

# Characterizing and Understanding Software Developer Networks in Security Development

Song Wang\* and Nachi Nagappan<sup>§</sup>

\*York University; <sup>§</sup>Microsoft Research  
wangsong@eecs.yorku.ca, nachin@microsoft.com

*Abstract*—To build secure software, developers often work together during software development and maintenance to find, fix, and prevent security vulnerabilities. Examining the nature of developer interactions during their security activities regarding security introducing and fixing activities can provide insights for improving current practices.

In this work, we conduct a large-scale empirical study to characterize and understand developers’ interactions during their security activities regarding security introducing and fixing, which involves more than 16K security fixing commits and over 28K security introducing commits from nine large-scale open-source software projects. For our analysis, we first examine whether a project is a hero-centric project when assessing developers’ contribution in their security activities. Then we study the interaction patterns between developers, explore how the distribution of the patterns changes over time, and study the impact of developers’ interactions on the quality of projects. In addition, we also characterize the nature of developer interaction in security activities in comparison to developer interaction in non-security activities (i.e., introducing and fixing non-security bugs).

Among our findings we identify that: most of the experimental projects are non hero-centric projects when evaluating developers’ contribution by using their security activities; there exist common dominating interaction patterns across our experimental projects; the distribution of interaction patterns has correlation with the quality of software projects. We believe the findings from this study can help developers understand how vulnerabilities originate and fix under the interactions of software developers.

*Index Terms*—security analysis, social network analysis, developer network, developer interaction

## I. INTRODUCTION

Building reliable and security software becomes more and more challenging in modern software development. As vulnerabilities can have catastrophic and irreversible impacts, e.g., the recent Heartbleed (CVE-2014-0160) cost more than US\$500 million to the global economy [1].

Developing secure software is a team effort, developers work together to find, fix, and prevent security vulnerabilities and during which they form implicit collaborative developer networks [2]–[18]. Understanding the structure of developer interaction in security assurance practices can be helpful for building more secure software. Along this line, many developer network-related analyses have been proposed to deal with problems in real-world security practice such as vulnerabilities prediction [2], [7], exploring the impact of human factors on security vulnerabilities [3], [5], [19], and monitoring vulnerabilities [10], [12].

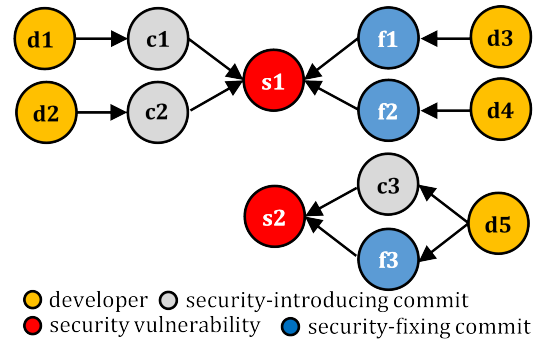


Fig. 1: An example developer security network.

Most of the existing approaches to exploring developer cooperation in security activities construct developer social network based on a single type of developer interaction, e.g., developers have co-changed/co-commented files that contain security vulnerabilities [2], [3], [5], [7], [10], [12], [19]. However, security vulnerabilities are introduced and fixed by developers. During the life cycle of a security vulnerability, developers interact with each other via multiple ways. For example, as shown in Figure 1, developers d1 and d2 introduced the security vulnerability s1 via commit c1 and c2; s1 was later fixed by developer d3 and d4 via commit f1 and f2. The security vulnerability s2, was introduced via commit c3 and fixed via commit f3 by the same developer d5. Examining the nature of developer interactions during their security activities including both introducing and fixing security vulnerabilities can provide insights for improving current security practices.

In this paper, we propose the first study to characterize and understand developers’ interactions in introducing and fixing security vulnerabilities by analyzing the developer networks built on their security activities. Our experiment dataset involves more than 16K security fixing commits and over 28K security introducing commits that ever appeared in nine large-scale open-source software projects including operation systems, compilers, PHP interpreter, Android platform, and JavaScript engine, etc. For our analysis, we first examine the heroism of software project when assessing developers’ contribution in developers’ security activities. As recent studies [20]–[25] showed that most software projects are hero-centric projects where 80% or more of the contributions (i.e., number of commits) are made by around

20% of the developers. Then we explore whether there exist dominating interaction patterns between developers across our experimental projects, after that we study how the distribution of developers' interaction patterns changes in different projects over time. Finally we explore the potential impact of the developer interaction patterns on the quality of software projects by measuring the correlation between the changes of developer interactions and security density (i.e., dividing the number of security vulnerability by the number of submitted commits) in a given period of time. In addition, we also characterize the nature of developer interaction in security activities in comparison to developer interaction in non-security activities (i.e., introducing and fixing non-security bugs).

This paper makes the following contributions:

- We conduct the first study to analyze developer interactions in security networks built on developers' security activities including both introducing and fixing security vulnerabilities.
- We confirm that all experimental projects are hero-centric projects when assessing developers' contribution with non-security activities. However, we also find that most (eight out of nine) experimental projects are non hero-centric projects when assessing developers' contribution by using security activities.
- We show that there exist dominating interaction patterns in both security and non-security activities across our experimental projects, while the distribution of developers' interaction in security and non-security activities are significantly different.
- We examine that developers' interaction is correlated with the quality of a software project regarding security vulnerability density.

The rest of this paper is organized as follows. Section II presents the background. Section III describes the methodology of our approach. Section IV shows the experimental setup. Section V presents the evaluation results. Section VI discusses the threats to the validity of this work. Section VII presents related studies. Section VIII concludes this paper.

## II. BACKGROUND

### A. Version-Control Systems

Version-control systems (VCS) are widely used in modern software development to coordinate developers' incremental contributions to a common software system. A VCS stores the entire source-code change history in the form of atomic change sets, called commits, which contain information about the changed code, the committers, and the timestamp of commits, etc. Git is one of the most popular VCSs, which has been adopted by more than 57M open-source projects and used by more than 20M developers<sup>1</sup> globally. Git's unique features make it especially appropriate for mining invaluable information to better understand software process [26], [27]. For example, Git can track the history of lines as they are

modified. By using the `git blame` feature, we can track the modification history of each line in a commit.

In this work, we collect software security history data from nine projects that are maintained by Git to explore the developer interaction structures during their security activities, details are showed in Section III.

### B. Developer Security Network

Developers interactions during their security activities including security fixing and introducing enable us to identify collaborative relationships between developers. The developer relationships can be described by a network, in which nodes represent developers and edges represent interactions between developers, in which nodes represent developers and edges represent interactions between developers.

In this study, a network can be formalized as a graph  $G = (V, E)$ , where  $V$  is a set of vertices and  $E$  is a set of edges, denoted by  $V(G)$  and  $E(G)$ , respectively. An edge  $e \in E$  is denoted as  $e = v, u$ , where  $v$  is the origin node and  $u$  is the destination node from  $V$ . Graph edges are directed with different meanings.

Different from most of existing developer social network studies [28]–[65], in which  $v \in V$  is a developer, and  $e \in E$  represents a particular form of developer interactions, e.g., fixed bugs together [40], [43], [48], [66], co-changed files [2], [3], [5], [7], [10], [12], [19], [38], [44], worked on the same project [46], or have communicated via email [50], etc., we consider a  $v \in V$  in a developer security network may have three different types, i.e., developer, security-fixing commit, and security-introducing commit. Consequently, a  $e \in E$  has also have three different types of meanings, i.e., a developer introduces a security vulnerability via a security-introducing commit, a developer fixes a security vulnerability via a security-fixing commit, a security-fixing commit fixes the vulnerability introduced by a security-introducing commit.

## III. DATA COLLECTION METHODOLOGY

### A. Subject Projects

We selected nine open-source projects from existing studies [23], [28], [67]–[69], listed in Table I, to explore developer interaction in security activities. The projects vary by the following dimensions: (a) size (lines of source code from 20 KLOC to over 17 MLOC, number of developers from 604 to 19K), (b) age (days since first commit), (c) programming language (C/C++, Java, PHP, and JavaScript), (d) application domain (operating system, compiler, PHP interpreter, Android platform, and JavaScript engine, etc.), and (e) VCS used (Git, Subversion). For each project, we extracted its code repository, and all the historical code commits hosted in GitHub on Nov. 5th 2018. Details of our approach to collecting the commits that introduce or fix security vulnerabilities and non-security bugs are as follows.

### B. Finding Public Vulnerabilities

#### 1) Collecting Security Vulnerability Fixing Commits:

Our data collection of security vulnerability fixing commits

<sup>1</sup><https://en.wikipedia.org/wiki/GitHub>

TABLE I: Experimental projects in this study. **Dev** is the number of developers. **Fix** is the number of commits that fixed security or non-security issues. **Intro** is the number of commits that introduced security or non-security issues.

Project	Language	LastCommitDate	#Commit	#Dev	#CVE	Security Vulnerability			Non-Security Bugs		
						Fix	Intro	Dev	Fix	Intro	Dev
FFmpeg	C/C++	2018/11/05	92,349	1,713	308	810	1,007	199 (11.62%)	16,024	26,592	1,138 (66.43%)
Freebsd	C/C++	2018/11/05	255,969	766	341	2,640	4,086	386 (50.39%)	35,776	66,490	604 (78.85%)
Gcc	C/C++	2018/11/05	165,475	604	6	575	1,341	200 (33.11%)	15,836	29,975	506 (83.77%)
Nodejs	JS	2018/11/05	24,401	2,640	48	252	402	105 (3.98%)	4,792	10,571	1,302 (49.32%)
Panda	C/C++	2018/11/05	52,580	1,220	24	557	1,072	230 (18.85%)	9,133	17,125	838 (68.69%)
Php	C/C++	2018/11/05	109,461	911	588	979	1,292	165 (18.11%)	25,610	48,296	663 (72.78%)
Qemu	C/C++	2018/11/05	64,840	1,459	261	789	1,459	263 (18.03%)	12,139	23,954	1,023 (70.12%)
Linux	C/C++	2018/11/05	796,003	19,362	2,207	10,316	17,126	3,686 (19.04%)	174,687	313,804	14,046 (72.54%)
Android	Java	2018/11/05	377,801	2,938	1,763	2,439	2,521	496 (16.88%)	70,157	128,893	2,132 (72.57%)

starts from the National Vulnerability Database (NVD) [70], a database provided by the U.S. National Institute of Standards and Technology (NIST) with information pertaining to publicly disclosed software vulnerabilities. NVD contains entries for each publicly released vulnerability. These vulnerabilities are identified by CVE (Common Vulnerabilities and Exposures) IDs [71]. When security researchers or vendors identify a vulnerability, they can request a CVE Numbering Authority to assign a CVE ID to it. Upon public release of the vulnerability information, the summarization the vulnerability, links to relevant external references (such as security fixing commits and issue reports), list of the affected software, etc., will be added to the CVEs. We first extracted all the public CVEs of each experimental subject on Nov. 5th 2018. We then crawled the Git commit links to identify and clone the corresponding Git source code repositories and collected security fixes using the commit hashes in the links. Note that, we also find that some of the external references only contain the bug/issue report links, e.g., the external reference of security vulnerability CVE-2018-14609<sup>2</sup> does not contain the security fixing commits instead it shows the bug report ID<sup>3</sup>. For these security vulnerabilities, we used the fixing commits of these bugs as the security fixing commits. To collect the fixing commits of these bugs, we consider commits whose commit messages contain the bug report ID as the fixing commits by following existing studies [69], [72].

As reported in existing studies [73], [74], not all security vulnerability have CVE identifiers, around 53% of vulnerabilities in open source libraries are not disclosed publicly with CVEs [75], [76]. To cover all possible vulnerabilities, we used the heuristical approaches proposed by Zhou et al. [75], to identify the security fixing commits. Specifically, we used the regular expression rules listed in their Table 1, which included possible expressions and keywords related to security issues.

2) *Grouping Security Fixing Commits*: In the above section we have described how to collect security fixing commits. We found that some of the security fixing commits are made for fixing the same security vulnerability. For example, to fix security vulnerability CVE-2018-10883<sup>4</sup>, developers have made

---

### Algorithm 1 Grouping fixing commits algorithm

---

**Require:**

- Fixing commit set  $C$ ;
- Query fixing commit  $q$ ;
- Commit message similarity threshold  $thres_s$ ;
- Fixing location overlap rate threshold  $modif_o$ ;

**Ensure:**

- A list of grouped fixing commit  $D$ ;
  - 1: **for** each commit  $r$  in  $C$  and  $q$  **do**
  - 2:   Extract commit messages and compute the similarity  $message_s$ ;
  - 3:   Extract modified files and compute the overlap rate  $modif_o$ ;
  - 4:   **if**  $message_s > thres_s$  and  $modif_o > 0$  **then**
  - 5:     put  $r$  in  $D$
  - 6:   **end if**
  - 7: **end for**
- 

two commits. Identifying fixing commits that belong to the same security vulnerability could provide us valuable information about how vulnerabilities are fixed through developer interactions. To group fixing commits, first, for fixing commits that have CVE identifiers in their commit messages, we consider fixing commits that contain the same CVE identifiers belong to the same security vulnerabilities. Second, for fixing commits that do not have CVE identifiers in their commit messages, we propose a heuristical algorithm to group them, which is described in Algorithm 1. Specifically, given two fixing commits, we group them together if the similarity of their commit messages is larger than a threshold (i.e.,  $message_s$ ) and the modification location has overlaps. Following existing study [77]–[80], we use the Cosine similarity to measure the similarity between two commit messages. We employ tf-idf [81], stop words removal (e.g., “is”, “are”, and “in” since these words are used in most commit messages and thus have little discriminative power) and stemming (e.g., “groups” and “grouping” are reduced to “group”.) to extract string vectors from the commit messages. For the threshold  $thres_s$ , we assume the ratios of collaborative fixing commits (i.e., fixing the same vulnerability) are similar between commits which have CVEs and commits that do not have CVEs. Thus for each project, we use the ratio of the collaborative fixing commits among the fixing commits that have CVEs to specify its threshold  $thres_s$ . In addition, we set the maximum interval between two collaborative fixing commits as six months, which is the typical length of fixing a security vulnerability [82].

<sup>2</sup><https://nvd.nist.gov/vuln/detail/CVE-2018-14609>

<sup>3</sup>[https://bugzilla.kernel.org/show\\_bug.cgi?id=199833](https://bugzilla.kernel.org/show_bug.cgi?id=199833)

<sup>4</sup><https://nvd.nist.gov/vuln/detail/CVE-2018-10883>




Fig. 2: The ratios (in percentage) of collaborative fixing commits grouped from security fixing and non-security fixing commits.

### 3) Collecting Security Vulnerability Introducing Commits:

With the above security-fixing commits, we further identify the security-introducing commits by using a blame technique provided by a Version Control System (VCS), e.g., git or SZZ algorithm [72]. Following existing studies [83]–[86], we assume the deleted lines in a security-fixing commit are related to the root cause and considered as faulty lines. The most recent commit that introduced the faulty line is considered a security-introducing commit.

Note that, different from security-fixing commits, we did not group security-introducing commits. This is because the security-introducing commits are identified by security-fixing commits. Since we have already grouped security-fixing commits, these security-introducing commits are grouped accordingly. The details of the security-introducing commits as listed in Table I. The average number of security-introducing commits of a security-fixing commit ranges from 1.03 (Android) to 2.41 (Nodejs).

### C. Finding Non-Security Bugs

In addition, to explore the difference of developer interaction structures between developers’ security activities and non-security activities, we also collect general bugs (i.e., non-security).

Typically software bugs are discovered and reported to an issue tracking system such as Bugzilla and later on fixed by the developers. A bug report usually records the description, the opening and fixing date, type (bug, enhancement, feature, etc.), etc. We consider a bug report in the Bugzilla database that is labelled as a “bug” to be a general bug. However, not all the projects have well-maintained bug tracking systems, in this work, following existing studies [84]–[86] if a projects bug tracking system is not well maintained and linked, we consider changes whose commit messages contain the word “fix” and “bug” as bug-fixing commits. If a projects bug tracking system is well maintained and linked, we consider commits whose commit messages contain a bug report ID as bug-fixing commits. For each of the bug-fixing commit, we adopt the same approach as we used to identify security-introducing commits in Section III-B3. Note that, if any of the non-security fixing commit appears in the security-fixing commit dataset, we will remove it from the non-security fixing commit dataset. The details of non-security fixing commits and their corresponding non-security introducing commits are showed in Table I. The average number of non-security introducing commits of a non-security fixing commit ranges from 1.66 (FFmpeg) to 2.21 (Nodejs).




Fig. 3: The overlap rate (in percentage) of non-security introducing commits and security introducing commits.

In Section III-B2, we group security-fixing commits that fix the same security vulnerability. For non-security bugs, we also found the same phenomenon, i.e., some of the non-security fixing commits are made for fixing the same non-security bugs. For grouping these non-security fixing commits, we reuse Algorithm 1. As described in In Section III-B2, for grouping security fixing commits, we use the ratio of collaborative fixing commits (i.e., fix the same security vulnerability) that have CVE identifiers to set the threshold  $thres_s$  of a specific project. However, for non-security fixing commits, not all projects have well-maintained bug tracking systems, for some projects (e.g., Linux), we cannot use bug report ID to specify  $thres_s$ . Thus, we randomly pick and manually check 100 pairs of collaborative fixing commits on each the subject project, we use the average Cosine similarity value to set  $thres_s$  in Algorithm 1 to group non-security fixing commits.

Figure 2 shows the ratios of collaborative fixing commits in the security fixing and non-security fixing commits. As we can see from the table, the ratios in non-security fixing commits are slightly higher than that of security fixing commits in most projects. On average, the ratio for security fixing commits is 7.2% and the ratio for non-security fixing commits is 8.3%, which is consistent with the finding from an existing study [87], that 9% of bug fixes were bad across three Java projects.

As we mentioned above, we have removed the non-security fixing commits from the security fixing commit dataset, while for non-security introducing commits and security introducing commits, we do not handle the overlaps, since it’s possible that a security vulnerability and non-security bug can be introduced by the same introducing commit. In this work, we use overlap rate to measure the overlap level between two datasets. We define the **overlap rate** between datasets  $A$  and  $B$  as  $\frac{A \cap B}{A \cup B}$ . Figure 3 shows the overlap rates of non-security introducing commits and security introducing commits in the experimental projects. As we can see from the figure, the overlap rates of all experimental projects are lower than 10%, which suggests that security vulnerability and non-security bugs usually have different introducing commits.

### D. Identifying Distinct Developers

To build the developer security network, we need to obtain the developer information of security-fixing and security-introducing commits. In Git, for every pushed commit, Git maintains the user who did the commit, i.e., committer. Git

computes the committer out of the Git configuration parameters ‘user.name’ and ‘user.email’. Thus, by retrieving a commit, we can easily obtain its committer information. However, Git also allows users to change their profiles, which introduces the alias issue of developers in mining open-source [50], [88], i.e., a developer may have different emails/names. To solve this challenge, we use the aliases unmasking algorithms proposed in [50] to identify distinct developers.

In total, we have around 45K distinct developers from the nine experimental projects, details are listed in Table I. Overall, the percentage of developer that involved in security activities ranges from 3.98% to 50.39%, while the percentage of developer that involved in non-security activities ranges from 49.32% to 83.77%.

#### IV. RESEARCH QUESTION

Our experimental study is designed to answer the following research questions.

##### **RQ1. What are the distributions of developers in security and non-security activities?**

Software security vulnerability and bugs are introduced and fixed by developers, in this RQ, we aim to explore the basic distribution of developers in security and non-security activities regarding fixing and introducing. For example, what is the overlap rate between developers that have ever involved in security activities and developers that have ever involved in non-security activities? What is the overlap rate between developers that have fixed security vulnerabilities and developers that have introduced security vulnerabilities?

##### **RQ2. How common are hero-centric projects regarding software security activities?**

Recent studies [20]–[25] show that most software projects are hero-centric projects where 80% or more of the contributions (e.g., the number of commits) are made by the 20% of the developers. While the above studies assess developers’ contribution from broad aspects, e.g., Agrawal et al. [20] used the number of commits made by each developer to represent its contribution to a project. Majumder et al. [21] built a social interaction graph from developers’ communication and used the node degree to represent a developer’s contribution. Most of existing studies explore the heroism of projects from developers’ code contribution and social communication perspectives. In this RQ, we aim to explore the heroism of a project when assessing developers’ contribution by using a specific type of commits, e.g., security fixing, security introducing, non-security fixing, non-security introducing.

##### **RQ3. What are the common interaction patterns between two developers in security activities?**

Developers interact with each other during the development of a software project. In software development, the social and organizational aspects have an impact on the individual and collective performance of the developers [89]. Along this line, in this RQ, we aim to explore the common interaction structures among developers during their security and activities

regarding security fixing and security introducing across different projects, which we believe can help us gain insight into distinct characteristics of developers’ security activities.

##### **RQ4. Are the distributions of developer interaction patterns in security and non-security activities different?**

In this RQ, we aim to characterize the nature of developer interaction in security vulnerabilities in comparison to other non-security bugs. Specifically, we compare the distributions of interaction structures among developers in security vulnerabilities and non-security bugs.

##### **RQ5. How do interaction structures among developers evolve over time?**

Software team organization evolves over time [40], [41], i.e., developers may leave a project and new developers may join during the life cycle of a project, which causes the evolution of developer community. Along this line, in this RQ, we aim to explore whether the interaction structure among developers changes over time and how it evolves during the life cycle of a project.

##### **RQ6. Does the change of interaction structures have an impact on the quality of software?**

Developer social network and its evolution information have been examined could be used to predict new vulnerabilities and bugs [2], [7]. Along this line, in this RQ, we investigate whether the change of interaction structure has a correlation with the quality of software regarding the density of security vulnerabilities.

#### V. ANALYSIS APPROACH AND RESULTS

##### *A. RQ1: Distributions of Developers in Security and Non-Security Activities*

To answer this RQ, we obtain unique developers from different activities, i.e., fixing security vulnerabilities, introducing security vulnerabilities, fixing non-security bugs, and introducing non-security bugs. Given the developer sets of two activities, we calculate their overlap rates via dividing the overlapping data points by all the unique data points. Table II shows the basic overlaps between developers that have been involved in different activities. As we can see from the table, in all the projects, developers from **secFix** and **secIntro** have higher overlap rates, i.e., range from 60.0% to 89.0% and on average is 71.1%, which indicates that most of the security vulnerabilities are introduced and fixed by a core group of developers. We can also see that the overlap rates of developers from security activities and non-security activities are lower, e.g., the overlap rate of developers from **secFix** and **nonSecFix** ranges from 5.0% to 30.2% and is 16.9% on average, the overlap rate of developers from **secIntro** and **nonSecIntro** ranges from 19.6% to 38.4% and on average is 28.8%. Overall, the overlap rate from **sec** and **nonSec** is 18.6% on average, which indicates that most of the developers that are involved in security activities are different from developers that are involved in non-security activities. This may be because security issues are critical to software that require non-trivial domain expertise. Thus only a small group of developers is capable

TABLE II: The overlap rates between developers that have been involved in different activities. **secFix** denotes developers that have made security fixing commits, **secIntro** denotes developers that have made security introducing commits, **nonSecFix** denotes developers that have made non-security fixing commits, **nonSecIntro** denotes developers that have made non-security introducing commits, **sec** denotes developers that have made security fixing or introducing commits, and **nonSec** denotes developers that have made non-security fixing or introducing commits. **secFix-secIntro** means the overlap rate between **secFix** and **secIntro**. The higher values with statistical significance ( $p$ -value  $< 0.05$ ) are shown with an asterisk (\*).

Project	secFix-secIntro (*)	secFix-nonSecFix	secFix-nonSecIntro	secIntro-nonSecFix	secIntro-nonSecIntro	sec-nonSec
FFmpeg	60.0	10.7	10.4	14.9	19.6	10.7
Freebsd	89.0	30.2	32.2	49.1	43.2	40.2
Gcc	88.1	25.6	24.5	38.8	38.4	25.6
Nodejs	63.0	5.0	5.8	7.6	10.5	5.0
Panda	65.9	17.0	18.9	21.5	32.1	17.0
Php	70.5	15.5	16.9	21.7	29.3	15.5
Qemu	67.1	16.6	18.1	20.6	28.9	16.6
Linux	66.6	16.1	18.0	21.6	29.6	16.1
Android	69.6	15.6	18.1	19.9	27.1	15.6
<b>Average</b>	71.1	16.9	18.1	24.0	28.8	18.0

TABLE III: The percentages of developers involved when contributing 80% of a specific type of commits. Values with a red diamond ( $\diamond$ ) indicate that a project is non hero-centric project. **All** denotes the combination of the four types of commits.

Project	secFix	secIntro	nonSecFix	nonSecIntro	All
FFmpeg	20.1 ( $\diamond$ )	17.1	3.6	5.5	3.5
Freebsd	32.1 ( $\diamond$ )	26.0 ( $\diamond$ )	13.4	11.3	11.1
Gcc	33.1 ( $\diamond$ )	21.1 ( $\diamond$ )	17.5	16.3	15.6
Nodejs	34.0 ( $\diamond$ )	24.7 ( $\diamond$ )	13.5	6.6	5.1
Panda	36.5 ( $\diamond$ )	25.8 ( $\diamond$ )	10.6	10.1	7.5
Php	21.6 ( $\diamond$ )	23.7 ( $\diamond$ )	6.3	8.2	5.7
Qemu	30.1 ( $\diamond$ )	22.3 ( $\diamond$ )	8.6	9.8	6.8
Linux	30.9 ( $\diamond$ )	21.6 ( $\diamond$ )	11.0	11.4	8.5
Android	32.7 ( $\diamond$ )	21.3 ( $\diamond$ )	11.5	11.0	8.3

of handling security vulnerabilities, which makes the overlap rates of developers from security activities and non-security activities lower. We further conduct the Wilcoxon signed-rank test ( $p < 0.05$ ) to compare the overlap rates among different pairs. The results suggest that the overlap rates of **secFix** and **secIntro** are significantly higher than those of other pairs.

Developers that are involved in security and non-security activities are different and have low overlap rates. For non-security activities, developers that introduced and fixed bugs have low overlap rates. However, for security activities, developers that introduced and fixed security vulnerabilities have higher overlap rates.

### B. RQ2: Heroism in Security and Non-Security Activities

Following existing studies [20], [21], we define a project to be hero-centric when 80% of the contributions are done by about 20% of the developers in this study. In addition, if the percentage of developers involved when contributing 80% of a specific type of commits is larger than 20%, we treat the project as non hero-centric projects.

In this RQ, we first examine whether a project is a hero-centric project with only considering a specific type of commits, e.g, security fixing, security introducing, non-security fixing, and non-security introducing. To assess the contribution

of a developer, following Agrawal et al. [20], we count the number of a specific type of commits (i.e., security fixing, security introducing, non-security fixing, non-security introducing) made by each developer to represent his/her contribution to a project. We then rank developers ascendingly based on their contributions. Finally, we accumulate developers' contributions and record developers involved until 80% of the contributions are done. In addition, we also evaluate a developer's contribution via the combination of the four types of commits.

Table III shows the percentages of developers involved when contributing 80% of a particular type of commits. As we can see from the table, all the projects are hero-centric projects, i.e., the percentages of developers involved are smaller than 20%, when assessing developers' contribution by using non-security fixing or non-security introducing commits or all commits together. However, most of the experimental projects are non hero-centric projects when assessing developers' contribution by using security fixing or security introducing commits, e.g., the percentage of developers involved are 36.5%, when evaluating developers' contribution by using security fixing commits in project Panda. Our finding indicates that although software development has "heroes", i.e., a small percentage of the staff who are responsible for most of the progress on a project, software security does not have typical "heroes".

We further calculate the overlap rates of "core developers" (i.e., contribute 80% of a specific type of commits) between different types of commits, which are shown in Table IV. Note that we use "core developers" since a project can be a non hero-centric project when assessing developers' contribution by using security activities. As we can see, the "core developers" from security fixing and security introducing have high overlap rates that range from 47.7% to 71.1% and on average is 63%, which is consistent with our findings in Sec V-A. The overlap rates of the "core developers" from security commits and non-security commits, i.e., "core developers" from **secFix** and **nonSecFix**, "core developers" from **secIntro** and **nonSecIntro** are lower than that of "core developers" only from security activities. This indicates that the "core developers" of

TABLE IV: The overlap rates of “core developers” (i.e., contribute 80% of a particular type of commits) between different types of commits. The higher values with statistical significance ( $p$ -value  $< 0.05$ ) are shown with an asterisk (\*).

Project	secFix-secIntro (*)	secFix-nonSecFix	secIntro-nonSecIntro
FFmpeg	71.1	50.0	68.9
Freebsd	69.9	55.6	67.8
Gcc	69.2	32.9	43.5
Nodejs	66.7	24.4	29.2
Panda	67.0	48.0	62.5
Php	48.5	54.3	66.7
Qemu	64.6	50.0	63.8
Linux	64.8	38.0	45.3
Android	47.0	34.3	45.8
Average	63.2	43.0	55.0

TABLE V: The distribution of developer interaction patterns during security activities (in percentage).

Project	CoIntro	CoFix	IntroFix	SelfIntroFix	SelfIntro	SelfFix
FFmpeg	52.3	1.9	36.6	5.2	3.0	1.0
Freebsd	66.1	0.2	28.8	3.1	1.6	0.1
Gcc	58.0	10.4	18.7	9.9	2.8	0.1
Nodejs	50.9	7.4	23.4	10.7	7.3	0.3
Panda	70.3	0.7	19.6	5.1	4.0	0.3
Php	73.1	1.4	18.0	3.8	3.3	0.4
Qemu	68.0	1.8	19.3	6.6	3.9	0.5
Linux	67.1	0.5	21.9	5.8	4.5	0.2
Android	55.4	8.1	25.3	3.2	7.4	0.6
Average	62.4	3.6	23.5	5.9	4.2	0.4

security and non-security activities are different in most of the experimental projects.

All the experimental projects are examined as hero-centric projects in non-security activities, while most of them (8 out of 9) are non hero-centric projects in security activities.

### C. RQ3: Common Developer Interaction Patterns in Developer Security Activities

In this RQ, we identify developer interactions during the security activities including both introducing and fixing security vulnerabilities. Specifically, in order to explore developer interactions, we capture three possible interactions between two developers, i.e., two developers introduce the same security vulnerability (CoIntro), two developers fix the same security vulnerability (CoFix), a security vulnerability is introduced by a developer and fixed by another developer (IntroFix), which are showed in Figure 4 from 4a to 4c. In addition, we also collect the interactions of a single developer, i.e., a security vulnerability is introduced and fixed by a single developer (SelfIntroFix), a security vulnerability is introduced by multiple commits of a single developer and fixed by other developers (SelfIntro), and a security vulnerability is fixed by multiple commits of a single developer and is introduced by other developers (SelfFix), which are showed in Figure 4 from 4d to 4f.

For each subject project, we first build a security activity network, then with these interaction patterns, we further collect the numbers and calculate the percentages of the six patterns, which are showed in Table V. As we can see from the

figure, the six developer interaction patterns exist in each of the experimental projects. The CoIntro and IntroFix patterns are dominating (i.e., the accumulated percentage is larger than 80%) across all the experimental projects. Other patterns take up around 20% of developer interactions, for example, the percentages of SelfFix are lower than 1% in all experimental projects. Although CoIntro and IntroFix are dominating, the percentages of them in different projects are different, i.e., range from 74.3% (Nodejs) to 94.9% (Freebsd). In addition, the percentage of interactions between developers (i.e., CoIntro, CoFix, and IntroFix) is much larger than that of interactions of the same developers (i.e., SelfIntro, SelfFix, and SelfIntroFix), which indicates the nature of software security development is teamwork.

The percentages of the developer interaction patterns vary dramatically in different projects. However, CoIntro and IntroFix patterns are dominating across all the experimental projects in developers’ security activities.

### D. RQ4: Comparison of Developer Interaction Patterns between Security and Non-Security Activities

In this RQ, we try to explore the difference of developer interactions between developers’ security activities and non-security activities, which we believe can help us gain insight into distinct characteristics of developers’ security activities. For each subject project, we first build a non-security activity network, then we further collect the ratios of the six patterns.

Table VI shows the distribution of the six developer interaction patterns in developers’ non-security activities. As we can see from the figure, although the six developer interaction patterns also exist in each of the experimental projects, the percentages of these patterns are different from that of security activities showed in Table V. In Figure 5, we show the detailed difference of interaction patterns between security activities and non-security activities. Specifically, the percentages of patterns CoIntro and CoFix, and SelfIntro vary dramatically across the experimental projects in this work.

Different from security activities, the dominating patterns (i.e., the accumulated percentage is larger than 80%) in non-security activities include three patterns, i.e., CoIntro, IntroFix, and CoFix. Note that in security activities, the percentage of CoFix pattern ranges from 0.2% to 10.4% and on average is 3.5%, while in non-security activities it ranges from 4.2% to 26.1% on average is 17.2%. This may indicate that security vulnerability fixing requires domain expertise than fixing non-security bug fixing and most developers are incapable to fix security vulnerabilities, thus results in less teamwork.

In addition, we also find that the dominating patterns are more balanced in developers’ non-security activities compared to security activities. For example, the difference of the percentages of dominating patterns in security activities ranges from 15.7% to 55.1% and on average is 38.8%, while in non-security activities, the difference ranges from 8.3% to 35.2% and on average is 21.2%.






Fig. 4: The potential interaction relationships between developers during their security activities.

TABLE VI: The distribution of developer interaction during non-security activities (in percentage).

Project	CoIntro	CoFix	IntroFix	SelfIntroFix	SelfIntro	SelfFix
FFmpeg	30.4	14.8	31.5	3.2	19.7	0.4
Freebsd	40.4	11.4	30.3	2.6	15.1	0.2
Gcc	37.3	26.1	22.4	2.4	11.7	0.2
Nodejs	59.0	26.3	9.7	1.5	3.4	0.1
Panda	39.4	4.2	35.6	3.9	16.6	0.3
Php	42.7	16.2	25.3	3.7	12.0	0.1
Qemu	35.4	19.7	28.4	4.5	11.9	0.1
Linux	32.1	15.5	33.0	3.5	15.8	0.1
Android	29.1	20.8	27.9	5.7	16.4	0.1
Average	38.4	17.2	27.1	3.4	13.6	0.2

TABLE VII: The correlated patterns in each project.

Project	Correlated Patterns
FFmpeg	P1, P3, P5
Freebsd	P1, P3, P4
Gcc	P1, P2, P3, P4
Nodejs	P1, P2, P3, P4
Panda	P1, P3
Php	P1, P3, P4, P5
Qemu	P1, P3, P5
Linux	P1, P3
Android	P1, P2, P3, P4, P5




Fig. 5: The difference of the percentages of interaction patterns between security activities and non-security activities.

Developers have different dominating patterns in security and non-security activities. In addition, the distribution of developers' interaction in security and non-security activities are different.

#### E. RQ5: Evolution of Developer Interaction in Developer Security Activities

To explore the evolution of developer interactions, for each project, we collect the numbers and calculate the percentages of the six patterns that only appear in a specific year from 2009 to 2018. Thus, for each pattern, we have 10 different percentage values in each project. Figure 6 shows the boxplots of the percentages of each interaction pattern in each project.

The figure shows that overall, the percentages of a pattern vary dramatically in a project over time, for example, in FFmpeg, the percentages of pattern CoIntro range from 22.7% to 63.1% in 10 years. Regarding the dominating patterns, we find that patterns CoIntro and IntroFix are dominating on each project over time. In addition, the same phenomenon is also observed in developers' non-security activities.

The percentages of developer interaction patterns vary over time. While all the projects do not witness a change in terms of the dominating patterns.

#### F. RQ6: Impact of Developer Interaction on Software Quality

To explore the relation between the changes of developers' interaction in security activity and the quality of the software, following existing studies [7], [90], we use the Spearman rank correlation [91] to compute the correlations between the percentages of patterns and the density of security vulnerability appeared in each year from 2010 to 2018. The closer the value of a correlation is to +1 (or -1), the higher two measures are positively (or negatively). A value of 0 indicates that two measures are independent. Values greater than 0.10 can be considered a small effect size; values greater than 0.30 can be considered a medium effect size [7]. In this work, we consider the values larger than 0.10 or smaller than -0.10 as correlated, others are uncorrelated.

Table VII shows the correlated patterns in each project. As we can see, five of the six patterns are selected as correlated in at least one project. In addition, the dominating patterns CoIntro and IntroFix are selected across all experimental projects.

Developers' interaction in security activities is correlated with the density of security vulnerabilities.

## VI. THREATS TO VALIDITY

a) *Internal Validity*: Threats to internal validity are related to experimental errors. Following previous work [72], [84]–[86], the process of collection security introducing or non-security introducing commits is automatically completed with the annotating or blaming function in VCS. It is known that this process can introduce noise [72]. The noise in the data can potentially affect the result of our study. Manual inspection of the process shows reasonable precision and recall on open



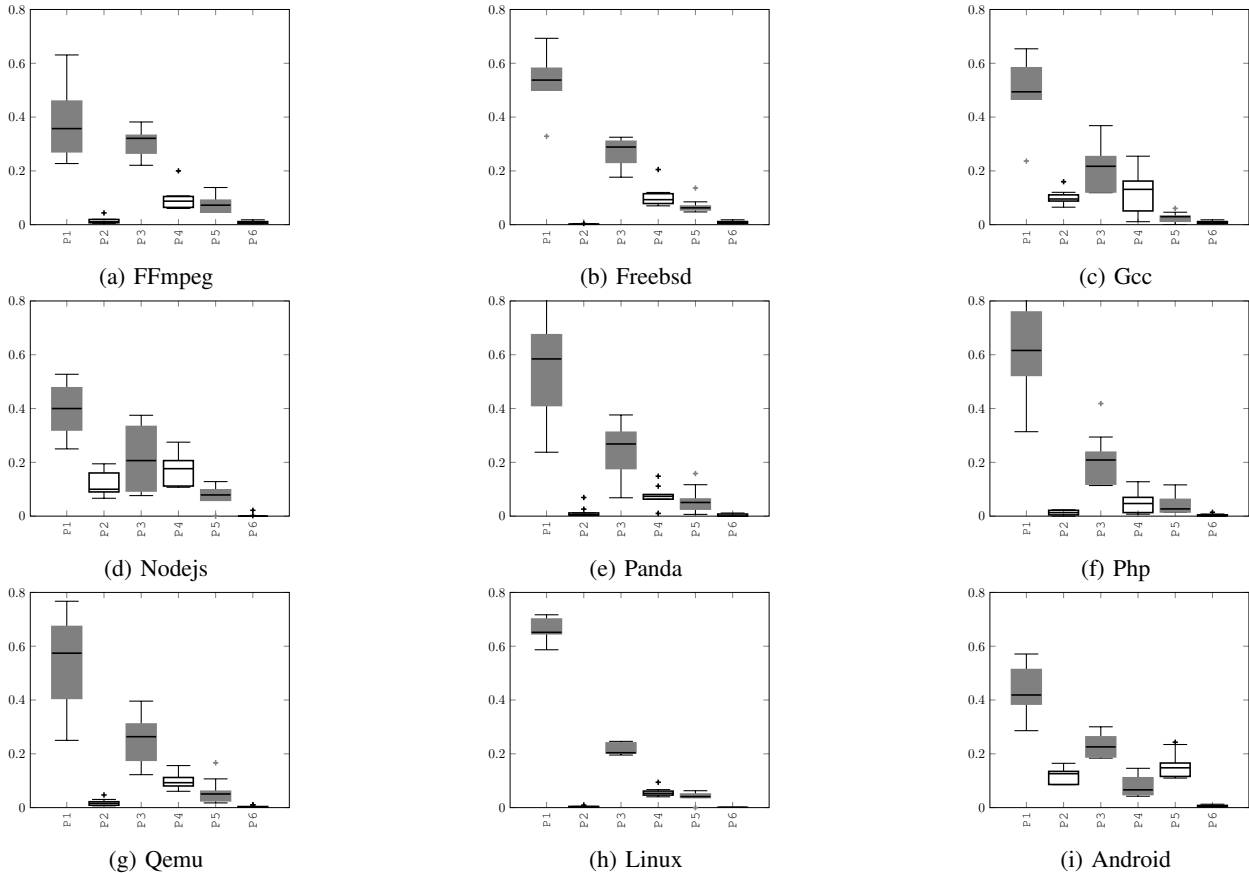


Fig. 6: The distribution of the number of different patterns in each project from 2009 to 2018. P1 denotes CoIntro, P2 denotes CoFix, P3 denotes IntroFix, P4 denotes SelfIntroFix, P5 denotes SelfIntro, and P6 denotes SelfFix.

source projects [86], [92]. To mitigate this threat, we use the noise data filtering algorithm introduced in [92].

*b) External Validity:* Threats to external validity are related to the generalization of our study. The examined projects in this work have a large variance regarding project types. We have tried our best to make our dataset general and representative. However, it is still possible that the nine projects used in our experiments are not generalizable enough to represent all software projects. Our approach might generate similar or different results for other projects that are not used in the experiments. We mitigate this threat by selecting projects of different functionalities (operating systems, servers, and desktop applications) that are developed in different programming languages (C, Java, and JavaScript).

In this work, all the experimental subjects are open source projects. Although they are popular projects and widely used in security research, our findings may not be generalizable to commercial projects or projects in other languages.

## VII. RELATED WORK

### A. Developer Social Network

There has been a body of work that investigated aspects of developer social networks built on developers' activities during software development [28]–[39], [39]–[65], [93].

Lopez-Fernandez et al. [56], [57] first examined the social aspects of developer interaction during development, where developers were linked based on contributions to a common module. Bird et al. [50] investigated developer organization and community structure in the mailing list of four open-source projects and used modularity as the community-significance measure to confirm the existence of statistically significant communities. Wolf et al. [16], [44] introduced an approach to mining developer collaboration from communication repositories and they further use developer collaboration to predict software build failures. Toral et al. [58] applied social-network analysis to investigate participation inequality in the Linux mailing list that contributes to role separation between core and peripheral contributors. Hong et al. [40] and Zhang et al. [30] explored the characteristics of developer social networks built on developers interactions in bug tracking systems and how these networks evolve over time. Surian et al. [42] extracted developer collaboration patterns from a large developer collaborations network extracted from SourceForge.Net, where developers are considered connected if both of them are listed as contributors to a project. Jeong et al. [48] and Xuan et al. [66] leveraged network metrics mined from social networks built in bug tracking systems

to recommend developers for fixing new bugs. Surian et al. [46] used developer collaboration network extracted from Sourceforge.Net to recommend a list of top developers that are most compatible based on their programming language skills, past projects and project categories they have worked on before for a developer to work with. Researchers have also built social networks based on developers' security activities, i.e., have co-changed files that contain security vulnerabilities to predict new vulnerabilities [2], [7], exploring the impact of human factors on security vulnerabilities [3], [5], [19], and monitoring vulnerabilities [10], [12].

Most of the above studies construct developer networks based on a particular form of developer collaboration e.g., co-changed files, co-commented bugs, and co-contributed projects, etc., from bug tracking systems, mailing lists, or project contribution lists. These developer networks are homogeneous, which have merely one type of node (developers) and one type of link (a particular form of developer collaboration). Wang et al. [43] and Zhang et al. [45] leveraged heterogeneous network analysis to mined different types of developer collaboration patterns in bug tracking system and further used these different collaborations to assist bug triage.

Our work differs in two ways from most of these prior studies: (1) We study developers' social interactions in security activities; (2) We explore different types of developer interactions during their security activities, which is more complex and with richer information.

### B. Security Vulnerability Analysis

There are many studies to explore, analyze, and understand software security vulnerabilities [82], [83], [94]–[106].

Frei et al. [98] examined how vulnerabilities are handled with regard to information about discovery date, disclosure date, as well as the exploit and patch availability date in large-scale by analyzing more than 80,000 security advisories published between 1995 and 2006. Walden et al. [104] provided a vulnerability dataset for evaluating the vulnerability prediction effectiveness of two modelling techniques, i.e., software metrics based and text mining based approaches. Medeiros et al. [105] examined the performance of software metrics on classifying vulnerable and non-vulnerable units of code. Yang et al. [106] leveraged software network to evaluate structural characteristics of software systems during their evolution. Decan et al. [97] and Shahzad [99] presented a large scale study of various aspects associated with software vulnerabilities during their life cycle. Ozment et al. [102] investigated the evolution of vulnerabilities in the OpenBSD operating system over time, observing that it took on average 2.6 years for a release version to remedy half of the known vulnerabilities. Perl et al. [83] analyzed Git commits that fixed vulnerabilities to produce a code analysis tool that assists in finding dangerous code commits. Xu et al. [103] developed a method for identifying security patches at the binary level based on execution traces, providing a method for obtaining and studying security patches on binaries and closed-source software.

Li et al. [82] conducted an analysis of various aspects of the patch development life cycle. There also existed some other studies that explored the characteristics of software general bugs [62]–[65], [100], [107].

In this work, we propose the first study to characterize and understand developers' interaction by considering their activities in introducing and fixing security vulnerabilities by analyzing developer networks built on their security activities.

### C. Heroism in Software Development

Heroism in software development is a widely studied topic. Various researchers have found the presence of heroes in software projects [20]–[25].

Koch et al. [22] studied the GNOME project and showed the presence of heroes throughout the project history. Krishnamurthy [24] conducted a case study on 100 projects and reported that a few individuals are responsible for the main contribution of the projects. Agarwal et al. [20] studied heroism in software development on 661 open source projects from Github and 171 projects from an Enterprise Github. They assess the contribution of a developer by the number of his/her commits submitted. Their experiment showed that 77% projects exhibit the pattern that 20% of the total contributors complete 80% of the contributions, which means hero-centric projects are very common in both public and enterprise projects. Majumder [21] studies the heroes developer communities in 1100+ open source GitHub projects. They built a social interaction graph from developers' communication and used the node degree to represent a developer's contribution. Based on the analysis, they found that hero-centric projects are majorly all projects.

The above studies explore the heroism in software development from developers' code contribution and social communication perspectives. In this work, we examine whether software projects are hero-centric projects when assessing developers' contribution in their security activities.

## VIII. CONCLUSION

This work conducts a large-scale empirical study to characterize and understand developers' interaction during developers' security activities including both security vulnerability introducing and fixing activities, which involves more than 16K security fixing commits and over 28K security introducing commits from nine large-scale open-source software projects. We first examine whether a project is a hero-centric project when assessing developers' contribution with developers' security activities. Then we examine the interaction patterns between developers in security activities, after that we show how the distribution of these patterns changes in different projects over time, finally we explore the potential impact of developers' interaction on the quality of projects by measuring the correlation between developers' interactions and the security density in a given period of time. In addition, we also characterize the nature of developer interaction in security activities in comparison to developer interaction in non-security activities (i.e., introducing and fixing non-security bugs).

Among our findings we identify that: most of the experimental projects are non hero-centric projects when assessing developers' contribution by using security activities; different projects have different dominating interaction structures; developers' interaction has correlation with the quality of software projects. We believe the findings from this study can help developers understand how vulnerabilities originate and fix under the interaction of developers.

## REFERENCES

- [1] "Heartbleed," <http://heartbleed.com/>.
- [2] Y. Shin, A. Meneely, L. Williams, and J. A. Osborne, "Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities," *TSE'11*, vol. 37, no. 6, pp. 772–787.
- [3] A. Meneely and L. Williams, "Secure open source collaboration: an empirical study of linus' law," in *CCS'09*, pp. 453–462.
- [4] —, "Socio-technical developer networks: Should we trust our measurements?" in *ICSE'11*, pp. 281–290.
- [5] —, "Strengthening the empirical analysis of the relationship between linus' law and software security," in *ESEM'10*, p. 9.
- [6] A. Meneely, H. Srinivasan, A. Musa, A. R. Tejada, M. Mokary, and B. Spates, "When a patch goes bad: Exploring the properties of vulnerability-contributing commits," in *ESEM'13*, pp. 65–74.
- [7] T. Zimmermann, N. Nagappan, and L. Williams, "Searching for a needle in a haystack: Predicting security vulnerabilities for windows vista," in *ICST'10*, pp. 421–428.
- [8] A. Meneely, P. Rotella, and L. Williams, "Does adding manpower also affect quality?: an empirical, longitudinal analysis," in *FSE'11*, pp. 81–90.
- [9] A. Meneely, L. Williams, W. Snipes, and J. Osborne, "Predicting failures with developer networks and social network analysis," in *FSE'08*, pp. 13–23.
- [10] A. Sureka, A. Goyal, and A. Rastogi, "Using social network analysis for mining collaboration data in a defect tracking system for risk and vulnerability analysis," in *ISEC'11*, pp. 195–204.
- [11] T. Zimmermann and N. Nagappan, "Predicting defects using network analysis on dependency graphs," in *ICSE'08*, pp. 531–540.
- [12] S. Trabelsi, H. Plate, A. Abida, M. M. B. Aoun, A. Zouaoui, C. Mis-saoui, S. Gharbi, and A. Ayari, "Mining social networks for software vulnerabilities monitoring," in *NTMS'15*, pp. 1–7.
- [13] C. Bird, N. Nagappan, H. Gall, B. Murphy, and P. Devanbu, "Putting it all together: Using socio-technical networks to predict failures," in *ISSRE'09*, pp. 109–119.
- [14] P. Bhattacharya, M. Iliofotou, I. Neamtii, and M. Faloutsos, "Graph-based analysis and prediction for software evolution," in *ICSE'12*, pp. 419–429.
- [15] A. Younis, Y. K. Malaiya, and I. Ray, "Assessing vulnerability exploitability risk using software properties," *SQJ'16*, vol. 24, no. 1, pp. 159–202.
- [16] T. Wolf, A. Schroter, D. Damian, and T. Nguyen, "Predicting build failures using social network analysis on developer communication," in *ICSE'09*, pp. 1–11.
- [17] A. Kumar and A. Gupta, "Evolution of developer social network and its impact on bug fixing process," in *ISEC'13*, pp. 63–72.
- [18] X. Zheng, D. Zeng, H. Li, and F. Wang, "Analyzing open-source software systems as complex networks," *Physica A'08*, vol. 387, no. 24, pp. 6190–6200.
- [19] A. Meneely, A. C. R. Tejada, B. Spates, S. Trudeau, D. Neuberger, K. Whitlock, C. Ketant, and K. Davis, "An empirical investigation of socio-technical code review metrics and security vulnerabilities," in *SSE'14*, pp. 37–44.
- [20] A. Agrawal, A. Rahman, R. Krishna, A. Sobran, and T. Menzies, "We don't need another hero?: the impact of heroes on software development," in *ICSE-SEIP'18*, pp. 245–253.
- [21] S. Majumder, J. Chakraborty, A. Agrawal, and T. Menzies, "Why software projects need heroes (lessons learned from 1100+ projects)," *arXiv preprint arXiv:1904.09954*, 2019.
- [22] S. Koch and G. Schneider, "Effort, co-operation and co-ordination in an open source software project: Gnome," *Information Systems Journal'02*, vol. 12, no. 1, pp. 27–42.
- [23] A. Mockus, R. T. Fielding, and J. D. Herbsleb, "Two case studies of open source software development: Apache and mozilla," *TOSEM'02*, vol. 11, no. 3, pp. 309–346.
- [24] S. Krishnamurthy, "Cave or community?: An empirical examination of 100 mature open source projects."
- [25] G. Robles, J. M. Gonzalez-Barahona, and I. Herraiz, "Evolution of the core team of developers in libre software projects," in *MSR'09*, pp. 167–170.
- [26] C. Bird, P. C. Rigby, E. T. Barr, D. J. Hamilton, D. M. German, and P. Devanbu, "The promises and perils of mining git," in *MSR'09*, pp. 1–10.
- [27] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian, "The promises and perils of mining github," in *MSR'14*, pp. 92–101.
- [28] M. Joblin, W. Mauerer, S. Apel, J. Siegmund, and D. Riehle, "From developer networks to verified communities: a fine-grained approach," in *ICSE'15*, pp. 563–573.
- [29] A. Jermakovics, A. Sillitti, and G. Succi, "Mining and visualizing developer networks from version control systems," in *CHASE'11*, pp. 24–31.
- [30] W. Zhang, L. Nie, H. Jiang, Z. Chen, and J. Liu, "Developer social networks in software engineering: construction, analysis, and applications," *SCIS'14*, vol. 57, no. 12, pp. 1–23.
- [31] Y. Tymchuk, A. Mocchi, and M. Lanza, "Collaboration in open-source projects: Myth or reality?" in *MSR'14*, pp. 304–307.
- [32] M. Joblin, S. Apel, C. Hunsen, and W. Mauerer, "Classifying developers into core and peripheral: An empirical study on count and network metrics," in *ICSE'17*, pp. 164–174.
- [33] B. Çağlayan and A. B. Bener, "Effect of developer collaboration activity on software quality in two large scale projects," *JSS'16*, vol. 118, pp. 288–296.
- [34] J. Ren, H. Yin, Q. Hu, A. Fox, and W. Koszek, "Towards quantifying the development value of code contributions," in *FSE'18*, pp. 775–779.
- [35] F. Palomba, D. A. Tamburri, A. Serebrenik, A. Zaidman, F. A. Fontana, and R. Oliveto, "How do community smells influence code smells?" in *ICSE-Companion'18*, pp. 240–241.
- [36] A. Jermakovics, A. Sillitti, and G. Succi, "Exploring collaboration networks in open-source projects," in *IFIP-ICOS'13*, pp. 97–108.
- [37] M. Joblin, S. Apel, and W. Mauerer, "Evolutionary trends of developer coordination: A network approach," *EMSE'17*, vol. 22, no. 4, pp. 2050–2094.
- [38] M. Pinzger, N. Nagappan, and B. Murphy, "Can developer-module networks predict failures?" in *FSE'08*, pp. 2–12.
- [39] F. Thung, T. F. Bissyande, D. Lo, and L. Jiang, "Network structure of social coding in github," in *CSMR'13*, pp. 323–326.
- [40] Q. Hong, S. Kim, S. Cheung, and C. Bird, "Understanding a developer social network and its evolution," in *ICSM'11*, pp. 323–332.
- [41] C. Bird, D. Pattison, R. D'Souza, V. Filkov, and P. Devanbu, "Latent social structure in open source projects," in *FSE'08*, pp. 24–35.
- [42] D. Surian, D. Lo, and E.-P. Lim, "Mining collaboration patterns from a large developer network," in *WCRE'10*, pp. 269–273.
- [43] S. Wang, W. Zhang, Y. Yang, and Q. Wang, "Devnet: exploring developer collaboration in heterogeneous networks of bug repositories," in *ESEM'13*, pp. 193–202.
- [44] T. Wolf, A. Schröter, D. Damian, L. D. Panjer, and T. H. Nguyen, "Mining task-based social networks to explore collaboration in software teams," *IEEE Software'09*, vol. 26, no. 1, pp. 58–66.
- [45] W. Zhang, S. Wang, Y. Yang, and Q. Wang, "Heterogeneous network analysis of developer contribution in bug repositories," in *ICSC'13*, pp. 98–105.
- [46] D. Surian, N. Liu, D. Lo, H. Tong, E.-P. Lim, and C. Faloutsos, "Recommending people in developers' collaboration network," in *WCRE'11*, pp. 379–388.
- [47] D. W. McDonald, "Recommending collaboration with social networks: a comparative evaluation," in *CHI'03*, pp. 593–600.
- [48] G. Jeong, S. Kim, and T. Zimmermann, "Improving bug triage with bug tossing graphs," in *FSE'09*, pp. 111–120.
- [49] M. S. Zanetti, I. Scholtes, C. J. Tessone, and F. Schweitzer, "Categorizing bugs with social networks: a case study on four open source software communities," in *ICSE'13*, pp. 1032–1041.
- [50] C. Bird, A. Gourley, P. Devanbu, M. Gertz, and A. Swaminathan, "Mining email social networks," in *MSR'06*, pp. 137–143.
- [51] M. S. Zanetti, I. Scholtes, C. J. Tessone, and F. Schweitzer, "The rise and fall of a central contributor: dynamics of social organization and performance in the gentoo community," in *CHASE'13*, pp. 49–56.

- [52] H. Jiang, J. Zhang, H. Ma, N. Nazar, and Z. Ren, "Mining authorship characteristics in bug repositories," *SCIS'17*, vol. 60, no. 1, p. 012107.
- [53] M. Zhou and A. Mockus, "Who will stay in the floss community? modeling participants initial behavior," *TSE'15*, vol. 41, no. 1, pp. 82–99.
- [54] —, "What make long term contributors: Willingness and opportunity in oss community," in *ICSE'12*, pp. 518–528.
- [55] M. Gharehyazie, D. Posnett, B. Vasilescu, and V. Filkov, "Developer initiation and social interactions in oss: A case study of the apache software foundation," *EMSE'15*, vol. 20, no. 5, pp. 1318–1353.
- [56] L. López-Fernández, G. Robles, J. M. Gonzalez-Barahona, and I. Her-raiz, "Applying social network analysis techniques to community-driven libre software projects," *IJITWE'06*, vol. 1, no. 3, pp. 27–48.
- [57] L. Lopez-Fernandez, G. Robles, J. M. Gonzalez-Barahona *et al.*, "Ap-plying social network analysis to the information in cvs repositories," in *MSR'04*, p. 101105.
- [58] S. L. Toral, M. d. R. Martínez-Torres, and F. Barrero, "Analysis of virtual communities supporting oss projects using social network analysis," *IST'10*, vol. 52, no. 3, pp. 296–303.
- [59] G. Canfora, L. Cerulo, M. Cimitile, and M. Di Penta, "Social in-teractions around cross-system bug fixings: the case of freebsd and openbsd," in *MSR'11*, pp. 143–152.
- [60] C. Bird, N. Nagappan, B. Murphy, H. Gall, and P. Devanbu, "Don't touch my code!: examining the effects of ownership on software quality," in *FSE'11*, pp. 4–14.
- [61] J. Tsay, L. Dabbish, and J. Herbsleb, "Influence of social and technical factors for evaluating contribution in github," in *ICSE'14*, pp. 356–366.
- [62] F. Rahman and P. Devanbu, "Ownership, experience and defects: a fine-grained study of authorship," in *ICSE'11*, pp. 491–500.
- [63] B. Zhou, I. Neamtiu, and R. Gupta, "A cross-platform analysis of bugs and bug-fixing in open source projects: Desktop vs. android vs. ios," in *EASE'15*, p. 7.
- [64] D. Izquierdo-Cortazar, A. Capiluppi, and J. M. Gonzalez-Barahona, "Are developers fixing their own bugs?: Tracing bug-fixing and bug-seeding committers," *IJOSSP'11*, vol. 3, no. 2, pp. 23–42.
- [65] D. M. German, "The gnome project: a case study of open source, global software development," *Software Process: Improvement and Practice*, vol. 8, no. 4, pp. 201–215, 2003.
- [66] J. Xuan, H. Jiang, Z. Ren, and W. Zou, "Developer prioritization in bug repositories," in *ICSE'12*, pp. 25–35.
- [67] T. T. Dinh-Trong and J. M. Bieman, "The freebsd project: A replication case study of open source development," *TSE'05*, vol. 31, no. 6, pp. 481–494.
- [68] C. Izurieta and J. Bieman, "The evolution of freebsd and linux," in *ISESE'06*, pp. 204–211.
- [69] Y. Tian, J. Lawall, and D. Lo, "Identifying linux bug fixing patches," in *ICSE'12*, pp. 386–396.
- [70] U. N. I. of Standards and Technology, "National vulnerability database," <https://nvd.nist.gov/home.cfm>.
- [71] M. Corporation, "Common vulnerabilities and exposures," <https://cve.mitre.org/>.
- [72] S. Kim, T. Zimmermann, K. Pan, E. James Jr *et al.*, "Automatic identification of bug-introducing changes," in *ASE'06*, pp. 81–90.
- [73] D. Wijayasekara, M. Manic, J. L. Wright, and M. McQueen, "Mining bug databases for unidentified software vulnerabilities," in *ICHSI'12*, pp. 89–96.
- [74] S. E. Ponta, H. Plate, A. Sabetta, M. Bezzi, and C. Dangremont, "A manually-curated dataset of fixes to vulnerabilities of open-source software," in *MSR'19*.
- [75] Y. Zhou and A. Sharma, "Automated identification of security issues from commit messages and bug reports," in *FSE'17*, pp. 914–919.
- [76] "Sourceclear," <https://www.sourceclear.com/>.
- [77] J. Wang, M. Li, S. Wang, T. Menzies, and Q. Wang, "Images dont lie: Duplicate crowdtesting reports detection with screenshot information," *IST'19*.
- [78] J. Wang, Q. Cui, Q. Wang, and S. Wang, "Towards effectively test report classification to assist crowdsourced testing," in *ESEM'16*, p. 6.
- [79] P. Runeson, M. Alexandersson, and O. Nyholm, "Detection of duplicate defect reports using natural language processing," in *ICSE'07*, pp. 499–510.
- [80] H. Rocha, M. T. Valente, H. Marques-Neto, and G. C. Murphy, "An empirical study on recommendations of similar bugs," in *SANER'16*, vol. 1, pp. 46–56.
- [81] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal, *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [82] F. Li and V. Paxson, "A large-scale empirical study of security patches," in *CCS'17*, pp. 2201–2215.
- [83] H. Perl, S. Dechand, M. Smith, D. Arp, F. Yamaguchi, K. Rieck, S. Fahl, and Y. Acar, "Vccfinder: Finding potential vulnerabilities in open-source projects to assist code audits," in *CCS'15*, pp. 426–437.
- [84] J. Śliwerski, T. Zimmermann, and A. Zeller, "When do changes induce fixes?" in *MSR'05*, vol. 30, no. 4, pp. 1–5.
- [85] D. A. da Costa, S. McIntosh, W. Shang, U. Kulesza, R. Coelho, and A. E. Hassan, "A framework for evaluating the results of the szz approach for identifying bug-introducing changes," *TSE'17*, vol. 43, no. 7, pp. 641–657.
- [86] T. Jiang, L. Tan, and S. Kim, "Personalized defect prediction," in *ASE'13*, pp. 279–289.
- [87] Z. Gu, E. T. Barr, D. J. Hamilton, and Z. Su, "Has the bug really been fixed?" in *ICSE'10*, vol. 1, pp. 55–64.
- [88] G. Robles and J. M. Gonzalez-Barahona, "Developer identification methods for integrated data from various sources," in *MSR'05*, pp. 1–5.
- [89] K. Ehrlich and M. Cataldo, "All-for-one and one-for-all?: a multi-level analysis of communication patterns and individual performance in geographically distributed software development," in *CSCW'12*, pp. 945–954.
- [90] E. Giger, M. Pinzger, and H. C. Gall, "Can we predict types of code changes? an empirical analysis," in *MSR'12*, pp. 217–226.
- [91] A. D. Well and J. L. Myers, *Research design & statistical analysis*. Psychology Press, 2003.
- [92] S. Kim, H. Zhang, R. Wu, and L. Gong, "Dealing with noise in defect prediction," in *ICSE'11*, pp. 481–490.
- [93] S. Astromskis, G. Bavota, A. Janes, B. Russo, and M. Di Penta, "Patterns of developers behaviour: A 1000-hour industrial study," *JSS'07*, vol. 132, pp. 85–97.
- [94] W. Bu, M. Xue, L. Xu, Y. Zhou, Z. Tang, and T. Xie, "When program analysis meets mobile security: an industrial study of misusing android internet sockets," in *FSE'17*, pp. 842–847.
- [95] N. Munaiah, "Assisted discovery of software vulnerabilities," in *ICSE'18*, pp. 464–467.
- [96] F. Camilo, A. Meneely, and M. Nagappan, "Do bugs foreshadow vulnerabilities?: a study of the chromium project," in *MSR'15*, pp. 269–279.
- [97] A. Decan, T. Mens, and E. Constantinou, "On the impact of security vulnerabilities in the npm package dependency network," in *MSR'18*, pp. 181–191.
- [98] S. Frei, M. May, U. Fiedler, and B. Plattner, "Large-scale vulnerability analysis," in *SIGCOMM'06*, pp. 131–138.
- [99] M. Shahzad, M. Z. Shafiq, and A. X. Liu, "A large scale exploratory analysis of software vulnerability life cycles," in *ICSE'12*, pp. 771–781.
- [100] H. Zhong and Z. Su, "An empirical study on real bug fixes," in *ICSE'15*, pp. 913–923.
- [101] D. Mu, A. Cuevas, L. Yang, H. Hu, X. Xing, B. Mao, and G. Wang, "Understanding the reproducibility of crowd-reported security vulner-abilities," in *USENIX Security'18*, pp. 919–936.
- [102] A. Ozment and S. E. Schechter, "Milk or wine: does software security improve with age?" in *USENIX Security Symposium'06*.
- [103] Z. Xu, B. Chen, M. Chandramohan, Y. Liu, and F. Song, "Spain: security patch analysis for binaries towards understanding the pain and pills," in *ICSE'17*, pp. 462–472.
- [104] J. Walden, J. Stuckman, and R. Scandariato, "Predicting vulnerable components: Software metrics vs text mining," in *ISSRE'14*, pp. 23–33.
- [105] N. Medeiros, N. Ivaki, P. Costa, and M. Vieira, "Software metrics as indicators of security vulnerabilities," in *ISSRE'17*, pp. 216–227.
- [106] Y. Yang, J. Ai, X. Li, and W. E. Wong, "Mhcp model for quality evaluation for software structure based on software complex network," in *ISSRE'16*, pp. 298–308.
- [107] J. Park, M. Kim, B. Ray, and D.-H. Bae, "An empirical study of supplementary bug fixes," in *MSR'12*, pp. 40–49.