

修士学位論文

題目

**Detection of License Inconsistencies
in Free and Open Source Software Projects**

指導教員

井上 克郎 教授

報告者

Yuhao Wu

平成 28 年 2 月 9 日

大阪大学 大学院情報科学研究科

コンピュータサイエンス専攻 ソフトウェア工学講座

内容梗概

Free and open source software (FOSS) plays an important role in source code reuse practice. They usually come with one or more software licenses written in the header part of source files, stating the requirements and conditions which should be followed when been reused.

Removing or modifying the license statement by re-distributors will result in the inconsistency of license with its ancestor, and may potentially cause license infringement. In this paper, we describe and categorize different types of license inconsistencies and propose a method to detect them. Then we apply this method to Debian 7.5 and a collection of 10,514 Java projects on GitHub and present the license inconsistencies found in these systems.

With a manual analysis, we summarized various reasons behind these license inconsistencies, some of which imply license infringement and require the attention from the developers. This analysis also exposes the difficulty to discover license infringements, highlighting the usefulness of finding and maintaining source code provenance.

主な用語

Software license

Code clone

License inconsistency

目次

1	Introduction	4
2	License Inconsistencies	7
2.1	Definition	9
2.2	Example	9
2.3	Categorization	10
3	Study Approach	12
3.1	License Inconsistency Metrics	13
3.2	Method of Detecting License Inconsistencies	14
3.3	Example	15
4	Empirical Study	17
4.1	Empirical Study on Debian 7.5	17
4.1.1	Results	17
4.1.2	Manual Analysis	22
4.2	Empirical Study on Java Projects	25
4.2.1	Results	25
4.2.2	Manual Analysis	26
5	Discussion	29
5.1	Discussion of the Results	29
5.2	Answering RQs	30
5.3	Improvement of the Method	31
5.3.1	Debian 7.5	31
5.3.2	Java Projects	32
5.4	What appears to be a copy might not be a copy	32
5.5	Threats to Validity	33
6	Related Work	35
7	Conclusion and Future Work	36
	謝辭	37

1 Introduction

Software reuse has long been advocated as a good practice to reduce development time and increase product quality [19, 22, 3]. The popularity of Free and Open Source Software (FOSS) has made software reuse a common practice. FOSS software can be defined as software that is licensed under a free and open source license. In a nutshell, a free and open source license allows the software to be freely used, modified, and redistributed (in modified or unmodified form) by anyone, as long as the conditions of its license are satisfied. The Open Source Initiative (OSI) has defined a set of characteristics that an open source license should have, and published a list of approved open source licenses¹. The Free Software Foundation defines a set of similar conditions that a license should satisfy in order to be considered a free software license².

Developers who reuse FOSS should pay special attention to the license under which a source file is made available, and make sure that they satisfy the conditions and limitations of its license. Otherwise they risk losing the right to reuse the software. Typically, the license of a file is located in the initial part of the file. We will refer to this area of the file as the *license statement* of the file.

The license of a file can only be changed by its copyright owner. In some special cases, the license terms allow others to change the license of the file. Otherwise, if the license is changed there is the potential for copyright infringement. For example in a case of XimpleWare Corp v. Versata Software Inc. et al³, Versata was sued for including GPL-licensed code into one of its products but removing the copyright and use notices required by GPL. This case was settled out of court in favor of Ximpleware.

For the purpose of this paper, we are interested in the situation where a copy of a file has a different license than the original file. If the new license has not been approved by the copyright owner we are confronted with a potential *license violation*. However, in many cases it is not clear whether the change in license has been approved by the copyright owner. For example, the copyright owner might have approved, via direct communication, a change in license. Under this scenario, the copy has a different license than the original, but it is not a license violation. For this reason, when a copy of a file has a different license than the original, we say that there is a *license inconsistency* between the licenses of the two files. Some license inconsistencies might turn to be license violations.

¹<https://opensource.org/licenses/alphabetical>

²<http://www.gnu.org/licenses/license-list.html>

³<http://www.ifross.org/en/artikel/versata-saga-settled-prejudice-1>

Anyone who wants to reuse FOSS software should concern that the software being reused is properly licensed. If the reused software contains files that have been copied from other sources, and these files have license inconsistencies, then it is important to resolve these inconsistencies. Otherwise the reuser of these files might be involved in legal disputes with the original copyright owner.

Previous study by Li et al. [16] shows that 36% of the developers who reused the OSS components changed the source code, but they did not point out whether these changes involve the license statement. In our study, we focus on the license statement changes and the license inconsistencies introduced between the different copies of the files.

To the best of our knowledge, no research has been done to discover and study the characteristics of license inconsistencies in software reuse. For example, how many types of license inconsistency are there? Do they exist in open source projects? If so, what is the proportion of each type? What caused these license inconsistencies?

Based on these questions, we set our research question as follows:

- **RQ1** *How can we categorize license inconsistencies?*
- **RQ2** *Do license inconsistencies exist in open source projects?*
- **RQ3** *What is the proportion of each type of license inconsistency?*
- **RQ4** *What caused these license inconsistencies? Are they potential license violations?*

The contributions of this paper are:

1. We describe and categorize different types of license inconsistencies.
2. We propose a method to detect license inconsistencies in large collections of open source projects, which can show the existence and number of each type of license inconsistency inside these projects.
3. We perform an empirical evaluation on our method using two sets of FOSS projects. This study reveals that license inconsistencies exist and proved the feasibility of our method.
4. We perform a manual analysis of some license inconsistency cases to understand the reasons behind them. We then summarized these reasons into 4 categories. Among them, two categories indicate license problems and require the developers' attention.

This paper is organized as follows. Section 2 describes background on FOSS licenses and license inconsistencies. Section 3 introduces our research method. An empirical study that uses this method is described in Section 4, followed by Section 5 with a discussion of the results. After a description of related work in Section 6, Section 7 concludes this paper and points out the future direction.

2 License Inconsistencies

A software license is a permission to reproduce, modify and redistribute a software, usually granted under certain conditions. An open source license is a software license that follows Open Source Definition⁴ and is approved by Open Source Initiative. As of today, only 70 licenses have been approved as Open Source License. However Black Duck Software claims that the Black Duck Knowledge Base includes over 2200 licenses⁵. Some licenses have been grouped under the same name as different versions. For example, the General Public License (GPL⁶) has versions 1, 2 and 3. Each version is, in legal terms, a totally independent license.

To reuse OSS source code files, developers must identify the license under which the files are made available, understand their terms, and satisfy their requirements. This is not a trivial task because one open source license does not usually allow easy integration with software under another license (German et al. gave a detailed discussion on this issue [6]). For example, software under the Apache Public License version 2 (APL-2.0) can be reused and integrated into software licensed under the GPL-3.0. On the other hand, software under the GPL-2.0 cannot be combined with software under the GPL-3.0 (however software under the GPL-2.0+, that is version 2 or any later version of the GPL, can be). Therefore, developers must know the licenses of files they reuse in order to avoid license violations.

It is also known that very frequently, the source code files in an application are under different licenses [18, 17]. In addition, copies of the same file might have different license because the copyright owner has licensed the file accordingly. For example, the copyright owner has decided to change the license from one version of the software to the other (even if the software did not have any changes).

Confusion can arise when a developer wishing to reuse a given file finds that two or more copies of it have different licenses. Let us assume that a developer wants to reuse two copies of the same file (not necessarily identical, due to their own evolution). The first copy, copy A, has license L_A , and copy B has license L_B . If the files both came directly from the copyright owner, then it can be assumed that both files have valid licenses; but if the

⁴<http://opensource.org/definition>

⁵<http://www.blackducksoftware.com/products/knowledgebase>

⁶In this paper we will use the abbreviations of FOSS licenses of the Software Package Data Exchange (SPDX), found at <http://spdx.org/licenses/>.

files came from third parties, one has to question if such parties have modified the licenses without the approval of the copyright owner (resulting in a potential license violation).

Usually, the license of an open source file is indicated in its license statement, found in the first comments of each source file. Here is an example of a license statement taken from `getopt.c` file in GNU library, which states that the file is under the GPL-3.0+:

```
/* Getopt for GNU.
 * NOTE: getopt is part of the C library, so if you don't
 * know what "Keep this file name-space clean" means, talk
 * to drepper@gnu.org before changing it!
 * Copyright (C) 1987-1996, 1998-2004, 2006, 2008-2012 Free
 * Software Foundation, Inc.
 * This file is part of the GNU C Library.
 *
 * This program is free software: you can redistribute it
 * and/or modify it under the terms of the GNU General Public
 * License as published by the Free Software Foundation;
 * either version 3 of the License, or (at your option) any
 * later version.
 *
 * This program is distributed in the hope that it will be
 * useful, but WITHOUT ANY WARRANTY; without even the implied
 * warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR
 * PURPOSE. See the GNU General Public License for more
 * details.
 *
 * You should have received a copy of the GNU General Public
 * License along with this program.
 * If not, see <http://www.gnu.org/licenses/>.
 */
```

Generally, the license statement of a source file can only be modified by its copyright owner. Reusers shall never modify the license statement unless it is under the permission of the copyright owner or allowed by the terms of the license.⁷ Otherwise, the reusers may incur a license violation.

In order to identify potential license violations, the first step is to identify license inconsistencies between files of different projects. In the following subsections, we introduce our definition of license inconsistency and give an examples of license inconsistencies we

⁷Some licenses, such as the Mozilla tri-license (which allowed the reuse of the file under either the MPL-1.0, the GPL-2.0+ or the LGPL-2.1) allows the user to remove one or two licenses. Similarly, files are frequently licensed with the ability to use newer versions of the license (corresponding to the + sign in the SPDX abbreviations of license names, such as GPL-2.0+).

have found in Debian 7.5. Finally we categorize them based on our analysis of our two target datasets.

2.1 Definition

For the purpose of this research, a *license inconsistency* refers to the situation when two source files that have evolved from the same original code have license statements which include different licenses.

2.2 Example

In the Debian 7.5 Linux distribution, two packages, `dpkg` and `anubis`, contain a file named `obstack.c`. Except for the license statement, these two files are identical. For this reason we assume that these two files share the same provenance.

From package `dpkg`, the license of this file is GPL-2.0+:

```
[...]  
This program is free software; you can redistribute it  
and/or modify it under the terms of the GNU General Public  
License as published by the Free Software Foundation; either  
version 2, or (at your option) any later version.  
[...]
```

While from package `anubis` the license is GPL-3.0+:

```
[...]  
This program is free software: you can redistribute it  
and/or modify it under the terms of the GNU General Public  
License as published by the Free Software Foundation; either  
version 3 of the License, or (at your option) any  
later version.  
[...]
```

As we can see, the licenses of the two files are different: GPL-2.0+ and GPL-3.0+. The first file can be combined with software under the GPL-2.0, but the second cannot (the GPL-2.0 is incompatible with the GPL-3.0 or any later version of the GPL). Based on our definition, this is a case of license inconsistency. Without tracing the history of each of these files, it is not possible to determine if both licenses are valid (i.e. if the copyright owner of made the file available under both licenses). The following are three of many potential scenarios that lead to this inconsistency:

- The first file is the original one and was copied to the second project, where the license was changed from GPL-2.0+ to GPL-3.0+. In this case, since the original license allows to use newer versions of the license, the change can be done by anybody, and it is not a potential license violation.
- The second file is the original one and was copied to the first project. The first project changed the license version from GPL-3.0+ to GPL-2.0+. This could be a potential violation if the change was made without the approval of the copyright owner of the file.
- Both of the files are copied from the same third-party project (who created the file). Each project made the copy at different times, one before, and one after the license of the file was changed by the original project copyright owner. In this case, there is no potential license violation.

To determine which one is the actual reason of the inconsistency, we need to examine the repository history of these two projects and try to determine the true origin and rationale for the of license. This topic will be discussed in Section 4.1.1.

2.3 Categorization

Based on the analysis of our two datasets, we observed 5 cases of license evolution. They are either executed by the original author or reuser:

1. **License Addition:** The source file was without a license, and a license is added in a later release.
2. **License Removal:** The source file was under a certain license, and the license is removed in a later release.
3. **License Upgrade:** The source file was under a certain version of the GPL license, and it is upgraded to a newer version of the GPL license.
4. **License Downgrade:** The source file was under a certain version of the GPL license, and it is downgraded to an older version of the GPL license.
5. **License Change:** The source file was under a certain license, and it is changed to another license (except for License Upgrade and License Downgrade).

Note that, in the case of license upgrade and downgrade, we only consider the GPL license. This is because currently only the GPL licenses have a “*or later*” option (e.g. GPL-2.0+, LGPL-2.1+) which allows the reuser to choose a later version of GPL as the license for redistribution (i.e. to *upgrade* to a newer version). Although some other licenses, such as Apache license, may have different versions, reusers are not allowed to choose an arbitrary version of the license. Thus it is reasonable to treat various versions of these licenses as completely different licenses. For such reason we treat the license evolution between different versions of licenses other than GPL as *license change* in our research.

License inconsistencies are naturally caused by changes in the license of the files. We use the following types to denote different types of license inconsistencies between two files:

LAR License Addition or Removal. One of the two files contains a license while the other file contains no license. This type of license inconsistency is usually caused either a License Addition or a License Removal.

LUD License Upgrade or Downgrade. One of the two files contains a certain version of a license while the other file contains a different version of the same license. This type of license inconsistency is caused by either upgrading or downgrading the license of the file.

LC License Change. Two files contain different licenses. This type of license inconsistency is usually caused because the license of the file was changed.

3 Study Approach

In our previous work [25], we have proposed a method that can efficiently detect license inconsistencies. However, a major issue with that method is that it only considers license inconsistencies among files that have the same base name in order to achieve a fast performance. Thus if files are renamed during the process of copy-and-own reuse, the license inconsistencies will not be detected. To solve this problem and make our result cover more license inconsistencies, we propose a new method in this paper. A detailed comparison of these two methods will be discussed in Section 5.

In our new approach, we focus on detecting license inconsistencies among file clones. In the scenario of source code reuse where source files are imported from an upstream project, the contents of reused source files remain almost the same, sometimes with small changes (such as modifying comments, renaming identifiers etc.) [21].

To decide whether source files are copies of each other—or in other words whether they share the same provenance—we compare their *normalized token sequences*[20]. Normalized token sequences are generated from the source file by removing the comments, redundant white spaces, new lines, carriage returns and then converting identifiers to normalized tokens. If two files have the same normalized token sequences, then it is likely that they are copies of each other and we call them *file clones*, which are actually Type-2 code clones[20, 21]. We use **CCFinder** [14], a code clone detection tool, to analyze and determine if files are file clones. **CCFinder** will generate a pre-process file which contains the normalized token sequences of the source file. For those file clones with the same normalized token sequences, we assume that they come from the same origin, and then gather them into the same *file group*. Files in the same file group might have different base names but similar program statements, possibly with different comments including license statement.

Once that we group these similar files, we identify the license of the files in each group. In our approach we adopted **Ninka** to detect the license of source files, since **Ninka** is reported to have the highest precision of all the license detection tools including **FOSSology**, **ohcount** and **OSLC** in the research by German et al. [7]. **Ninka** is a sentence-based license detection tool which can identify 110 different licenses with 93% accuracy, and it can handle more than 600 files per minute. There are two special results from **Ninka**: one is *UNKNOWN*, which represents that **Ninka** has found a license but does not recognize it. The other one is *None*, which states that the source file has no license.

We then compare the licenses of each file in the license list of each group. If all the files have no license, or all of them have the same license, then there is no license inconsistency. Otherwise, the group is likely to contain license inconsistencies. And then, based on the relation between licenses, our approach identifies the type of license inconsistency. Note that a group may have multiple license inconsistencies. For example, if a group include a file under GPL-2.0+, another file under GPL-3.0+ and the other file under Apache-2.0, the group has two license inconsistencies: *LUD* between GPL-2.0+ and GPL-3.0+, *LC* between GPL-2.0+/GPL-3.0+ and Apache-2.0. For such reason, we calculate *License Inconsistency Metrics* for each of these groups, from which we can measure what type of license inconsistencies and how many of each type exist in the groups.

3.1 License Inconsistency Metrics

The following 5 metrics are introduced to help measure the license inconsistencies for a file group:

#File: Number of files in this group.

#Lic: Number of different licenses in this group. If there are two or more licenses found, then it is likely that there is a license inconsistency. If no license, or only one license is found, then all the files are either without license, or they have the same license.

#Unknown: Number of files with an unknown license in this group. For our purposes we consider all the files with unknown licenses as if they have the same license (this might under-estimate the number of inconsistencies).

#None: Number of files without any license in this group. If $\#None > 0$ and $\#Lic > 0$ then it is possible that at least one file in the group had its license added or removed (i.e. *LAR* inconsistency).

#GPL: Number of licenses in GPL family. This metric allows us to identify *LUD* in the GPL family.

These metrics are calculated for each file group based on their license lists. The strategies shown in Table 1 enable us to decide whether a certain type of license inconsistency exists in this group.

Specifically, if we query the metrics result for those groups with $\#None > 0$ and $\#Lic > 0$, which means there are one or more files with no license, and also one or more files

Table 1: Strategies to decide whether a certain type of license inconsistency exists in a group.

Inconsistency Type	Strategy
<i>LAR</i>	$\#None > 0$ and $\#Lic > 0$
<i>LUD</i>	$\#GPL \geq 2$
<i>LC</i>	$\#GPL \leq 1$ and $\#Lic \geq 2$

contain a license. According to our definition, this is *LAR*; If we query for those whose $\#GPL \geq 2$, which tells us that there are two or more different licenses in GPL family (such as GPL-2.0+ and GPL-3.0+), and should be *LUD*; If we query for those items with $\#GPL \leq 1$ and $\#Lic \geq 2$, which means there are more than two licenses appear in this group and no more than one GPL license exists (to exclude *LUD* case), it seems to be the case that one license is changed to another one, which should be a *LC*.

3.2 Method of Detecting License Inconsistencies

As a summary, our method is divided into 3 steps:

1. **Create groups of file clones:** For all the source files in the target projects, we apply *CCFinder* to extract the normalized token sequences of each file. By computing and categorizing the hash value of these token sequences, we then create a *group* for files that have the same normalized token sequences. Each group contains at least two files; i.e. a single unique file is not contained in any group.
2. **Identify licenses for files in each group:** For each group of file clones, *Ninka* is employed to identify the license(s) of each file. The result is a list of licenses for each file group.
3. **Report groups that contain license inconsistencies and calculate the inconsistency metrics:** We compare the license list of each file group. File groups are reported to have license inconsistencies unless all the licenses on the list are exactly the same. The result is a list of file groups that contain one or more types of license inconsistencies.

3.3 Example

We illustrate our method with a simple imaginary project shown in Figure 1. This project consists of 4 packages. The source code of `foo.c` file in `Pkg2` is exactly the same with the one in `Pkg1`, but the license statement is changed from `GPL-2.0+` to `GPL-3.0+`; The source code of `foo.c` in `Pkg3` is different from the one in `Pkg1`, i.e. they happen to have the same base name. It is reused in `Pkg4` with its name changed to `foo100.c` and license statement removed.

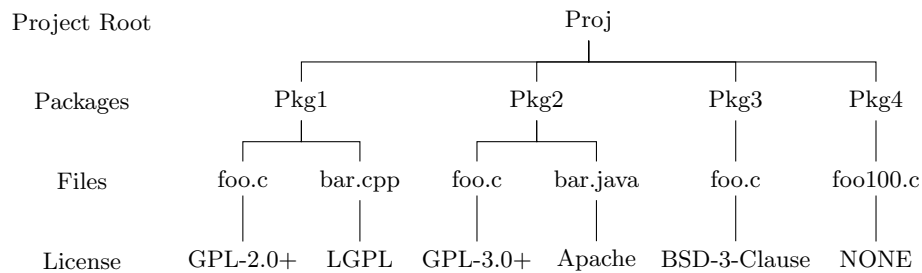


Figure 1: Hierarchy of an imaginary project and the license of each source file. Note that the `foo.c` file in `Pkg1` was imported to `Pkg2` with the license changed to `GPL-3.0+`; The `foo.c` in `Pkg3` contains totally different source code than the one in `Pkg1`, and was imported to `Pkg4` with its name changed to `foo100.c` and license removed.

1. **Create groups of file clones:** In this step, we use `CCFinder` to generate token files for each source file. Since the `foo.c` file from `Pkg1` and `Pkg2` have the same source code (except for their code comments which include license statement), `CCFinder` treats them the same, and generate the same token file. This also applies to `foo.c`




Figure 2: Hierarchy of the grouped files.

Table 2: License list of the selected files from the imaginary project.

Base name	GroupID	Package name	License
foo.c	1	Pkg1	GPL-2.0+
foo.c	1	Pkg2	GPL-3.0+
foo.c	2	Pkg3	BSD-3-Clause
foo100.c	2	Pkg4	NONE

Table 3: List of the license inconsistency metrics for each file group in the imaginary project.

GroupID	#File	#Lic	#None	#Unknown	#GPL
1	2	2	0	0	2
2	2	1	1	0	0

file from Pkg3 and foo100.c from Pkg4. Thus we can compare the hash value of the token files and group them into two groups, as shown in Figure 2.

- Identify licenses for files in each group:** For each file in the group, we use `Ninka` to detect their licenses and make a list of the base name, group index and the licenses, as shown in Table 2. *Base name* is the name of the source file. *GroupID* indicates the index we use to identify file groups.
- Report groups that contain license inconsistencies and calculate inconsistency metrics:** We examine the licenses of each group and found that both of these groups contain license inconsistencies. Thus we report both of these groups and compute the inconsistency metrics for each of them, as shown in Table 3.

According to our rule, $\#GPL \geq 2$ in Group 1 indicates a case of *LUD* in this group, while $\#None > 0$ and $\#Lic > 0$ in Group 2 indicates a case of *LAR* in this group. This conclusion is consistent to the scenario in our imaginary project, since the two foo.c files in Pkg1 and Pkg2 contain GPL-2.0+ and GPL-3.0+ respectively which is *LUD*, and the file foo.c in Pkg3 and foo100.c in Pkg4 contain BSD-3-Clause and no license respectively which is *LAR*.

Table 4: Main characteristics of Debian 7.5.

Characteristics	Number
Source Packages	17,160
Total files	6,136,637
.c files	472,861
.cpp files	224,267
.java files	365,213

4 Empirical Study

We have selected two target datasets for analysis: Debian 7.5 Linux distribution⁸ and a large number of Java projects downloaded from GitHub⁹. We then conducted our method on both datasets respectively. Since it is hardly feasible to determine how many and what types of license inconsistencies are there in the target projects, it is difficult to get an oracle data set and to perform a quantitative evaluation of our method, specially regarding its recall. However, a qualitative evaluation of this method is discussed in Section 5.5.

The following subsections will present the results obtained from the two datasets, respectively.

4.1 Empirical Study on Debian 7.5

We conducted our study using a large open source Linux distribution, Debian 7.5. The source code was downloaded from its official site and its main characteristics are shown in Table 4. Only .cpp, .c and .java files are used, since they account for the majority of source code in the Debian distributions and are the only file formats that are supported by CCFinder.

4.1.1 Results

In the first step, we grouped the files under each set by their normalized token sequences and resulted in 125,092 groups in total. And the number of files within one group ranges from 2 to 160, and the average number of files per group is 2.8 with a median value of 2. The breakdown of each file type is shown in Table 5.

⁸<https://www.debian.org/>

⁹<https://github.com/>

Table 5: Breakdown of number of groups and files for each type in analyzing Debian 7.5.

File type	Group count	File count	#Files(mean)	#Files(median)
.c	68,568	207,620	3.0	2
.cpp	16,202	38,617	2.4	2
.java	40,322	108,868	2.7	2
Total	125,092	355,105	2.8	2

Table 6: Partial list of the license inconsistency metrics for each file group in detecting Debian 7.5.

Base name*	GroupID	#File	#Lic	#None	#Unknown	#GPL
obstack.c	6645	19	2	0	0	2
getopt.c	46474	6	2	3	0	0
getopt.c	52662	9	2	1	7	1
...

* Each group may contain files with different file names, here we show the most frequent *base name* in that group.

Completing the following two steps, 6,763 groups were reported to have at least one type of license inconsistency, which is 5.4% of the 125,092 groups in total. For the sake of space, we show only part of them in Table 4.1.1.

Then we calculate the number of each type of license inconsistency and their proportion. The result is shown in Table 7. From this table, we can see that from the total of 6,763 groups that contain one or more license inconsistencies, 67.5% of them contain *LC*, followed by *LUD* and then *LAR*. As it can be seen, *LC* is the most common license inconsistency, suggesting that developers are likely to change the license of the source file to another one. For such reason, further study is urged to investigate the legality of these modifications.

In the following paragraphs, we show examples for each type of license inconsistencies.
– *LAR*:

Examining the `getopt.c` in second line from the inconsistency result list in Table 4.1.1, we get the license list of that group in Table 8. The rest files that contain the same licenses are omitted from this list.

Table 7: Number of different license inconsistency types and their proportion in Debian 7.5. Note that one group may contain more than one inconsistency types, so that the total percentage can exceed 100%.

Inconsistency type	Frequency	Perc.
<i>LC</i>	4,562	67.5%
<i>LUD</i>	2,137	31.6%
<i>LAR</i>	883	13.1%

Table 8: Example of *LAR* inconsistency, in `getopt.c`

Package name	License
icedove	NONE
iceweasel	MPL-2.0

We can see that the license of the `getopt.c` file from the `iceweasel` package has an MPL-2.0 license while the one from package `icedove` has no license (marked as NONE). The contents of each file is as follows.

`getopt.c` from `icedove` package:

```
#include <stdio.h>
#include <string.h>
[...]
int main(int argc, char **argv)
{
    PLOptState *opt;
    PLOptStatus ostate;
    [...]
    return 0;
}
```

`getopt.c` from `iceweasel` package:

```
/* This Source Code Form is subject to the terms of the
 * Mozilla Public License, v. 2.0. If a copy of the MPL
 * was not distributed with this file, You can obtain one
 * at http://mozilla.org/MPL/2.0/.
 */
#include <stdio.h>
#include <string.h>
[...]
int main(int argc, char **argv)
{
```

```

    PLOptState *opt;
    PLOptStatus ostate;
    [...]
    return 0;
}

```

As we can see in the file from `icedove` package, there is no license statement at all, while the file `getopt.c` from `iceweasel` package contains a MPL-2.0 license. Meanwhile, the other parts of these two files are exactly the same, hence we consider it safe to assume that the origin of both files is the same. There are several possible explanations to this case of license inconsistency:

1. The file from `icedove` package is the original one, and the developers of `iceweasel` project reused the file and added a license to it.
2. The file from `iceweasel` package is the original, and developers of `icedove` project reused this file and removed the license statement.
3. Both of the files in these two projects reused different versions of this file from another project (where the license was added or removed).

One way to try to discover which one is the true explanation is to look at the history of the files in their corresponding version control repositories. By tracing the revision history of both files, we found that the third possible explanation reflects the actual history of the files: the files in these two projects were imported from a third project named `nspr`, where the `getopt.c` file was created without a license in version 4.7.1, and, for version 4.9.1 the license was changed to the MPL-2.0. It seems that `icedove` reused this file before the license statement was added, while `iceweasel` imported the version after the license was added, thus caused the inconsistency of license.

– *LUD*:

To exemplify this inconsistency, we will use `obstack.c`, which is the first in Table 4.1.1. Table 9 shows two packages that reuse this file. As we can see from this table, the first file is licensed under GPL-2.0+ while the second one is under GPL-3.0+.

The license statements of the files from `dpkg` and `anubis` package were listed in Section 2.2. Both of these files contain more than 400 lines of code, and they are exactly the same except for their license statements. Tracing the file history in both projects we found that this file was originally created in `gnulib`. The license of this file was upgraded-in

Table 9: License list of group 6645 of `obstack.c` where *LUD* exists.

Package name	License
<code>dpkg</code>	GPL-2.0+
<code>anubis</code>	GPL-3.0+

Table 10: License list of group 52662 of `getopt.c` where *LC* and *LAR* exist.

Package name	License
<code>p0f</code>	NONE
<code>snort</code>	GPL-2.0
<code>sofia-sip</code>	UNKNOWN (IBM)

`gnulib`—from GPL-2.0+ to GPL-3.0+. By examine the commit log of `dpkg`, we found that the developers of `dpkg` intentionally reused the older version of the file from `gnulib` project (they wanted the file to be licensed GPL-2.0+, not GPL-3.0+), which caused the license inconsistency.

– *LC*:

We demonstrate this inconsistency using `getopt.c`, the third line from the Table 4.1.1.

As shown in Table 10, `getopt.c` from `snort` package contains GPL-2.0 while the license of the one from `sofia-sip` could not be recognized.

The contents of these files are as follows.

`getopt.c` file from `snort` package:

```
[...]
** it under the terms of the GNU General Public License
** Version 2 as published by the Free Software Foundation.
** You may not use, modify or
[...]
```

`getopt.c` file from `sofia-sip` package:

```
[...]
* COPYRIGHTS:
*This module contains code made available by IBM
*Corporation on an AS IS basis. Any one receiving the
*module is considered to be licensed under IBM copyrights
*to use the IBM-provided source code in any way he or she
*deems fit, including copying it, compiling it, modifying
[...]
```

Table 11: The count and percentage of each category for the 25 investigated license inconsistency cases.

Category	Count	Perc.
Safe changes	14	56%
Unsafe changes	5	20%
Uncertain cases	6	24%
Total	25	100%

From the header we know that the second file is licensed under IBM copyrights, but this is not a standard version of IBM Public License, thus `Ninka` reported it as UNKNOWN. Since both these files contain the same program code, we may assume that someone changed the license from one to the other. We tried to find out the direction of this change, but due to lack of history it was not possible to do so. This shows that determining the true provenance of a file is difficult in general.

4.1.2 Manual Analysis

To decide whether these license inconsistencies may indicate legal problems or not, we have conducted a manual analysis on the history of a subset of the files.

We randomly chose the samples using trial-and-error methodology, that is, first we randomly select a case of license inconsistency and investigate whether it is legally safe or not, then we randomly select the next case and repeat the process. Due to the difficulties and the time invested, we have only sampled 25 cases in total. Then we tried to categorize them according to the reason that caused such inconsistencies. They are divided them into three categories: safe changes (no violation is found), unsafe changes (there appears to be a violation) and uncertain (it was not possible to determine whether it was safe or unsafe.). The results are shown in Table 11, and the a detailed explanation of each category is as follows:

- **Safe Changes:** In this category, either the original author or the developers who reused the file changed the license statement, but the change they made is based on the terms described in the license thus we classify it as a safe change. They are further divided into 2 groups:

Original author modified/upgraded the license. In this case, the author of that file modified the license statement (either by upgrading or totally changing it to another license), while the reusers still use the old version of the file (either intentionally or unintentionally).

For example, we examined a file named `obstack.c` in our inconsistency result. This file originates from `gnulib` project, and its license is upgraded from GPL-2.0+ to GPL-3.0+ in a commit on 10/7/2007. This file was reused in the `dpkg` project but with a GPL-2.0+ license, and in the last commit on 9/25/2011 the log is as follows:

```
libcompat: Update obstack module from gnulib. The version taken is the one
before the switch to GPL-3.0+. With a slight code revert to not have to include
exitfail.c and exitfail.h.
[...]
```

We can see that in this case, the reuser intentionally takes an older version from the original project, which caused the inconsistency of license.

In another example, there is a file named `paintwidget.cpp`, which originates from `Qt` project with BSD-3-Clause license. In another project called `PySide`, this same file is licensed under LGPL-2.1/GPL-3.0 dual license. Since these two projects both belong to Digia plc, which were acquired from Nokia, this shall be a legal license modification.

The file was originally multi-licensed and reusers chose either one. The author of the file licensed the file under two or more licenses, and the reusers can choose either one of them.

There is a file named `SimpleXMLParser.java` which originates from `iText` project and was under the Mozilla MPL-1.1/LGPL-2.0+ dual license. This license allows the removal of one license. Developers in `pdftk` project reused this file removing the MPL-1.1 license and chose LGPL-2.0+ as its license.

– **Unsafe Changes:** Under this category, developers who reused the source file seemed to have modified the license statement which is not allowed by the original license terms. This change may lead to legal disputes, thus we say it is an unsafe change. We should clarify that we have reached this conclusion based on the historical evidence available. The consequence is that anybody who would like to reuse these files should pay special attention to these cases, and do due diligence to determine what is the appropriate licensing of the file, and if it indeed poses a legal risk.

Reuser replaced the original license, and changed the copyright owner. The file is under a certain license in the original project and developers who reused the file changed the license statement and the copyright owner.

From our inconsistency list, we examined a file named `SpringUtilities.java`. According to the copyright year, Oracle is the copyright owner, and licensed the file under BSD-3-Clause. When reused in `freemind` project, developers changed the license to GPL-2.0+ and the copyright header, which is not allowed in BSD-3-Clause. This kind of changes to the license statement by the reuser may lead to license infringement, and may involve the reuser into legal disputes.

Reuser added one or more licenses. The original file is under some licenses, and the reuser added one or more licenses to it while retaining the original license.

From the result we examined a file named `DOMException.java`. This author of this file is World Wide Web Consortium (W3C), and was licensed under W3C Software License. When developers reused this source file in `ikvm` project, they added a GPL-2.0 License to it resulting a composition of these two licenses. Meanwhile, the program code of this file was not changed at all. We consider this case as unsafe, since this type of license modification makes it unclear which part contains the original license and which part contains the new license, since they added the license without adding any source code changes to the file.

– **Uncertain Cases:** This category contains the license inconsistency cases which are difficult to determine whether they are legally safe or not due to several reasons:

Source files are too small. Some files contain the same source code, but due to their small size it is difficult to decide whether one is reused by the other or they just happen to be the same. This problem is discussed in Section 5.5.

Files can not be found in the upstream repositories. We found many cases of licenses inconsistencies in the projects in Debian 7.5 that, when investigated the upstream project's repository, the file no longer existed.

For example, our method reported a file named `jim-win32.c` in `jimtlc` package with BSD-2-Clause license and in `openocd` package with Apache-2.0 license. When we tried to look for this file in the repository of `openocd` project, it was not found. One explanation is that the file was removed in the project, but was not yet updated in Debian 7.5.

Project repository not available. Some project repositories could not be found due to lack of documentation, while some could not be accessed due to server error.

One example is, when we tried to checkout the source code of `axis` project using the SVN command on its official website¹⁰, the command returned an error that the URL does not exist.

¹⁰<https://axis.apache.org/axis/cvs.html> (Last access: Oct. 2nd, 2015)

Table 12: Main characteristics of Java projects cloned from GitHub.

Characteristics	Number
Projects	10,514
Total files	3,374,164
.c files	15,627
.cpp files	21,176
.java files	3,337,361

Table 13: Number of groups and files in each group in analyzing Java projects.

File type	Group count	File count	#Files(mean)	#Files(median)
.java	199,284	769,220	3.9	2

4.2 Empirical Study on Java Projects

The other data set we studied is a large collection of Java projects cloned from GitHub. The snapshot was taken in Mar. 2015, and only those projects that consist of at least 100 commits are selected. Table 12 shows the characteristics of these projects. Since .java files are 98.9% of all the files, we will focus our following analysis on them only.

4.2.1 Results

In the first step, source files are grouped by their parameterized token sequences. The result was 199,284 groups. The number of files within each group ranges from 2 to 1514, and the average number is 3.9 with a median value of 2, as shown in Table 13.

With the following steps being done, 13,916 groups are reported to contain license inconsistencies, which is 7.0% of the 199,284 groups in total.

Furthermore, the number and proportion of each type of license inconsistency is shown in Table 14.

Table 14: Number of different license inconsistency types and their proportion in Java projects.

Inconsistency type	Number	Perc.
<i>LC</i>	12,653	90.9%
<i>LAR</i>	6,179	44.4%
<i>LUD</i>	1,316	9.5%

Table 15: The count and percentage of each category for the 17 investigated license inconsistency cases in the Java projects.

Category	Count	Perc.
Safe changes	11	65%
Unsafe changes	1	6%
Uncertain cases	5	29%
Total	17	100%

4.2.2 Manual Analysis

As we did in the Debian study, we examined a random sample of the inconsistent groups. We sampled 17 cases, and tried to categorize them according to the reason that caused such inconsistencies. As described before, they are divided into three categories, the percentage of each category is shown in Table 15, and the explanation to each category is as follows:

– **Safe Changes:**

Source files are in the same project but with different licenses. Some projects were imported from other version control systems, such as SVN, where branching and tagging makes copies of the whole project. When the license of source files in the main branch (trunk) changes, license inconsistency occurs among these branches.

For example, there is a project named `weka` which was imported from SVN. In this project, files were originally licensed under GPL-2.0+ and then upgraded to GPL-3.0+. Developers made a series of tags in the SVN repository, leaving several copies of the whole project. Thus license inconsistency exists between the files under the tags which were made before the license upgrade and those in the trunk.

Some other cases are, the source files are in the same project but exist under different directories with different licenses.

Duplicated projects are not up-to-date. Some entire GitHub projects—or subdirectories in other cases—act as a backup (or a copy) of another project, and their license of source code is not updated while the original project changed its license.

We examined two projects: `JCryptTool`¹¹ and `JCT-CA`¹². A file named `ResizeHelper.java` exists in both projects with the same normalized token sequences. The one in `JCT-CA`

¹¹<https://github.com/jcryptool/cryptool>

¹²<https://github.com/Kalliope/minica>

is without a license, while the one in `JCryptTool` was originally with no license but then added with a EPL-1.0. The readme file from JCT-CA states:

```
JCT-CA is going to be a plugin for the JCryptTool regarding Public
Key Infrastructure. Main development is done in the master branch,
others (if any) are just for backing up older parts of the project
and keeping master clean.
```

From this notice we can see that, this project is a partial backup of the `JCryptTool` project, but its license is not up-to-date when the original copy has changed, resulting in license inconsistencies.

Reusers added a same license to the source file. One rare case we found is, the developers of a reused a source file—under Apache-2.0— added another exactly same Apache-2.0 license description in the header. One explanation is that the developers are using automated tools to manage the licenses, but did not check whether the file already contains a license. Though it does not conflict with the license terms, we consider it as a bad smell.

– **Unsafe Changes:**

Reusers modified the license terms. Some developers reused the code from other projects but made some modifications to the license terms. In this case, if it is not with the permission from the original author, these modifications are unsafe.

There is a file named `MNP.java` in both `kawa-fork`¹³ and `classpath`¹⁴. This file is originally from `Kawa`, and `kawa-fork` is a fork of the this project. The license of this file in `Kawa` is MIT, while the one in `classpath` was changed to GPL-2.0+ with link exception.

– **Uncertain Cases:**

Licenses are modified outside the scope of their repositories. There are cases that, the source files in different projects are with different licenses, but their license statements have never changed since they were imported into these repositories. Another alternate explanation is that developers downloaded the software and modified the license before the first commit into the new repository, making it impossible to track the point where the license was changed.

Source files are too small. This case is same as the one in Debian data set. They will be discussed in Section 5.5.

For example, a file named `ReaderInputStream.java` was found in `bingo-core` project with an Apache-2.0 license and in `hibernate-orm` project with an LGPL license. However,

¹³<https://github.com/maoueh/kawa-fork>

¹⁴<https://github.com/penberg/classpath>

the source code contents of these files are quite small, which merely contains two empty constructor methods. The source code part excluding the comments is shown as following:

```
[...]  
import java.io.IOException;  
import java.io.InputStream;  
import java.io.Reader;  
public class ReaderInputStream extends InputStream {  
    private final Reader reader;  
    public ReaderInputStream(Reader reader){  
        this.reader = reader;  
    }  
    @Override  
    public int read() throws IOException {  
        return reader.read();  
    }  
}
```

It is possible that different developers write the same code like this from scratch, thus it is difficult to judge whether these files are copies of each other.

5 Discussion

In this section, we first analyze the results from the previous section, and present some problems we met with in the process of data analysis. Then we summarize the answers to our research questions. Next, we show the improvement we made to the research method with a comparison between these two methods. Finally, we present a reply from a developer team we contacted.

5.1 Discussion of the Results

From these results we can see that the license inconsistencies are not uncommon: in Debian 7.5, out of 125,092 file groups, 6,763 (5.4%) of them contain one or more license inconsistencies: *LC* has the highest proportion with 67.5%, followed by *LUD* with 31.6%, *LAR* comes next with 13.1%. While in Java projects, out of 199,284 file groups, 13,916 (7.0%) of them contain one or more license inconsistencies: *LC* has the highest proportion with 90.9%, followed by *LUD* with 44.4%, *LAR* comes next with 9.5%.

The manual analysis of several cases of license inconsistencies gives us a rough understanding of the how many of these cases are safe or not. From Table 11 and Table 15 we can see that, both in Debian 7.5 and Java projects we selected, unsafe and uncertain cases take up 44% and 35% respectively. This shows that it is not uncommon that license inconsistencies might lead to potential license violation problems.

During this process of analysis, we also found several challenges that prevent us from automatically analyzing the history of files.

Many files in an open source project are frequently imported from other projects. It is not a trivial task to find the repositories of these upstream projects. Take the Debian distribution as an example: some of the packages contain a file indicating the repository URL of that package, but some do not. For such packages, we needed to search for the official site of the upstream project and try to find its repository URL. There are packages that appear not to use version control systems. They simply provide source code tarballs for each version on their server. In this case, we have to download each tarball and track the license change manually. This makes provenance tracing more difficult.

In some cases the change of the license statement is not recorded in the revision history because the license statement is changed (we presume) before the file is added to the repository's project. In this case, we have to check other information (e.g. on the official

site of the project or in the commit comment where the file was added) to find out the reason why developers changed the license.

Also, after we found out that the files with same normalized token sequences in different packages contain different licenses, we have to determine where the file comes from, i.e. the original project of that file, in order to decide the direction of the license change. But to the best of our knowledge, there is no good way to find the true origin of a certain file. We address this problem by using the date of the first commit of that file as a reference. When we have two copies in different repositories, we assume that the file with the oldest commit is the original, and files with newer dates are copies of it. If the commit date is not available—e.g. when not using a version control system—we have to manually check the comments of the source file to see if it contains information about its true origin or its license. If not, then we are not able to decide which file comes first.

5.2 Answering RQs

Revisiting the research questions:

- **RQ1:** *How can we categorize license inconsistencies?* We categorize license inconsistencies into these 3 types: *i) LAR*, which is typically caused by license addition or removal; *ii) LUD*, which is related to license upgrade or downgrade in the GPL family; *iii) LC*, which is usually caused by license change in the process of license evolution.
- **RQ2:** *Do license inconsistencies exist in open source projects?* Yes, license inconsistencies exist in open source projects. As we have shown in our empirical studies of Debian 7.5 and a large collection of Java projects on GitHub, various types of license inconsistencies were detected.
- **RQ3:** *What is the proportion of each type of license inconsistency?* In the case study of Debian 7.5, out of 125,092 file groups, 5.4% of them contain one or more license inconsistencies. The proportion of each type is: *LAR* (13.1%), *LUD* (31.6%) and *LC* (67.5%). In the case study of Java projects, out of 199,284 file groups we selected, 7.0% of them contain one or more license inconsistencies. The proportion of each type is: *LAR* (9.5%), *LUD* (44.4%) and *LC* (90.9%).

- **RQ4:** *What caused these license inconsistencies? Are they legally safe?* The reasons that caused license inconsistencies can be summarized into these groups according to our observation:

- i)* Original author modified/upgraded the license.
- ii)* The file was originally multi-licensed and reusers chose either one.
- iii)* Reuser added one or more licenses.
- iv)* Reuser appear to have replaced the original license, and changed the copyright owner.

We consider the last two types of modification as unsafe, which would require further analysis to determine the legal risk associated with using them.

5.3 Improvement of the Method

As described in Section 3, our previous method [25] omits the cases if the files are renamed during the process of copy-and-paste reuse to achieve higher performance.

In the old method, we assume that many copy-and-paste reuse are conducted without renaming the source files. Thus we first create file sets where each set contains source files with the same base name. And then, under each file set, we then group the files by their normalized token sequences. Finally, we detect the licenses for each file in every file clone group and calculate the license inconsistency metrics.

In this paper, however, the new method treats all the source files as a whole set, and groups them by their normalized token sequences. Thus it should obtain a more comprehensive result of license inconsistencies.

The following two subsections compare the two methods on the two data sets we used, respectively.

5.3.1 Debian 7.5

Table 16 shows the comparison of results obtained by the two methods, for Debian 7.5.

As we can see from the table, the new method covers all the groups that the old method reported. Besides, it also reported 1419 (21.0%) more license inconsistency groups. As a conclusion: the result from the new method is a superset of the one from the old method, which is consistent with our expectation.

Table 16: Comparison of two methods on Debian 7.5.

Number of groups	New method	Old method
Total	6763	5344
Intersection ⁱ	5344	5344
Relative complement ⁱⁱ	1419	0

ⁱ Intersection indicates the groups both method reported.

ⁱⁱ Relative complement indicates the groups reported in one method but not the other.

Table 17: Comparison of two methods on Java projects.

Number of groups	New method	Old method
Total	13,916	13,894
Intersection	13,894	13,894
Relative complement	22	0

5.3.2 Java Projects

Table 17 shows the comparison of results obtained by the two methods when applied to the Java projects in GitHub.

Again we can see from this table, the new method covers all the groups that the old method reported. However, there are merely 22 more groups reported by the new method, from which we can infer that the renaming operations are not frequently conducted in the process of copy-and-paste code reuse in these Java projects. This also proves that our old method is able to produce a good result in detecting license inconsistencies where rename operation are not often conducted during the process of code reuse.

5.4 What appears to be a copy might not be a copy

We sent emails to the 3 development teams of the projects where unsafe license modification were found, to understand why they modified the license and whether they consider it as an illegal modification. However, only one of them replied us, claiming that they wrote the source code all from scratch, and denied that this source file was copied from somewhere else. This source file was so small which contains merely two empty constructors, thus we believe it is possible that different developers happen to create the same file.

Note that, this is not a false positive case of our method, since our method is designed to detect the *license inconsistencies* in the target projects, not the *license violation* cases. However, it stresses the need to consider a minimum size threshold, in order for these small files not be considered in the analysis.

5.5 Threats to Validity

In this study, we use `CCFinder` to detect file clones which are exactly identical to each other regarding their normalized token sequences. However, source code files are evolving: those that come from the same provenance may differ from each other dramatically, making their normalized token sequences different after being modified by developers. But since we can still get large numbers of file groups that contain license inconsistencies using the proposed method, we believe that it is enough for this exploratory study. To mitigate this problem, we can use similarity metrics instead.

On the other hand, during our manual analysis we found files clones that contain the same normalized token sequences, but due to their small size and simplicity, it is difficult to decide whether they are copies of each other or they were written from scratch by independent developers. If the later one is the real case, then it would be a false positive of our result. But we believe it might be good practice to report these cases, have a manual investigation on them and ask the developers directly.

One aspect that is important to highlight is that the our method relies on the ability to detect copies of files. In our previous paper [25], we found copies of files by analyzing files with the same name. In this paper we compared the normalized token sequences of files. We could also do full clone detection and consider two files to be copies of each other only if they were above certain threshold. This process would have been significantly more time consuming. Ultimately, detecting license inconsistencies is a balance between performance of the detection vs. recall. If necessary, step one of our method can be replaced with other methods that provide better recall, at the expense of being slower, and potentially require more manual analysis to filter false positives.

It is also important to highlight that the ability to detect license inconsistencies relies heavily on having a comprehensive corpus to compare against. In this study we have used two collections of source code: Debian 7.5 and Java GitHub projects. License inconsistencies in the source code that an organization reuses can only be found if the original code is in the corpus that is being compared against.

In the process of license identification, as we employed `Ninka` to identify the license of source files, its accuracy should also be considered. German et al. reported that the accuracy of `Ninka` is 93% [7]. We believe this is sufficiently high, so that the license detection result is good enough to support our analysis. In addition, we regard UNKNOWN licenses as the same license within each group, different from any other licenses. If these UNKNOWN licenses in a same group are actually different from each other, we may underestimate license inconsistencies. But this concern is mitigated according to our observation to these UNKNOWN licenses: most of those in the same group actually contain the same license statement, either a license that is not approved by OSI or a user modified version of an OSI-approved license. On the other hand, if these UNKNOWN licenses are actually the same as those recognized ones (e.g. GPL-2.0, BSD-3-Clause etc.) in the same group, this could be considered as a false positive. In this case, these UNKNOWN licenses are not exactly the same as the original license (since `Ninka` has reported a different license), meaning that someone must have modified the license statement. We believe that it is necessary to check whether these changes are legal or not. Thus it is reasonable to treat them as license modifications, which is consistent with our assumption. To obtain more precise results, it is necessary to improve license identification.

6 Related Work

Many studies address inconsistent changes among code clones. Krinke [15] studied on changes applied to code clones in open source software systems and showed that half of the changes to code clone groups are inconsistent changes and these changes are not solved if they occurred in a near version. Göde et al. [11] studied patterns of consecutive changes to code clone in real software systems. Some approach to find inconsistent changes are proposed [5, 13]. On the other hand, Bettenburg et al. [2] showed that only 1% ~ 4% of inconsistent changes to code clone introduce software defects. In addition, Göde et al. [12] showed that most code clones do not evolve and the number of inconsistent changes is small. Our work does not address inconsistency in changes to code clones but inconsistency among licenses under which source files including code clones are distributed.

In addition, many studies in software engineering investigated software license. Some approaches for software license identification are proposed [7, 10, 23]. Using these approaches, some researches analyzed software licenses in open source projects and revealed some license issues. Di Penta et al. [4] provided an automatic method to track changes occurring in the licensing terms of a system and did an exploratory study on license evolution in six open source systems and explained the impact of such evolution on the projects. German et al. [8] proposed a method to understand licensing compatibility issues in software packages. They mainly focused on the compatibility between license declared in packages and those in source files. In another research by Di Penta et al. [9], they analyzed license inconsistencies of code siblings (a code clone that evolves in a different system than the code from which it originates) between Linux, FreeBSD and OpenBSD, but they did not explain the reasons underlying these inconsistencies. Alspaugh et al. [1] proposed an approach for calculating conflicts between licenses in terms of their conditions. However, our work proposed an approach to find license inconsistencies in similar files. By investigating the revision history of these files, we summarized the factors that caused these license inconsistencies and tried to decide whether they are legally safe or not. Recently Vendome et al. [24] performed a large empirical study of Java applications and found that changing license is a common event and a lack of traceability between when and why the license of a system changes.

7 Conclusion and Future Work

This paper describes and categorizes different types of license inconsistencies, some of which might lead to potential license violations. We also proposed a method to identify files that might have license inconsistencies. With the proposed method, we managed to detect all these types of license inconsistency from two data sets of open source projects: a Linux distribution Debian 7.5 and Java projects selected from GitHub. These results show the existence of license inconsistency in open source projects and proves the feasibility of our method.

With a manual analysis on some license inconsistency cases, we discovered that there are several reasons behind license inconsistencies: in some cases the copyright owner changed the license statement; sometimes the reuser exercise the permission that the file license gave her to remove one or more licenses from the file; in other cases, the reuser added another license to the file, and finally, the reuser modified the license. Among them, the last two categories are potentially unsafe and require further investigation.

In the process of our manual analysis, we came across a great difficulty to find out the reason behind each license inconsistency case. On one hand, it is difficult to find out from where a certain file in a project is imported when lacking enough information. On the other hand, it is also not a trivial task to decide which file is the original work when they are found in multiple projects. These problems highlight the need for a method to find and maintain the provenance between applications.

For future work, we will apply our tool to large numbers of open source projects and examine the proportion of each type of license inconsistency. With the increased number of projects, we believe that much more license inconsistency cases will be found. And we will try to make a quantitative evaluation of this tool. Furthermore, we will try to develop a method to help us analyze the history of each file, so that we can decide the safety of these inconsistencies efficiently.

謝辞

Firstly, I would like to express my sincere gratitude to my supervisor Prof. Katsuro Inoue, who gave me this precious opportunity to study in Japan. He continuously helps me and leads me to the right direction whenever I have any concerns with my research or writing this thesis. It would not have been possible for me to finish my this thesis without his support and help.

Besides my supervisor, I would like to thank the rest of the members in my research group: Prof. Daniel M. German, Prof. Yuki Manabe and Dr. Tetsuya Kanda, for their numerous invaluable suggestions and comments in my research.

I would also like to thank the rest of the members in my laboratory, who helped me a lot in learning Japanese and getting used to the life in Japan. I would give my special thanks to Mrs. Mizuho Karube, who is always ready to support me whenever I meet with any trouble in my stay in Japan.

Last but not least, I would like to thank my family for their constant support of my study and living in Japan.

参考文献

- [1] T.A. Alspaugh, H.U. Asuncion, and W. Scacchi. Intellectual property rights requirements for heterogeneously-licensed systems. In *Proceedings of the 17th International Requirements Engineering Conference (RE2009)*, pp. 24–33, 2009.
- [2] Nicolas Bettenburg, Weyi Shang, W. Ibrahim, B. Adams, Ying Zou, and A.E. Hassan. An empirical study on inconsistent changes to code clones at release level. In *Proceedings of the 16th Working Conference on Reverse Engineering (WCRE2009)*, pp. 85–94, 2009.
- [3] Barry W. Boehm. Improving software productivity. *Computer*, Vol. 20, No. 9, pp. 43–57, September 1987.
- [4] Massimiliano Di Penta, Daniel M. German, Yann-Gaël Guéhéneuc, and Giuliano Antoniol. An exploratory study of the evolution of software licensing. In *Proceedings of the 32nd International Conference on Software Engineering (ICSE2010)*, pp. 145–154, 2010.
- [5] Mark Gabel, Junfeng Yang, Yuan Yu, Moises Goldszmidt, and Zhendong Su. Scalable and systematic detection of buggy inconsistencies in source code. In *Proceedings of the 25th International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA2010)*, pp. 175–190, 2010.
- [6] Daniel M German and Ahmed E Hassan. License integration patterns: Addressing license mismatches in component-based development. In *Software Engineering, 2009. ICSE 2009. IEEE 31st International Conference on*, pp. 188–198. IEEE, 2009.
- [7] Daniel M German, Yuki Manabe, and Katsuro Inoue. A sentence-matching method for automatic license identification of source code files. In *Proceedings of the 25th International Conference on Automated Software Engineering (ASE2010)*, pp. 437–446, 2010.
- [8] D.M. German, M. Di Penta, and J. Davies. Understanding and auditing the licensing of open source software distributions. In *Proceedings of the 18th International Conference on Program Comprehension (ICPC2010)*, pp. 84–93, 2010.
- [9] D.M. German, M. Di Penta, Y.-G. Gueheneuc, and G. Antoniol. Code siblings: Technical and legal implications of copying code between applications. In *Proceedings*

- of the 6th Working Conference on Mining Software Repositories (MSR2009), pp. 81–90, 2009.
- [10] Robert Gobeille. The FOSSology project. In *Proceedings of the 5th Working Conference on Mining Software Repositories (MSR2008)*, pp. 47–50, 2008.
- [11] Nils Göde and Jan Harder. Oops! . . . I changed it again. In *Proceedings of the 5th International Workshop on Software Clones (IWSC2011)*, pp. 14–20, 2011.
- [12] Nils Göde and Rainer Koschke. Frequency and risks of changes to clones. In *Proceedings of the 33rd International Conference on Software Engineering (ICSE2011)*, pp. 311–320, 2011.
- [13] Yoshiki Higo and Shinji Kusumoto. MPAnalyzer: A tool for finding unintended inconsistencies in program source code. In *Proceedings of the 29th International Conference on Automated Software Engineering (ASE2014)*, pp. 843–846, 2014.
- [14] Toshihiro Kamiya, Shinji Kusumoto, and Katsuro Inoue. CCFinder: A multilingual token-based code clone detection system for large scale source code. *IEEE Transactions on Software Engineering*, Vol. 28, No. 7, pp. 654–670, 2002.
- [15] Jens Krinke. A study of consistent and inconsistent changes to code clones. In *Proceedings of the 14th Working Conference on Reverse Engineering (WCRE2007)*, pp. 170–178, 2007.
- [16] Jingyue Li, R. Conradi, C. Bunse, Marco Torchiano, O. Slyngstad, and Maurizio Morisio. Development with off-the-shelf components: 10 facts. *IEEE Software*, Vol. 26, No. 2, pp. 80–87, March 2009.
- [17] Yuki Manabe, DanielM. German, and Katsuro Inoue. Analyzing the relationship between the license of packages and their files in free and open source software. In *Proceedings of the 10th International Conference on Open Source Systems (OSS2014)*, pp. 51–60, 2014.
- [18] Yuki Manabe, Yasuhiro Hayase, and Katsuro Inoue. Evolutional analysis of licenses in FOSS. In *Proceedings of the Joint ERCIM Workshop on Software Evolution and International Workshop on Principles of Software Evolution (IWPSE-EVOL2010)*, pp. 83–87, 2010.

- [19] M Douglas McIlroy, JM Buxton, Peter Naur, and Brian Randell. Mass-produced software components. In *Proceedings of the 1st International Conference on Software Engineering (ICSE1968)*, pp. 88–98, 1968.
- [20] Chanchal K Roy, James R Cordy, and Rainer Koschke. Comparison and evaluation of code clone detection techniques and tools: A qualitative approach. *Science of Computer Programming*, Vol. 74, No. 7, pp. 470–495, 2009.
- [21] Yusuke Sasaki, Tetsuo Yamamoto, Yasuhiro Hayase, and Katsuro Inoue. Finding file clones in FreeBSD ports collection. In *Proceedings of the 7th Working Conference on Mining Software Repositories (MSR2010)*, pp. 102–105. IEEE, 2010.
- [22] Thomas A. Standish. An essay on software reuse. *IEEE Transactions on Software Engineering*, Vol. SE-10, No. 5, pp. 494–497, Sept 1984.
- [23] Timo Tuunanen, Jussi Koskinen, and Tommi Kärkkäinen. Automated software license analysis. *Automated Software Engineering*, Vol. 16, No. 3-4, pp. 455–490, 2009.
- [24] Christopher Vendome, Mario Linares-Vásquez, Gabriele Bavota, Massimiliano Di Penta, Daniel M. Germán, and Denys Poshyvanyk. License usage and changes: A large-scale study of java projects on github. In *The 23rd IEEE International Conference on Program Comprehension, ICPC 2015*, 2015.
- [25] Yuhao Wu, Yuki Manabe, Tetsuya Kanda, Daniel M German, and Katsuro Inoue. A method to detect license inconsistencies in large-scale open source projects. In *Proceedings of the 12th Working Conference on Mining Software Repositories (MSR2015)*, pp. 324–333, 2015.