.doi.org/10.1007/s10664-008-9077-5>
- Rayside D, Mendel L (2007) Object ownership profiling: A technique for finding and fixing memory leaks. In: Proceedings of the Twenty-second IEEE/ACM International Conference on Automated Software Engineering, ASE '07, pp 194–203
- Seaman CB (1999) Qualitative methods in empirical studies of software engineering. IEEE Trans Softw Eng 25(4):557–572, DOI 10.1109/32.799955, URL <http://dx.doi.org/10.1109/32.799955>
- Seaman CB, Shull F, Regardie M, Elbert D, Feldmann RL, Guo Y, Godfrey S (2008) Defect categorization: Making use of a decade of widely varying historical data. In: Proceedings of the Second ACM-IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM '08, pp 149–157, DOI 10.1145/1414004.1414030, URL <http://doi.acm.org/10.1145/1414004.1414030>

- Selakovic M, Pradel M (2016) Performance issues and optimizations in javascript: An empirical study. In: Proceedings of the 38th International Conference on Software Engineering, ACM, New York, NY, USA, ICSE '16, pp 61–72, DOI 10.1145/2884781.2884829, URL <http://doi.acm.org/10.1145/2884781.2884829>
- Shaham R, Yahav E, Kolodner EK, Sagiv M (2003) Establishing local temporal heap safety properties with applications to compile-time memory management. In: Proceedings of the 10th International Conference on Static Analysis, SAS'03, pp 483–503, URL <http://dl.acm.org/citation.cfm?id=1760267.1760304>
- Sommerville I (2010) Software Engineering, 9th edn. Addison-Wesley Publishing Company, USA
- Song L, Lu S (2014) Statistical debugging for real-world performance problems. In: Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '14, pp 561–578, DOI 10.1145/2660193.2660234, URL <http://doi.acm.org/10.1145/2660193.2660234>
- Sor V, Oü P, Treier T, Srirama SN (2013) Improving statistical approach for memory leak detection using machine learning. In: 2013 IEEE International Conference on Software Maintenance, pp 544–547, DOI 10.1109/ICSM.2013.92
- Tan L, Liu C, Li Z, Wang X, Zhou Y, Zhai C (2014) Bug characteristics in open source software. *Empirical Softw Engg* 19(6):1665–1705, DOI 10.1007/s10664-013-9258-8, URL <http://dx.doi.org/10.1007/s10664-013-9258-8>
- van Tonder R, Goues CL (2018) Static automated program repair for heap properties. In: Proceedings of the 40th International Conference on Software Engineering, ACM, New York, NY, USA, ICSE '18, pp 151–162, DOI 10.1145/3180155.3180250, URL <http://doi.acm.org/10.1145/3180155.3180250>
- Torlak E, Chandra S (2010) Effective interprocedural resource leak detection. In: Proceedings of the 32Nd ACM/IEEE International Conference on Software Engineering - Volume 1, ICSE '10, pp 535–544, DOI 10.1145/1806799.1806876, URL <http://doi.acm.org/10.1145/1806799.1806876>
- Weimer W, Necula GC (2004a) Finding and preventing run-time error handling mistakes. In: Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications, ACM, New York, NY, USA, OOPSLA '04, pp 419–431, DOI 10.1145/1028976.1029011, URL <http://doi.acm.org/10.1145/1028976.1029011>
- Weimer W, Necula GC (2004b) Finding and preventing run-time error handling mistakes. In: Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications, OOPSLA '04, pp 419–431, DOI 10.1145/1028976.1029011, URL <http://doi.acm.org/10.1145/1028976.1029011>
- Weimer W, Necula GC (2005) Mining temporal specifications for error detection. In: Proceedings of the 11th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, Springer-Verlag, Berlin, Heidelberg, TACAS'05, pp 461–476, DOI 10.1007/978-3-540-31980-1_30, URL http://dx.doi.org/10.1007/978-3-540-31980-1_30
- Weimer W, Nguyen T, Le Goues C, Forrest S (2009) Automatically finding patches using genetic programming. In: Proceedings of the 31st International Conference on Software Engineering, ICSE '09, pp 364–374, DOI 10.1109/ICSE.2009.5070536, URL <http://dx.doi.org/10.1109/ICSE.2009.5070536>
- Wu R, Zhang H, Kim S, Cheung SC (2011) Relink: Recovering links between bugs and changes. In: Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering, ESEC/FSE '11, pp 15–25, DOI 10.1145/2025113.2025120, URL <http://doi.acm.org/10.1145/2025113.2025120>
- Xie Y, Aiken A (2005) Context- and path-sensitive memory leak detection. In: Proceedings of the 10th European Software Engineering Conference Held Jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering, ESEC/FSE-13, pp 115–125, DOI 10.1145/1081706.1081728, URL <http://doi.acm.org/10.1145/1081706.1081728>
- Xu G, Rountev A (2008) Precise memory leak detection for java software using container profiling. In: Proceedings of the 30th International Conference on Software Engineering, ICSE '08, pp 151–160, DOI 10.1145/1368088.1368110, URL <http://doi.acm.org/>

- 10.1145/1368088.1368110
- Xu G, Bond MD, Qin F, Rountev A (2011) Leakchaser: Helping programmers narrow down causes of memory leaks. In: Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '11, pp 270–282, DOI 10.1145/1993498.1993530, URL <http://doi.acm.org/10.1145/1993498.1993530>
- Xu GH, Rountev A (2013) Precise memory leak detection for java software using container profiling. *ACM Trans Softw Eng Methodol* 22(3):17:1–17:28, DOI 10.1145/2491509.2491511, URL <http://doi.acm.org/10.1145/2491509.2491511>
- Yan D, Xu G, Yang S, Rountev A (2014) Leakchecker: Practical static memory leak detection for managed languages. In: Proceedings of Annual IEEE/ACM International Symposium on Code Generation and Optimization, CGO '14, pp 87:87–87:97, DOI 10.1145/2544137.2544151, URL <http://doi.acm.org/10.1145/2544137.2544151>
- Yan H, Sui Y, Chen S, Xue J (2016) Automated memory leak fixing on value-flow slices for c programs. In: Proceedings of the 31st Annual ACM Symposium on Applied Computing, SAC '16, pp 1386–1393, DOI 10.1145/2851613.2851773, URL <http://doi.acm.org/10.1145/2851613.2851773>
- Yin Z, Yuan D, Zhou Y, Pasupathy S, Bairavasundaram L (2011) How do fixes become bugs? In: Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering, ACM, New York, NY, USA, ESEC/FSE '11, pp 26–36, DOI 10.1145/2025113.2025121, URL <http://doi.acm.org/10.1145/2025113.2025121>
- Zhong H, Su Z (2015) An empirical study on real bug fixes. In: Proceedings of the 37th International Conference on Software Engineering - Volume 1, ICSE '15, pp 913–923, URL <http://dl.acm.org/citation.cfm?id=2818754.2818864>