Does Return Null Matter?

Shuhei Kimura, Keisuke Hotta, Yoshiki Higo, Hiroshi Igaki, and Shinji Kusumoto
Graduate School of Information Science and Technology, Osaka University
1-5, Yamadaoka, Suita, Osaka 565-0871, Japan
{s-kimura, k-hotta, higo, igaki, kusumoto}@ist.osaka-u.ac.jp

***Abstract*—Developers often use null references for the returned values of methods (return null) in object-oriented languages. Although developers often use return null to indicate that a program does not satisfy some necessary conditions, it is generally felt that a method returning null is costly to maintain. One of the reasons for is that when a method receives a value returned from a method invocation whose code includes return null, it is necessary to check whether the returned value is null or not (null check). As developers often forget to write null checks, null dereferences occur frequently. However, it has not been clarified to what degree return null affects software maintenance during software evolution. This paper shows the influences of return null by investigating return null and null check in the evolution of source code. Experiments conducted on 14 open source projects showed that developers modify return null more frequently than return statements that do not include null. This result indicates that return null has a negative effect on software maintenance. It was also found that the size and the development phases of projects have no effect on the frequency of modifications on return null and null check. In addition, we found that all the projects in this experiment had from one to four null checks per 100 lines.**

***Index Terms*—software evolution; return null; software maintenance**

I. Introduction

A null reference is a common mechanism in object-oriented languages, such as Java and C++. Developers use $null$ in various roles, for instance, as an initializer or sentinel. In particular, it is general practice that a method returns $null$ values when some necessary conditions are not satisfied (hereafter we refer to return statements whose operand is $null$ as "**return null**").

Developers often use a return null rather than error constants or exceptions for the following reasons:

- As $null$ is a unique value, we know some errors have occurred, even if there are no error messages.
- Returning $null$ is easier to write than handling exceptions.

Although return null supports developers for the above reasons, it is generally felt that return null is one of the factors that increase software maintenance efforts [1].

A null dereference, a bug by which a program dereferences a null value, is the main bug caused by a return null. To prevent a null dereference from occurring, returned values from a method that includes a return null need to be checked (hereafter we refer to such checking as "**null check**"). Accordingly, when a developer adds return null to a method, a null check needs to be added for all of its invocations. If a developer forgets to add null check, the program will try to dereference null, and NullPointerException will be thrown. In addition, $null$ contains no information about an error; whereas an exception contains information on errors that have occurred. For this reason, it is very difficult for caller methods that have received $null$ to know what errors have occurred. This characteristic sometimes hides the root causes of errors that have occurred, because a null dereference will occur at places remote from the point of origin of $null$ if a program has methods that include return null in a null check (e.g., Figure 1). For these reasons, return null is regarded as a factor that increases the cost of software maintenance.

```
1 final RevCommit base = walk.next();
2 if(base == null)
3   return null;
```

Fig. 1: Example of passing return null

However, it has not been clarified to what extent a return null affects software maintenance during software evolution. Therefore, in this research, we conducted an experimental study on 14 open source projects to ascertain whether, and to what extent, the presence of return null and null check affect software maintenance.

We found that the presence of return null and null check increases the cost of software maintenance, and that the source code modifications that are related to them have the following characteristics:

- Developers modified return nulls more frequently than return statements that do not include $null$.
- The size and the development phases of projects did not affect the frequency of modifications on return nulls and null checks.
- The density of null checks in source code was between one and four in 100 lines, in any project. If the density of null checks in a project exceeds 0.04, the project is presumed to have a problem. None of the projects had a density less than 0.01, and none of the projects had special measures in place for return nulls and null checks.
- There were modifications that replaced return null with exception handling or an erased null check. These phenomena indicate that programmers considered that the presence of return null and null check were harmful to software evolution.

The contributions of this paper include the following:

- This paper revealed that return null affect software maintenance during software evolution.
- The result indicated that developers should avoid writing return null without given a great deal of thought to doing so.
- Measurement of return null and null check showed that return null and null check are costly to maintain throughout the entire development periods.
- The result showed the average density of null checks. This value can be used as a criterion (e.g., when the density of null check in a project exceeds the average, developers should apply refactoring to return null and null check).

A developer could replace return null with other mechanisms, for example, exception handling, a proper object such as an empty array, or a NullObject pattern. However, this paper presents whether return null matter. Providing how to replace it is our future work.

The rest of this paper is organized as follows. Section II presents previous work by way of introducing our research. Section III defines our research questions and describes the experimental design. Section IV shows the result of the experiment and answers the research questions. In Section V, we discuss our answers to the research questions. Section VI discusses threats posed to validity. Section VII provides various approaches to resolving problems related to $null$, and Section VIII concludes this paper.

II. Background

Null dereference often occurs in programs written in object-oriented programming languages. For this reason, many researchers have proposed techniques to detect null dereference [2]–[7], and many tools have been released [8], such as FindBugs [9], SALSA [6], JLint, and ESC/Java [10]. However, previous studies did not mention whether, and to what extent, the presence of return null and null check affected software maintenance during software evolution.

A factor of null dereference is that a program unintentionally refers $null$ when programmers forget to write null check, despite the callee method possibly returning $null$. Figure 2(b) is a code fragment that appears in JGit. command.call(),in the fourth line, invokes the method shown in Figure 2(a). This method may return $null$, as written in the ninth line. The variable named "ref," defined in the fourth line in Figure 2(b), may be assigned $null$ by calling command.call(). As a result, ref.getName(), in the fifth line, possibly causes null dereference. Figure 2(c) is a bug fix for this null dereference. Prefix "+" in a line means that the line was added in this commit. This fix added null check in order to avoid the occurrences of null dereference. The commit message, "Do not fail when checking out HEAD," shows the purpose of this commit is to fix the null dereference. As mentioned above, return null, null check, and null dereference are closely related to each other. Thus, we can show how null dereference affects software maintenance by investigating the effect of return null and null check on software maintenance.

```
1  public Ref call() throws GitAPIException,
       RefAlreadyExistsException, RefNotFoundException,
       InvalidRefNameException, CheckoutConflictException
       {
2    checkCallable();
3    processOptions();
4    try {
5      if (checkoutAllPaths || !paths.isEmpty()) {
6        checkoutPaths();
7        status = new CheckoutResult(Status.OK, paths);
8        setCallable(false);
9        return null;
10   }
11 ...
```

(a) Method including return null

```
1  ...
2  try {
3    String oldBranch = db.getBranch();
4    Ref ref = command.call();
5    if (Repository.shortenRefName(ref.getName()).equals(
       oldBranch)) {
6      outw.println(MessageFormat.format(
7        CLIText.get().alreadyOnBranch,
8  ...
```

(b) Invokes the method shown in (a)

```
1  ...
2  try {
3    String oldBranch = db.getBranch();
4    Ref ref = command.call();
5  + if (ref == null)
6  +   return;
7    if (Repository.shortenRefName(ref.getName()).equals(
       oldBranch)) {
8      outw.println(MessageFormat.format(
9        CLIText.get().alreadyOnBranch,
10 ...
```

(c) After adding null check

Fig. 2: Example code fragments of return null, null dereference, and null check

To reveal the effect of return null, the experiment in this paper covers not only return null but also null check. Error constants, such as -1, are used as having the same meaning as return null. However, mainly in Java, as all the variables — except primitive ones — can be assigned $null$, $null$ probably has a greater negative effect on software maintenance than error constants. For this reason, this research focuses on return null and null check.

In the remainder of this paper, we use the phrase "a statement is costly to maintain" to indicate "developers modify the statement many times" and "the statement has a trend that new bugs are likely to be introduced into it."

III. Experimental Design

In this section, we explain the design of our experiment on open source projects. Herein, we use the abbreviations listed in Table I for convenience. In addition, $|S|$ means the number of elements in a given set S. Let c be a commit between revision r and $r+1$, and then the set of added/deleted ret_{null}, ret_{not}, $cond_{null}$, and $cond_{not}$ in c are defined as ΔRet^c_{null}, ΔRet^c_{not}, $\Delta Cond^c_{null}$, and $\Delta Cond^c_{not}$.

TABLE I: Abbreviations

Abbreviation	Explanation of Abbreviation
ret_{null}	return statements whose operands are $null$
ret_{not}	return statements whose operands are NOT $null$
$cond_{null}$	conditional expressions having comparison with $null$
$cond_{not}$	conditional expressions NOT having comparison with $null$
Ret^r_{null}	a set of ret_{null} in revision r
Ret^r_{not}	a set of ret_{not} in revision r
Ret^r	$Ret^r_{null} \cup Ret^r_{not}$
$Cond^r_{null}$	a set of $cond_{null}$ in revision r
$Cond^r_{not}$	a set of $cond_{not}$ in revision r
$Cond^r$	$Cond^r_{null} \cup Cond^r_{not}$
$Desc^r_{null}$	$Ret^r_{null} \cup Cond^r_{null}$
$Desc^r_{not}$	$Ret^r_{not} \cup Cond^r_{not}$
$Desc^r$	$Desc^r_{null} \cup Desc^r_{not}$
C	a set of all commits in a target project
R	a set of all revisions in a target project
loc^r	lines of code in revision r
$latest$	the latest revision in the specified term

A. Research Questions

We investigated the following research questions.

RQ1: Were ret_{null} and $cond_{null}$ modified more frequently than ret_{not} and $cond_{not}$?

RQ2: Did the size of projects affect the frequency of modifications to ret_{null} and $cond_{null}$?

RQ3: Did the development phases of the projects affect the frequency of modifications to ret_{null} and $cond_{null}$?

RQ4: Did the density of $cond_{null}$ increase as projects proceeded?

RQ1 is the question on whether statements including $null$ are, in fact, costly to maintain. In this research, "a statement is costly to maintain" means that the statement is modified frequently. RQ2 and RQ3 are questions on whether the frequency of modifications depends on the characteristics and development periods of projects. RQ4 is the question on how the density of $cond_{null}$ changes during software evolution. If the density of $cond_{null}$ increases as a project proceeds, the occurrences of $cond_{null}$, which is not the functionality we want to realize, are so many that $cond_{null}$ and ret_{null} are costly to maintain. ret_{null} is also costly because it is the cause of writing $cond_{null}$.

B. Target Projects

Table II shows the list of target projects. All the target projects are Java systems managed using git, and all had a high number of revisions and the size of each revision was not low. They were used in previous research [2], [3], [8], [9].

C. Experimental Method

In this experiment, the input was a target repository, and the outputs were the following:

- Ret^r_{null}, Ret^r_{not}, $Cond^r_{null}$, $Cond^r_{not}$, loc^r, for $\forall r \in R$
- ΔRet^c_{null}, ΔRet^c_{not}, $\Delta Cond^c_{null}$, $\Delta Cond^c_{null}$, for $\forall c \in C$

In order to collect the necessary data, we proceeded as follows. Figure 3 is an overview of the steps.

Step 1) We analyzed the source code in each revision to obtain the number of instances of $ret_{null/not}$ and $cond_{null/not}$ for each method. ret_{null} is only the literal sequence of tokens "return null," and $cond_{null}$ are conditional predicates comparing a variable to null, such as "a == null," and "null != a."

Step 2) By using the data obtained in Step 1, we found additions/deletions of $ret_{null/not}$ and $cond_{null/not}$. We regarded the changed number of instances of $ret_{null/not}$ and $cond_{null/not}$ in a method between two revisions to be additions/deletions, respectively. As we used the number of instances of $ret_{null/not}$ and $cond_{null/not}$, some changes were ignored in a case where some $ret_{null/not}$ were modified, but the number of instances of $ret_{null/not}$ in the method was not changed.

Step 3) We filtered out unnecessary additions/deletions. If $ret_{null/not}$ and $cond_{null/not}$ were added/deleted as a part of module additions/deletions, they were filtered out. This was done because the changes in numbers arising from adding/deleting modules are not commonly caused by bug fixes. In this experiment, when a method was added/deleted,

TABLE II: Target projects and their size

Project	loc^{latest}	$\|R\|$
ant	131,265	12,783
commons-io	25,031	1,526
eclipse.jdt.core	1,155,484	19,140
egit	92,305	3,126
jEdit	115,842	6,221
jboss-as	551,426	10,764
jetty	207,517	6,082
JGit	124,662	2,321
log4j	30,010	3,226
lucene-soir	537,150	8,026
maven	72,201	9,312
org.eclipse.cdt	1,029,497	21,157
org.eclipse.hudson.core	81,876	1,008
tomcat	240,086	9,172
Total	4,394,352	122,116

we considered it to be an addition/deletion of the modules. In summary, our filtering omitted the additions/deletions of $ret_{null/not}$ and $cond_{null/not}$ from the experimental target if they satisfied any of the following conditions: their number in their owner method was not changed, their owner method disappeared, or their owner method appeared anew at a given commit.

We obtained the remains of this filtering as necessary modifications. In the example shown in Figure 3, we obtained only modifications where the numbers were changed "0 to 1," "2 to 1," and "1 to 0."

D. Calculating the Frequency of Modifications

To answer the research questions, it is necessary to compare the frequencies of modifications between $ret_{null/not}$ and $cond_{null/not}$.

$f_{ret_{null}}(C)$, $f_{ret_{not}}(C)$, $f_{cond_{null}}(C)$, and $f_{cond_{not}}(C)$ are the frequency of modifications to ret_{null}, ret_{not}, $cond_{null}$, and $cond_{not}$. They were calculated as follows:

$$f_{ret_{null}}(C) = \frac{\sum_{c \in C} |\Delta Ret^{c}_{null}|}{|Ret^{latest}_{null}|} \quad (1)$$

$$f_{ret_{not}}(C) = \frac{\sum_{c \in C} |\Delta Ret^{c}_{not}|}{|Ret^{latest}_{not}|} \quad (2)$$

$$f_{cond_{null}}(C) = \frac{\sum_{c \in C} |\Delta Cond^{c}_{null}|}{|Cond^{latest}_{null}|} \quad (3)$$

$$f_{cond_{not}}(C) = \frac{\sum_{c \in C} |\Delta Cond^{c}_{not}|}{|Cond^{latest}_{not}|} \quad (4)$$

For example, in Eq. 1, the denominator is the total number of ret_{null} in the latest revision, and the numerator is the number of added/deleted ret_{null} throughout all the commits. The result is the frequency of modifications to one ret_{null}. By dividing $|Ret^{latest}_{null}|$, $|Ret^{latest}_{not}|$, $|Cond^{latest}_{null}|$, and $|Cond^{latest}_{nut}|$, we can reduce the influence arising from the difference in the number of modifications simply caused by the difference of their number.

E. Unit of Analysis

In this experiment, by using the caller/callee relationships of methods, we obtained the modifications to $cond_{null}$ in a caller method when ret_{null} was added to a callee method. As this research focuses on methods, $|\Delta Desc^{c}|$ for $\forall c \in C$ were calculated for each method.

F. Definition of Development Phases

We could not divide the development phases clearly because the experimental targets were open source projects. In this experiment, we considered a period between major version releases as one software development period, and we divided each of the development periods into an anterior half and a posterior half. Figure 4 shows the division of the periods between version 1.0.0 and version 1.2.0. We considered that modules had been added in the anterior half and bugs had been fixed in the posterior half. We performed such divisions, and calculated the frequency of modifications in the anterior half and the posterior half.

Division and statistical testing were conducted on 12 projects listed in Table II except for tomcat and eclipse.jdt.core. As the borders among major version releases were vague, we excepted those projects.

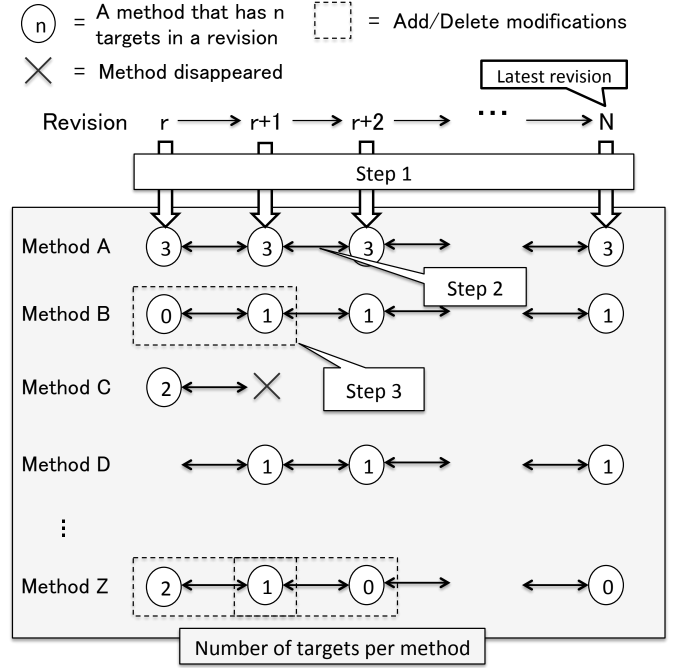

Fig. 3: Overview of steps

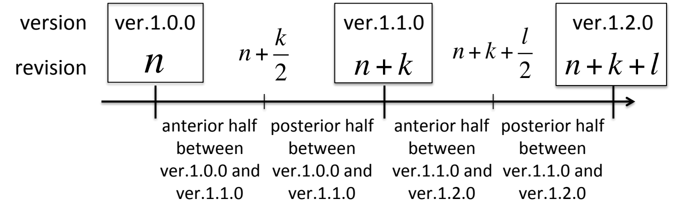

Fig. 4: Overview of division

G. Statistical Testing Methodology

In order to answer the research questions, we need to check whether two samples have a significant difference or a statistical correlation.

We used the Wilcoxon signed rank test [11] to determine whether two paired-samples had a significant difference. If the obtained p value is low, there is a low probability that the observed differences are accidental. In this experiment, the significance level is 1%. This means that there is a significant difference if $p \leq 0.01$, and there is no significant difference if $0.01 < p$.

Similarly, we calculated the Spearman's rank correlation coefficient (ρ value) to determine whether two paired-samples had a statistical correlation. After that, we calculated the p value corresponding to the ρ value. The two samples have a positive correlation if the ρ value is positive, and a negative correlation if the ρ value is negative. The meaning of the p value is the same as for the Wilcoxon signed rank test. Two samples have no correlation if $|\rho| < 0.5$, and they have a correlation if $0.5 \leq |\rho|$.

IV. Answers to Research Questions

A. Answer to RQ1

To ascertain whether ret_{null} and $cond_{null}$ were modified more frequently than ret_{not} and $cond_{not}$, we tested the difference between them. Figure 5 shows a comparison between $f_{ret_{null}}(C)$ and $f_{ret_{not}}(C)$, $f_{cond_{null}}(C)$ and $f_{cond_{not}}(C)$. In the figure, black regions are $Desc_{null}$, and gray regions are $Desc_{not}$. In addition, the figure does not show a value just as it is, but the percentage for the value. For example, the percentage of $f_{ret_{null}}(C)$ is calculated as follows: $f_{ret_{not}}(C)$, $f_{cond_{null}}(C)$, and $f_{cond_{not}}(C)$ are the same as $f_{ret_{null}}(C)$.

$$\frac{f_{ret_{null}}(C)}{f_{ret_{null}}(C) + f_{ret_{not}}(C)} \tag{5}$$

The p value obtained from the Wilcoxon signed rank test between $f_{ret_{null}}(C)$ and $f_{ret_{not}}(C)$ is 2.44×10^{-4} (Figure 5(a)). That means $f_{ret_{null}}(C) - f_{ret_{not}}(C)$ in almost all the projects are positive values. As $p \leq 0.01$, ret_{null} was modified more frequently than ret_{not}. On the other hand, the p value between $f_{cond_{null}}(C)$ and $f_{cond_{not}}(C)$ is 0.714 (Figure 5(b)). Thus, $cond_{null}$ were not modified more frequently than $cond_{not}$.

Therefore, our answer to RQ1 is that ret_{null} was modified more frequently than ret_{not}, and $cond_{null}$ was not modified more frequently than $cond_{not}$.

B. Answer to RQ2

To ascertain whether the size of projects and the frequency of modifications have any correlation, we hypothesized that loc^{latest} indicated the size of projects, and tested the correlation between them.

Figure 6 shows the scattergram between $f_{ret_{null}}(C)$, $f_{cond_{null}}(C)$ and loc^{latest}. The x-axis is loc^{latest}, and the y-axis is $f_{ret_{null}}(C)$ and $f_{cond_{null}}(C)$. We calculated that the Spearman's rank correlation coefficient (ρ value) and the p value corresponds to the ρ value.

The result from the Spearman's rank correlation test between $f_{ret_{null}}(C)$ and loc^{latest} was $\rho = -0.0330$, and $p = 0.916$ (Figure 6(a)). As $0.01 < p$, there was no significant correlation. Similarly, the result between $f_{cond_{null}}(C)$ and loc^{latest} was $\rho = 0.169$, and $p = 0.563$ (Figure 6(b)). As it was also $0.01 < p$, there was no significant correlation.

As a result, our answer to RQ2 is that the size of the projects did not have a significant effect on the frequency of modifications to ret_{null} and $cond_{null}$.

C. Answer to RQ3

There were 109 major versions among the target projects. After dividing them, we calculated $f_{ret_{null}}(C)$ and $f_{cond_{null}}(C)$ for the divided 218 periods. Figure 7 shows the calculated values.

We tested $f_{ret_{null}}(C)$ and $f_{cond_{null}}(C)$ for differences between the anterior half and the posterior half. The result from the Wilcoxon signed rank test between $f_{ret_{null}}(C)$ in the anterior half and the posterior half was $p = 0.126$ (Figure 7(a)). As $0.01 < p$, there was no significant difference. Similarly, the result between $f_{cond_{null}}(C)$ in the anterior half and the posterior half was $p = 0.383$ (Figure 7(b)). There was also no significant difference.

As a result, our answer to RQ3 is that the development phases did not have a significant effect on the frequency of modifications on ret_{null} and $cond_{null}$.

D. Answer to RQ4

We calculated $density(Cond^{r}_{null})$, representing the number of $cond_{null}$ per line of code, for each revision r in the target projects. The value was calculated as follows:

$$density_{cond_{null}}(r) = \frac{|Cond^{r}_{null}|}{loc^{r}} \tag{6}$$

For each project, we tested whether there was a positive correlation between the revision number (r) and $density_{cond_{null}}(r)$. A positive correlation means that the density of $cond_{null}$ is increasing as a project proceeds. Table III shows the test result. The ρ value indicates the strength of the correlation, and the p value indicates whether there is a significant correlation or not.

As a result, four projects showed significant positive correlations, five projects showed significant negative ones, and three projects had no correlation. Two projects did not have significant correlation because the p value of two projects is $p > 0.01$. This shows there was no regular pattern in the results.

Consequently, our answer to the RQ4 is that the density of $cond_{null}$ did not increase as the projects proceeded.

TABLE III: Spearman's rank correlation coefficient (ρ value) and p value for revision number r and $density_{cond_{null}}(r)$

Software	ρ value	p value
ant	0.432	under 2.2×10^{-16} [a]
commons-io	0.733	under 2.2×10^{-16}
jdt.core	-0.974	under 2.2×10^{-16}
egit	0.223	under 2.2×10^{-16}
jEdit	0.668	under 2.2×10^{-16}
jboss-as	-0.765	under 2.2×10^{-16}
jetty	-0.992	under 2.2×10^{-16}
jgit	0.598	under 2.2×10^{-16}
log4j	0.006	0.740
lucene-solr	-0.947	under 2.2×10^{-16}
maven	0.934	under 2.2×10^{-16}
cdt	0.156	under 2.2×10^{-16}
hudson.core	-0.151	2.898×10^{-4}
tomcat	-0.986	under 2.2×10^{-16}

[a] This indicates the p value is too small to calculate a precise value.

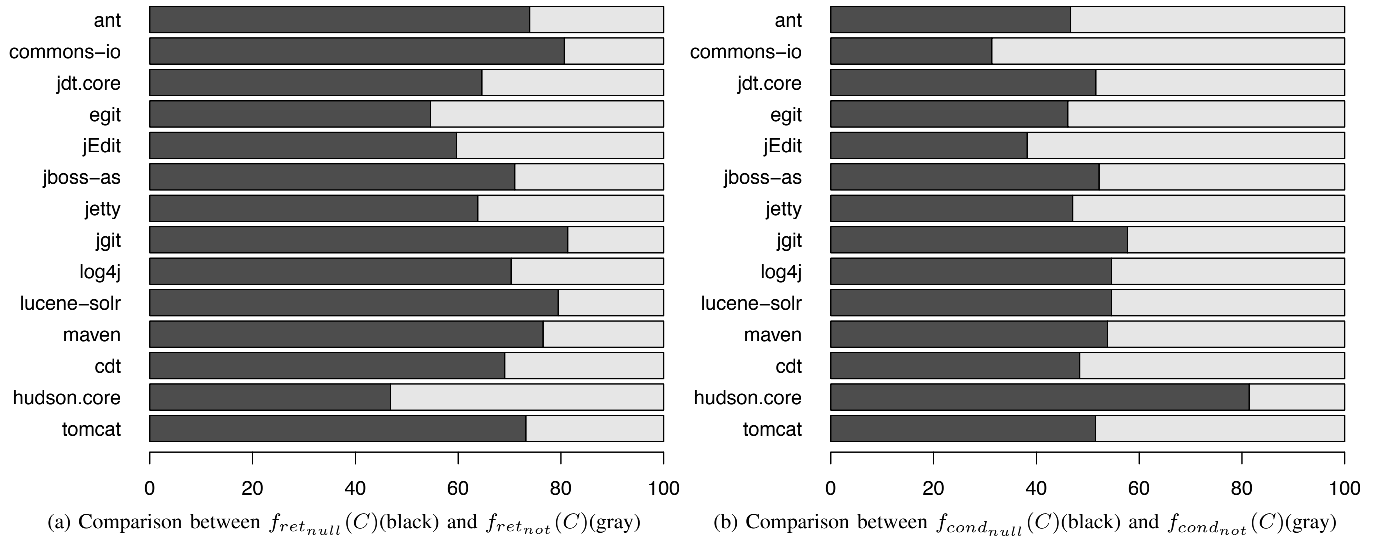

(a) Comparison between $f_{ret_{null}}(C)$(black) and $f_{ret_{not}}(C)$(gray)

(b) Comparison between $f_{cond_{null}}(C)$(black) and $f_{cond_{not}}(C)$(gray)

Fig. 5: Comparison between $f_{Desc_{null}}(C)$(black) and $f_{Desc_{not}}(C)$(gray)

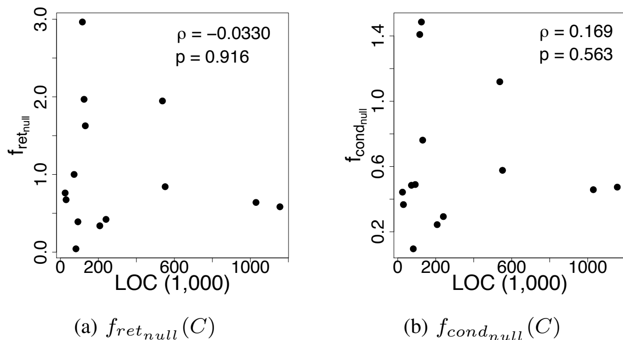

(a) $f_{ret_{null}}(C)$

(b) $f_{cond_{null}}(C)$

Fig. 6: Scattergrams of loc^{latest} and $f_{ret_{null}}(C)$, $f_{cond_{null}}(C)$

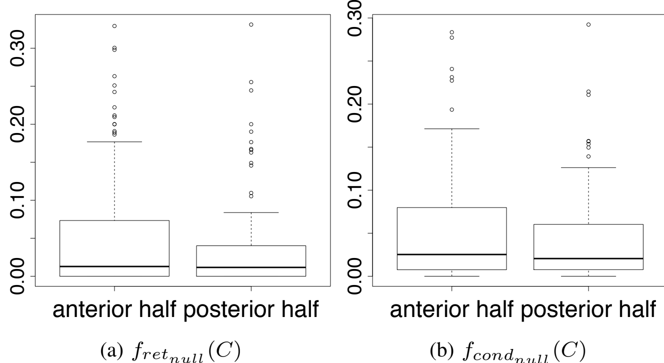

(a) $f_{ret_{null}}(C)$

(b) $f_{cond_{null}}(C)$

Fig. 7: Box plot showing $f_{ret_{null}}(C)$ and $f_{cond_{null}}(C)$ of the anterior half and the posterior half

V. Discussion

A. Discussion on the Results of RQ1, RQ2, and RQ3

Our answer to RQ1 is that ret_{null} was modified more frequently than ret_{not}. As ret_{null} was frequently modified, we can say that their presence is probably costly to maintain. Our answers to RQ2 and RQ3 are that the size and the development phases of projects did not affect the frequency of modifications to ret_{null} and $cond_{null}$. This means a developer modified ret_{null} and $cond_{null}$ throughout the entire development period of a project. Our answers above show that ret_{null} affects software maintenance throughout the entire development period of a project.

B. Discussion on RQ4

In our answer to RQ4 (see IV-D), we showed that many projects had positive or negative correlations between revision number r and $density_{cond_{null}}(r)$. We conducted additional experiments to find common characteristics in the density on $cond_{null}$.

Figure 8 shows the relationship between revision numbers and $density_{cond_{null}}(r)$ in four projects. In Figure 8, the x-axis represents the revision number, and the y-axis represents $density(Cond^{x}_{null})$. These projects show that the continuation of the change has been maintained. This means that the density will have been decreasing/increasing, and it will become a very low/high value. However, inspecting the overall $density_{cond_{null}}(r)$ gives us another idea.

Figure 9 shows $density_{cond_{null}}(r)$ of $\forall r \in R$ in all the projects. Position x in the x-axis means the revision in the top x% of revisions sorted by ascending order of revision numbers. The y-axis is $density_{cond_{null}}(x)$. The thick bar shows the median of all the projects. This graph shows that $density_{cond_{null}}(r)$ of all the projects is in the range 0.01 to 0.04, and the median is stable regardless of the progress of their projects. In other words, the number of $cond_{null}$ is from one to four lines per 100 lines, and this density will be kept in the future of any project. This result also means that countermeasures to reduce ret_{null} and $cond_{null}$ have not been taken in all the projects. After we see Figure 9, we can recognize that the density reaches the area between 0.01 and

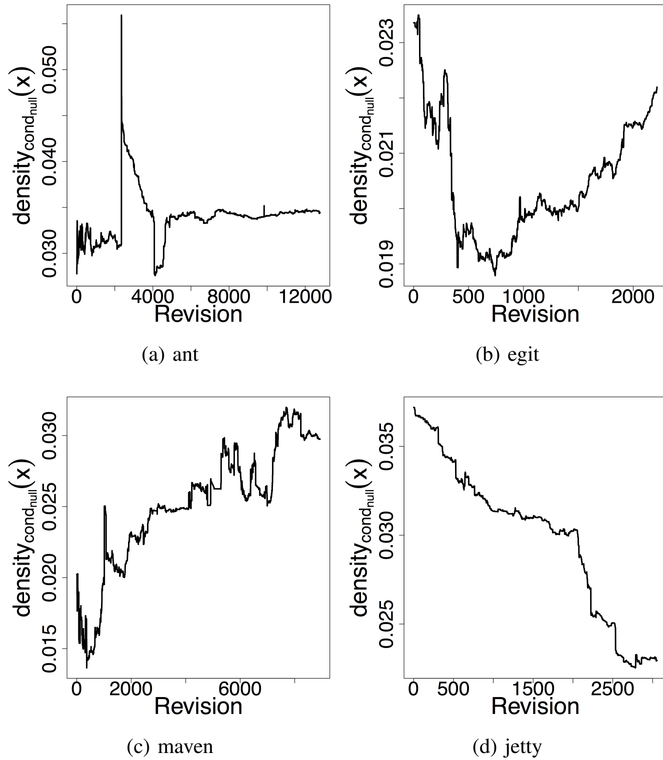

(a) ant

(b) egit

(c) maven

(d) jetty

Fig. 8: Change in frequency of modifications to $cond_{null}$ in some projects. Tendencies of $cond_{null}$ differ in each project

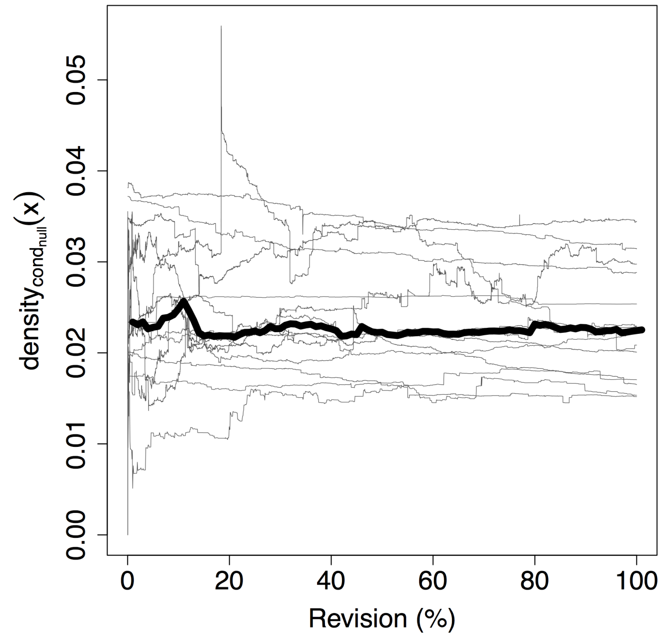

Fig. 9: $density_{cond_{null}}(r)\ \forall r \in R$ of all projects. This figure shows that $density_{cond_{null}}(latest)$ of all projects are between 0.01 and 0.04

```
1     } else {
2  -    return null;
3  +    if (merger.failedAbnormally())
4  +      return new CherryPickResult(merger.getFailingPaths
          ());
5  +    // merge conflicts
6  +    return CherryPickResult.CONFLICT;
7     }
```

Fig. 10: Example of replacing ret_{null} into an error constant

0.04 in all projects as shown in Figure 8. Once a density reaches the area, it is unlikely that the density will go out from the area. This is presumably a common characteristic of the density in $cond_{null}$.

In addition, if $density_{cond_{null}}(latest)$ in a project is high ($>$ 0.04), the project presumably has some problems. This means that the number of $cond_{null}$, which is not a functionality that we want to realize, has increased too much. Thus, if $density_{cond_{null}}(latest)$ is high, we suggest that the developer should refactor (e.g., check the existence of unnecessary $cond_{null}$, returning suitable values instead of $null$, and remove $cond_{null}$ by using a NullObject pattern).

C. Discussion on Modifications to ret_{null} and $cond_{null}$

We select modifications from $\Delta Desc^c$ in $\forall c \in C$ and we discuss the effect on the density of $cond_{null}$. In this study, we discuss the modifications to ret_{null} and $cond_{null}$ in JGit. The reason we targeted JGit is that the size was not too large, meaning we could check all the modifications, in addition the JGit project recorded information on bugs using a bug control system, Bugzilla [12]. We show the characteristic modifications and discuss them.

Figure 10 shows an example of replacing ret_{null} into an error constant. This is a representative example of replacing ret_{null} with other mechanisms. Such a replacement often causes software design problems because the return value is changed. It is difficult to perform such a replacement without causing a problem in a fully automated way.

Although such a replacement is necessary to identify the errors that occurred in the same method, performing such a replacement is difficult. Therefore, the presence of ret_{null} is costly to maintain.

Figure 11 shows the modification occurred in a commit whose commit message is "[findBugs] Don't pass $null$ for non-null parameter in RebaseCommand." This commit message shows that developers detected a null dereference by using FindBugs, and fixed the null dereference by adding $cond_{null}$. Many instances of null dereference occurred because developers had forgotten to write $cond_{null}$, but developers could easily detect this by using null dereference detection tools.

In Figure 12, as there was no $cond_{null}$ for the variable "objectID," developers added $cond_{null}$ and the error handling when "objectID" was $null$. The same problem was fixed in two places, but the developers forgot to fix all of them. In such cases, null dereference detection tools are effective. In addition, we need to fix bugs in multiple places not only when a developer forgets to write $cond_{null}$, but also when a developer adds ret_{null} to existing methods.

As described above, developers can detect null dereferences by using null dereference detection tools and fix them. On the

```
1    DirCacheCheckout dco;
2  - RevCommit commit = walk.parseCommit(repo.resolve(
        commitId));
3  + if (commitId == null)
4  +   throw new JGitInternalException(
5  +     JGitText.get().abortingRebaseFailedNoOrigHead);
6  + ObjectId id = repo.resolve(commitId);
7  + RevCommit commit = walk.parseCommit(id);
8    if (result.getStatus().equals(Status.FAILED)) {
9      RevCommit head = walk.parseCommit(repo.resolve(
        Constants.HEAD));
10     dco = new DirCacheCheckout(repo, head.getTree(),
```

Fig. 11: Example showing lack of $cond_{null}$

```
1  @@ -108,7 +110,10 @@
2    ObjectId objectId = repository.resolve(revstr);
3  - tree = new RevWalk(repository).parseTree(objectId);
4  + if (objectId != null)
5  +   tree = new RevWalk(repository).parseTree(objectId);
6  + else
7  +   tree = null;
8    this.initialWorkingTreeIterator = workingTreeIterator;
9
10 @@ -125,9 +130,13 @@
11   this.repository = repository;
12 - tree = new RevWalk(repository).parseTree(objectId);
13 + if (objectId != null)
14 +   tree = new RevWalk(repository).parseTree(objectId);
15 + else
16 +   tree = null;
17   this.initialWorkingTreeIterator = workingTreeIterator;
```

Fig. 12: Example of $cond_{null}$ being fixed in many places

other hand, sometimes we can remove $null$ instead of simply adding $cond_{null}$. Figure 13 shows refactoring that removes $cond_{null}$ by assigning a suitable value "repo.lockDirCache()" to variable "dc" as an initialization. However, it is difficult to perform such refactorings because we need to analyze all expressions that include a variable so as to be able to assign a suitable value to a variable. In addition, there is no popular tool that assigns suitable values to variables automatically. Therefore, the number of instances of $cond_{null}$ rarely decreased. We can regard such a change as a factor of maintenance stem from $null$. Our experiment took such a change into consideration because our experiment did not focus on only addition of $cond_{null}$.

In summary, there were ret_{null} and $cond_{null}$ replaced with other mechanisms in order to avoid bad effects of ret_{null}. We can detect null dereference automatically by using null dereference detection tools, and developers add $cond_{null}$ when null dereference is detected. However, as it requires complex analyses, reducing the number of instances of ret_{null} and $cond_{null}$ is difficult both when using manual methods and when using automatic methods.

D. Replacing Return Null with Other Mechanisms

Exception handling is one of the mechanisms for replacing ret_{null}. We measured the frequency of modifications on throw statements and try statements. Then, we compared them with $f_{ret_{null}}(C)$ and $f_{cond_{null}}(C)$. Figure 15 shows the comparison result. The result of the Wilcoxon signed rank test between $f_{ret_{null}}(C)$ and throw statements was $p = 0.0245$. As $0.01 < p$, there was no significant difference. Similarly, the result

```
1    private void resetIndex(RevCommit commit) throws
        IOException {
2  -   DirCache dc = null;
3  +   DirCache dc = repo.lockDirCache();
4      try {
5  -     dc = repo.lockDirCache();
6        dc.clear();
7        DirCacheBuilder dcb = dc.builder();
8        dcb.addTree(new byte[0], 0, repo.newObjectReader
        (), commit.getTree());
9        dcb.commit();
10 -   } catch (IOException e) {
11 -     throw e;
12     } finally {
13 -     if (dc != null)
14 -       dc.unlock();
15 +     dc.unlock();
16     }
17   }
```

Fig. 13: Example of removing $cond_{null}$ by modifying the default value

```
1  Object referencedObject = getTaskContext().getDataValue(
        reference);
2  if (referencedObject == null) {
3    throw new ExecutionException("Unable to locate the
        reference specified by refid '" + getReference() +
        "'");
4  }
5  if (!this.getClass().isAssignableFrom(referencedObject.
        getClass())) {
6    throw new ExecutionException("The object referenced by
        refid '" + getReference() + "' is not compatible
        with this element ");
7  }
```

Fig. 14: Example of throw statements being written in multiple places

between $f_{ret_{null}}(C)$ and try statements was $p = 0.390$. As it was also $0.01 < p$, there was no significant difference.

However, developers sometimes write multiple throw statements, all of which correspond to a single ret_{null}. This is because throw statements have error messages and exception types. Developers can change error messages and exception types depending on the reasons of errors (Figure 14 shows an example of such a situation). For this reason, the frequency of the modifications on throw statements became a larger value than the actual one. Therefore, it is not fair to compare the frequencies of the modifications between ret_{null} and exception handling. Exception handling is better than ret_{null} from the qualitative points of view, such as the examples in Section V-C, and the characteristic that ret_{null} have no error message.

Replacing ret_{null} and $cond_{null}$ is related to research that is intended to detect null dereference, because detecting null dereference is the same as detecting passing $null$ from one method to another method and a variable can be assigned $null$. The information as to which variables can be assigned $null$ is very useful when it comes to replacing ret_{null} and $cond_{null}$. Therefore, the results of this research reinforce the motivation of those studies, e.g., "As null dereference occurs frequently, it is useful to detect them," with the motivation "Supporting replacing ret_{null} and $cond_{null}$ that affect software maintenance."

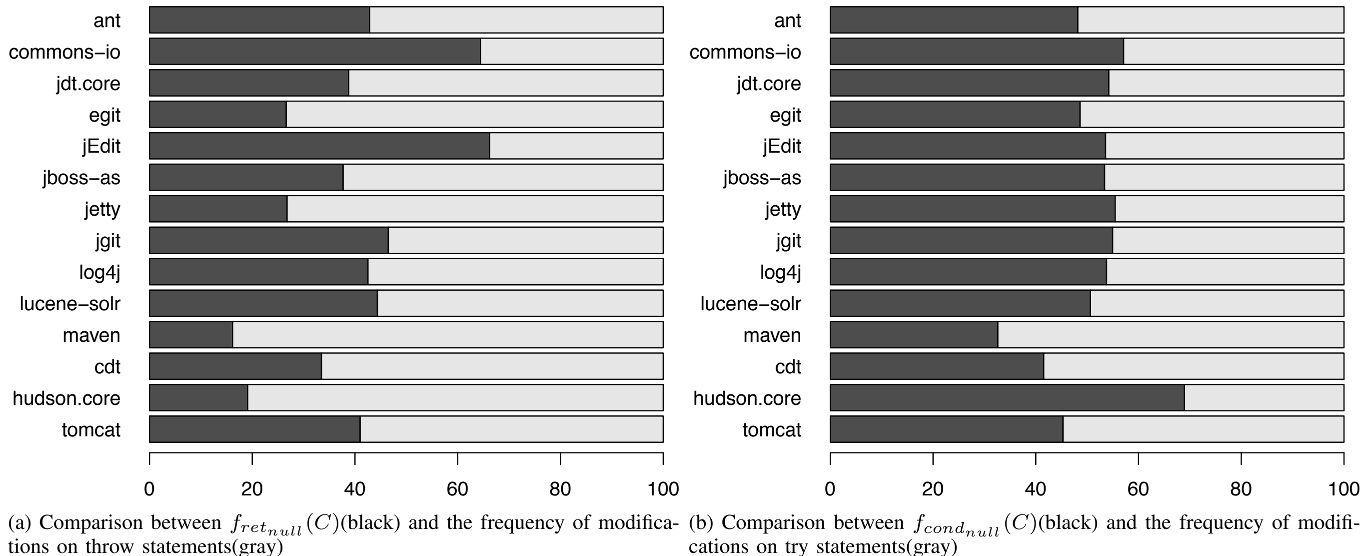

(a) Comparison between $f_{ret_{null}}(C)$(black) and the frequency of modifications on throw statements(gray)

(b) Comparison between $f_{cond_{null}}(C)$(black) and the frequency of modifications on try statements(gray)

Fig. 15: Comparison between $f_{Desc_{null}}(C)$(black) and the frequency of modifications on exception handling mechanisms(gray)

VI. Threats to Validity

The fairness of comparison on RQ1: As stated in Section V-D, it is not appropriate to compare ret_{null} with throw statements, because the frequency of modifications on throw statements was larger than the actual one. The unfairness also happened to ret_{not}. Although $f_{ret_{not}}(C)$ was presumably larger than the actual one, $f_{ret_{null}}(C)$ was larger than $f_{ret_{not}}(C)$ under adverse conditions (as shown in Section IV-A). Thus, our experimental results were not altered by the unfairness.

Code outside the methods: In this experiment, $\Delta Desc^c$ for $\forall c \in C$ were calculated for all methods. This calculation had the issue that $ret_{null/not}$ and $cond_{null/not}$ were ignored if $ret_{null/not}$ and $cond_{null/not}$ were outside the methods. However, $ret_{null/not}$ and $cond_{null/not}$ were rarely written outside the methods because $ret_{null/not}$ and $cond_{null/not}$ were return statements or conditional predicates. Thus, we excluded from consideration $ret_{null/not}$ and $cond_{null/not}$ that were outside the methods.

Changed/Moved code elements: $\Delta Desc^c$ for $c \in C$ consisted only of additions/deletions. Thus, we did not detect "Change" modifications, for example, a variable compared to $null$ had been changed to another variable. This may have an effect on the experimental result and the discussion. In addition, as our experimental steps could not detect a code move of $ret_{null/not}$ and $cond_{null/not}$, a code move caused by refactoring, such as Extract Method, was recognized as modifications of $ret_{null/not}$ and $cond_{null/not}$. A code move caused by refactoring should be filtered out, similarly to an addition/deletion of modules. However, in this experiment, we considered that the density of code moves was low, and ignored them.

Variables whose values are null: ret_{null} is only a return statement whose operand is $null$. If a return statement has a variable whose value is $null$, the statement is not counted. The situation is the same for $cond_{null}$. For this reason, the number of ret_{null} and $cond_{null}$ could differ from the actual number.

The division of the development phases: In the comparison of the development phases, we divided the periods between major versions into two periods, and tested whether they had a significant difference. However, as this division was performed based on the number of revisions, we could not obtain a precise division of the development phases. As a result, the experimental result could differ from the result when the periods were divided more precisely.

VII. Related Work

Many previous studies have tackled null dereference problems from many viewpoints. Currently, there are many tools to detect null dereference automatically, which were introduced in Section II. Those tools target existing null dereference.

By contrast, there are techniques that focus on preventing null dereference from occurring in compile time. One such technique is that developers add the constraint "null value is never assigned to this variable" (non-null) to some variables. This technique enables compilers to verify whether the variables could be $null$. As a result, these techniques can reduce mistakes in which a $null$ value is unintentionally assigned to some variables. In formal methods, developers represent a program as a formal specification using formal specification languages, such as JML (Java Modeling Language) [13], and check whether the constraints are violated or not by performing verifications. Developers can obtain proof that the variable will never become $null$ by using these methods. However, writing formal specifications needs a great deal of effort on the part of developers. Therefore, researchers tend to implement non-null features as the preferred type system and annotations. Fähndrich et al. proposed a NonNull type of system and implemented it on C# as an annotation [14]. Papi

et al. extended the Java-type system and enabled checks to be made as to whether variables annotated $@NonNull$ can be $null$ or not [15]. There are many researchers that are trying to make an environment in which non-null can be used in a practical environment.

There also are methods that enforce null check on variables that can be $null$. Hovemeyer et al. developed $@CheckForNull$ annotation, and proposed a method that enforces null check on variables on Java [5]. Haskell [16] and Scala [17] have variables of *Maybe type* or *Option type* that represent variables that can be $null$. Variables of these types contain actual values or *Nothing* values, which are equivalent to a $null$ value. When developers use such variables, they spontaneously handle cases in which the variables have *Nothing* values. Such types of systems make developers aware of exception handling, and can detect the lack of a null check in compile time. As a result, such systems prevent null dereference from occurring.

In addition, $null$ as an initializer is a subject of research, because such an initializer is another factor in the occurring of null dereference. By using initialization analysis, which analyzes whether variable initializations have been done or not, developers can know which variables are uninitialized [18]. This decreases the possibility of null dereference occurring. In addition, there are techniques that track the origins of $null$ values [19]. As described above, many approaches have been proposed.

VIII. Conclusion

In this paper, we investigated whether, and to what extent, the presence of return null and null check is costly to maintain by mining modifications during software evolution. This research focused on return null and null check, and an experiment conducted on 14 projects showed that return null were modified more frequently than statements that do not have $null$. In addition, we suggested that developers should avoid writing return null by reason of our qualitative analysis. Moreover, it was found that the frequencies of modifications on return null and null check bore no relation to the size and the development phases of projects. The density of null checks was from one to four in 100 lines, in each project. If the density of null check in a project exceeds 0.04, we can say that developers should refactor in order to decrease the density. A developer could replace return null with other mechanisms, for example, exception handling, a proper object such as an empty array, or a NullObject pattern.

We plan to research the reasons for writing return null and null check, and see to what extent we can lower the frequency of modifications. We can focus our experiment on "addition of null check." Results of such an experiment will be worth discussing. In addition, it is our future work to support developers by replacing return null and null check automatically. We also plan to find a method in order to compare return null with other error handling mechanisms quantitatively, and conduct the same experiment for other programming languages such as C.

Acknowledgment

This study was supported by Grants-in-Aid for Scientific Research (S) (25220003), Grant-in-Aid for Exploratory Research (24650011) from the Japan Society for the Promotion of Science, and Grant-in-Aid for Young Scientists (A) (24680002) from the Ministry of Education, Culture, Sports, Science and Technology.

References

[1] T. Hoare, "Historically Bad Ideas: "Null References: The Billion Dollar Mistake," in *QCon*, 2009.
[2] M. G. Nanda and S. Sinha, "Accurate Interprocedural Null-Dereference Analysis for Java," in *ICSE*, May 2009, pp. 16–24.
[3] N. Ayewah and W. Pugh, "Null Dereference Analysis in Practice," in *PASTE*, Jun. 2010, pp. 65–72.
[4] W. R. Bush, J. D. Pincus, and D. J. Sielaff, "A static analyzer for finding dynamic programming errors," *Software: Practice and Experience*, vol. 30, no. 7, pp. 775–802, Jun. 2000.
[5] D. Hovemeyer, J. Spacco, and W. Pugh, "Evaluating and tuning a static analysis to find null pointer bugs," in *PASTE*, Sep. 2005, pp. 13–19.
[6] A. Loginov, E. Yahav, S. Chandra, S. Fink, N. Rinetzky, and M. Nanda, "Verifying Dereference Safety via Expanding-Scope Analysis," in *ISSTA*, Jul. 2008, pp. 213–224.
[7] D. Hovemeyer and W. Pugh, "Finding More Null Pointer Bugs, But Not Too Many," in *PASTE*, Jun. 2007, pp. 9–14.
[8] M. G. Nanda, M. Gupta, S. Sinha, S. Chandra, D. Schmidt, and P. Balachandran, "Making Defect-Finding Tools Work for You," in *ICSE*, May 2010, pp. 99–108.
[9] D. Hovemeyer and W. Pugh, "Finding Bugs is Easy," in *OOPSLA*, Oct. 2004, pp. 24–28.
[10] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata, "Extended Static Checking for Java," in *PLDI*, Jun. 2002, pp. 17–19.
[11] M. Hollander and D. A. Wolfe, *Nonparametric Statistical Methods, 2nd Edition*. Wiley-Interscience, 1999.
[12] Bugzilla, "http://www.bugzilla.org/."
[13] G. T. Leavens, A. L. Baker, and C. Ruby, "Preliminary Design of JML: A Behavioral Interface Specification Language for Java," *SIGSOFT Software Engineering Notes*, vol. 31, no. 3, pp. 1–38, May 2006.
[14] M. Fähndrich and K. R. M. Leino, "Declaring and checking non-null types in an object-oriented language," in *OOPSLA*, Oct. 2003, pp. 302–312.
[15] M. M. Papi and M. D. Ernst, "Compile-Time Type-Checking for Custom Type Qualifiers in Java," in *OOPSLA*, Oct. 2007, pp. 809–810.
[16] S. P. Jones, J. Hughes, L. Augustsson, D. Barton, B. Boutel, W. Burton, J. Fasel, K. Hammond, R. Hinze, P. Hudak, T. Johnsson, M. Jones, J. Launchbury, E. Meijer, J. Peterson, A. Reid, C. Runciman, and P. Wadler, "The Haskell 98 Language Report," Dec. 2002, http://www.haskell.org/onlinereport/.
[17] M. Odersky, "The Scala Language Specication Version 2.9," May 2011, http://www.scala-lang.org/docu/files/ScalaReference.pdf.
[18] F. Spoto and M. D. Ernst, "Inference of Field Initialization," in *ICSE*, May 2011, pp. 231–240.
[19] M. D. Bond, N. Nethercote, S. W. Kent, S. Z. Guyer, and K. S. McKinley, "Tracking Bad Apples: Reporting the Origin of Null and Undefined Value Errors," in *OOPSLA*, Oct. 2007, pp. 405–422.