

# Mining the Temporal Evolution of the Android Bug Reporting Community via Sliding Windows

Feng Jiang      Jiemin Wang      Abram Hindle  
Mario A. Nascimento

September 2013

Department of Computing Science  
University of Alberta  
Edmonton, Alberta, Canada

© University of Alberta



## **Abstract**

The open source development community consists of both paid and volunteer developers as well as new and experienced users. Previous work has applied social network analysis (SNA) to open source communities and has demonstrated value in expertise discovery and triaging. One problem with applying SNA directly to the data of the entire project lifetime is that the impact of local activities will be drowned out. In this paper we provide a method for aggregating, analyzing, and visualizing local (small time periods) interactions of bug reporting participants by using the SNA to measure the betweenness centrality of these participants. In particular we mined the Android bug repository by producing social networks from overlapping 30-day windows of bug reports, each sliding over by day. In this paper we define three patterns of participant behaviour based on their local centrality. We propose a method of analyzing the centrality of bug report participants both locally and globally, then we conduct a thorough case study of the bug reporters' activity within the Android bug repository. Furthermore, we validate the conclusions of our method by mining the Android version control system and inspecting the Android release history. We found that windowed SNA analysis elicited local behaviour that were invisible during global analysis.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Background</b>	<b>3</b>
2.1	Betweenness Centrality . . . . .	3
2.2	Overlapping Time Windowing . . . . .	4
2.3	Clustering . . . . .	5
<b>3</b>	<b>Methodology</b>	<b>6</b>
3.1	Data . . . . .	6
3.2	Windowing Bug Reports and Extracting Social Networks Methodology . . . . .	6
3.3	Validation using the Android Release History and the Git . . . . .	8
<b>4</b>	<b>Results and Analysis</b>	<b>8</b>
4.1	Global Analysis . . . . .	8
4.2	Local Analysis . . . . .	11
<b>5</b>	<b>Validation</b>	<b>12</b>
5.1	Activity pattern validation . . . . .	12
5.1.1	Participants that appeared only once tend to be pure users . . . . .	12
5.1.2	Participants showed up periodically should be a combination of users and developers . . . . .	12
5.1.3	Participants who were continuously central for a long time pe- riod could have multiple areas of expertise . . . . .	13
5.2	Cluster validation . . . . .	15
5.2.1	Participants clustered together share similar areas of expertise and tasks . . . . .	16
5.2.2	Clusters' working areas of expertise are in accordance with the release contents along the time line . . . . .	16
<b>6</b>	<b>Limitations</b>	<b>17</b>
<b>7</b>	<b>Conclusion and Future Work</b>	<b>17</b>



# 1 Introduction

Global analysis provides us with easy to interpret data that gives us an overview of the entire system. It simplifies complicated dimensions like time and provides us with an easy way to explain results. Unfortunately, for tools like *Social Network Analysis* (SNA), a global analysis can miss a lot of important interactions, especially between stakeholders, thus we propose a method of using SNA to study bug repositories and tease out local collaborations.

SNA is a powerful tool that helps practitioners and researchers study the complicated interactions of participants within communities; SNA is well accepted in the area of software maintenance and mining software repositories communities [1, 2, 3]. The bug repository records interactions among software developers and users in a software project's community. With SNA, we are able to study the structure of the interactions by analysing the graph constructed through the interaction of bug reporters in the bug repository. The results can be used in expertise elicitation and triaging in order to suggest which participants have expertise relevant to an issue [3]. Usually SNA is run globally across all day, over a single period, or over an entire project lifetime. In this paper we argue that using SNA in a more local manner provides valuable insights into interactions between stakeholders during the development and maintenance of a software system.

Open-source communities are amenable to social network analysis as they are open to user interaction and participation. At the same time there is a lack of imposed organizational structures found within corporate organizations [4]. Because open source projects often lack strict centralized control and requirements [5], developers often choose their tasks instead of being assigned one [6]. This fact suggests that local structure of interactions among users and developers who express an interest in one part of the project tend to self organize and produce interesting collaboration structures (networks).

Bug repositories are also amenable to social network analysis as bug repositories host and record discussions regarding issues or bugs relevant to the development and the use of a software development project [6, 7]. Bug repositories are also heavily used by open-source projects. Collaboration among developers has been studied in various aspects about how the communication introduces or avoids bugs, and further influences the software quality, [8], [9], [10], [11]. Besides the collaboration among developers, collaboration between users and developers is evident in bug reports since the discussions and communications are recorded as reported bugs, and posted comments on bug reports. One point here is that, both users and developers are often periodic, and their activities or collaborations can be local and thus missed out in global analysis.

In the case of the Android bug repository, provided by the 2012 MSR Mining Challenge [12], a reporter would report a bug, which might attract comments from bug commenters; the commenters discuss the reasons and possible resolution of the bug. The bug reporting community members are usually comprised of both bug reporters and bug commenters who are either Android developers or Android users. From the perspective of the bug repository, unlike the version control system, there is actually no obvious boundary between a user and a developer. We refer to these different participants as *bug participants*.

In order to apply SNA to the bug repository, we first create the graph based on the interactions. We pose that each node of the graph represents one bug participant and each edge represents the connection between two participants who have communicated on the same bug. We will introduce the network graphs in detail in Section 3.

We use betweenness centrality to quantify the importance of a participant in the community [13] (betweenness centrality will be better explained in Section 3). The betweenness centrality could reveal two aspects of a participant in a community network: 1) the quantity of bug reports (which attract at least one comment) or comments they have made and 2) the importance of the content of their reports or comments. When participants have high betweenness, they might have: 1) reported quantities of bugs with at least one comment on them, 2) made lots of comments, 3) reported a very critical bug which attracts comments, 4) or made a very interesting comment which attracts comments from other participants.

However, the previous work [2, 6] applied SNA on the entire lifetime of a project, such that only a single community network was constructed. Some of collaborations might not be evident if one were to analyze a large single network. That is because certain structures will not be observable on the global scale. In order to peer into these local self organized structures using social network analysis, we felt it is better to choose a *windowed* approach, [3, 14]. Windowing allows us to look at network during a slice of time and then relate our measures (betweenness centrality per author) to the next window and beyond. This *sliding window* view of centrality allows us to see those developers and users who are constantly at the forefront of discussion or those who ebb and flow between issues and tasks. Moreover, by sliding windows, each pair of adjacent windows would have an overlap, which results in smoother trends, and more importantly, helps to maintain context. Other benefits provided by time windowed analysis is that it gives a more accurate and nuanced view of the data as locally central participants then will not be “drowned out”.

In summary, we use SNA to study the activities of bug participants based on the Android bug reports and comments repository. We apply the sliding window method to observe smooth change trends in the collaboration graph across time. With these mining results, we seek to analyze bug participants’ interactions, activity trends and patterns. We then demonstrate our analysis results via answering the following research questions about local and global behaviours.

Global research questions:

**RQ1.** How does the number of active bug participants change over time? Why?

**RQ2.** How does the betweenness centrality of a participant change over time? What are the reasons when they have a certain activity pattern?

Local research questions:

**RQ3.** Are there special time ranges during which participants are more/less active or central than normal? Why?

**RQ4.** What are the possible scenarios for a very sharp change of the participants’ centrality? Why?

We also validate if this windowed methodology actually highlights relevant behaviour by inspecting the Android release history<sup>1</sup> and the Android version control

---

<sup>1</sup>Android release history: <http://developer.android.com/sdk/index.html>

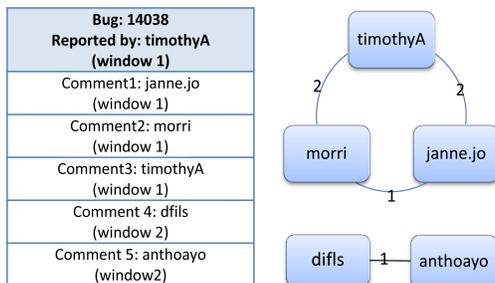


Figure 1: An example: bug 14038 is reported by timothyA, and there are five comments on this bug. When time window applied, comments are plotted into two windows, and the bug report of this example forms two networks with the weight noted on their edges

system. The validation would be discussed in Section 5.

The rest of our paper is organized as follows. Section 2 introduces basic concepts and techniques we used in this study. The specific steps and the methodology will be discussed in Section 3. Section 4 describes the details of our mining results. The analysis of the results and its corresponding validation is provided in Section 5. Section 6 presents the limitations of our mining process and Section 7 summarizes the paper and discusses the future work.

## 2 Background

### 2.1 Betweenness Centrality

The betweenness centrality of a vertex is the number of geodesic paths in a graph that includes this vertex; the geodesic path is defined as the shortest path which has the minimum weight between two nodes. Defined by Freeman [15], the betweenness can be represented as:

$$\sum_{i=1}^{j-1} \sum_{j=1}^n \frac{g_{ij}(k)}{g_{ij}}, i \neq j \neq k \quad (1)$$

where  $k$  is a vertex of the graph,  $n$  is the total number of vertices,  $i$  and  $j$  are vertices other than  $k$ ,  $g_{ij}$  is the number of geodesic paths between vertex  $i$  and  $j$ , and  $g_{ij}(k)$  is the number of geodesic paths that include  $k$ .

It is used as a measurement of a person's importance in a network. A person would be regarded as central if he is on the geodesic path between two other persons. As proposed by Freeman [15], if a person is located on the geodesic path between two other persons, he becomes one of the key persons who connects the others. That is, the more a person connects to the other people in a network, the more important or central he is [16].

In our work, we normalize the betweenness centrality values to eliminate the effect of different sizes of the networks. The betweenness is normalized as:

$$\text{Normalized } B = \frac{B}{\frac{(n-1)(n-2)}{2}} \quad (2)$$

where  $B$  represents the original betweenness value and  $n$  is the number of nodes in the graph being calculated.

Compared with simply counting the total number of comments or total bug reports of a participant, betweenness acts better to reflect the interactions among people. For example, when a person reports lots of bugs but none of them attract any comment, it is very likely that his bug reports are not interesting or important. In this case, if we merely counted the number of their reports or comments, we would possibly increase their importance in the network artificially. Therefore, we choose to use betweenness centrality to eliminate this unfair counting [13].

## 2.2 Overlapping Time Windowing

When SNA is applied in other papers [2, 3], it is typically applied to the entire history or one period of the partial history and all the bug reports within that period. Windowed analysis instead repeats social network analysis across 100s of windows (in our case, as many windows as we have days). These windows overlap and often the analysis of one window results in the same analysis as the previous window due to the overlap. We slid our windows by 1 day and for two adjacent windows  $A$  and  $B$ ,  $B$  starts on the second day of  $A$ , and they would have an overlap of 29 days, that is each window does some redundant analysis but produces smoother transitions in analysis between windows. Thus 1 comment in a bug report will have an effect on the graphs of 30 windows. This is similar to Hindle et al.’s [14] analysis of topics using windows but they did not use an overlap. We could thus see the changes in the trend of a participant’s activity.

Moreover, time windowed analysis could give a more accurate and nuanced view of the data [3, 14], as locally central participants would not be “drowned out”. For instance, if a bug participant participates in many bug reports and bug comments during one month, he would be one of the most central participants with a high betweenness within this window. However, if he appeared only for that month, globally, he would have low betweenness and would not show up as central, even though during a shorter period he played a vital role. As we can see in Figure 2, the left column graph shows the betweenness values of participants over the entire time period; local details are missed and we get nothing about the trend, compared to the right part of the results from overlapping windowing. For example, cluster 8 on Figure 2 is bright and important at the start of our analysis but does not appear in the global graph on the left. Also, if there is a very sharp drop of values of a certain participant, the overlapping windows would give a more nuanced view of the change and what was happening.

Another point is that, comments on the same bug might not be globally temporally relevant [17, 18] thus a global time analysis would not make much sense in this case. This could happen if new changes induce new bugs or modify the behaviour of a reported bug.

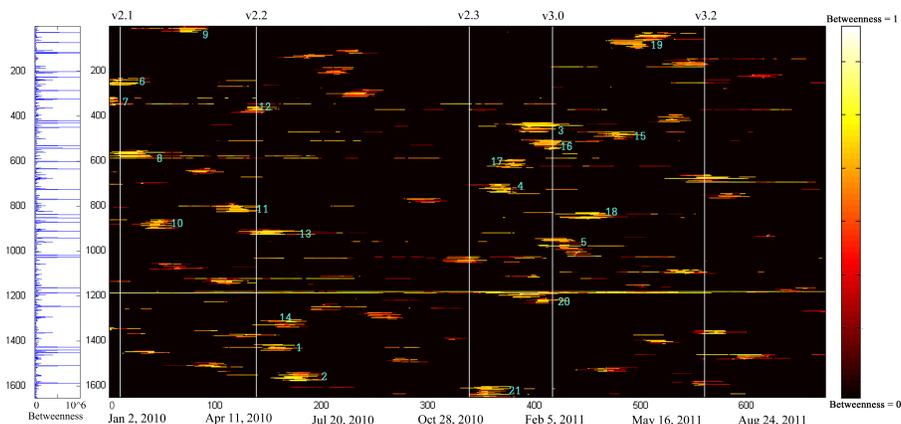


Figure 2: Betweenness centrality along time line: the x-axis represents the number of time windows and the starting dates are denoted every 100 windows. The y-axis represents the number of bug participants who have ever been central in the bug community with betweenness centrality valued greater than 0 for some time period. The color represents the value of betweenness centrality, with darker colors corresponding to lower betweenness and lighter colors for higher betweenness. We used K-means clustering with cosine distance where  $K = 100$

### 2.3 Clustering

Figure 3(1) orders participants by their betweenness values. We can indeed find that there are participants of low overall betweenness but being very active (show up bright) at some time points, and this supports the necessity of windowing, as stated in Section 2.2. However, we need more information about participants' working patterns and get an idea about their being interacting groups.

In order to perceive clusters, that are local groups of interactions, we clustered the bug participants using K-means by their betweenness centrality distribution along the time line. K-means is one of the most popular clustering methods which aims to partition  $n$  data items into  $k$  clusters that each data item belongs to the cluster with the nearest mean [19]. We choose to use K-means with cosine distance. The cosine distance between two vectors is defined as,

$$\text{cosine\_dist}(A, B) = 1 - \frac{A \cdot B}{\|A\| \|B\|} \quad (3)$$

where  $A$  and  $B$  are two vectors,  $\cdot$  represents the inner dot operation and  $\|\cdot\|$  indicates the module of the vector. With clustering, authors with similar temporal centrality would be grouped together so that bug participants with similar activity patterns are also grouped together.

In this paper, we choose the K-means with cosine distance because it gives a better

visual clustering result, as compared in the plots of Figure 3. In this case, cosine distance calculates the similarity between each pair of participants in terms of temporal centrality whereas Euclidean distance focuses on the magnitude of data, the size and frequency of centrality.

Moreover, we used  $k = 100$  for the K-means, to cluster authors. There is a trade off between the size of clusters and the variance within each cluster. As we can see from Figure 3(3),  $k = 10$  also gives good visualization result, but considering the number of more than 1600 participants, we should get a larger  $k$  to keep the diversity of working groups in similar. We set  $k = 100$  in this case, since we get the aesthetically best visualization (our subjective opinion based on visible clusters) of all the data; we had tried other values of  $k$  such as 5, 10, 25, 50, 75, 200, 300.

### 3 Methodology

Our methodology consists of six steps that deal with raw data, construct graphs and apply social network analysis with sliding windows. We conduct a thorough case study of the Android bug repository with the proposed method and validate the conclusions from the results by mining the Android version control system and inspecting the release history.

#### 3.1 Data

With the provided Android bug repository 2012 and the Android version control system from the MSR Challenge [12], we converted and stored the XML format data into a database for efficient analysis using Microsoft SQL Server Business Intelligence. Our analysis focused on the bug records of the previous two years from January 1st 2010 to December 4th 2011 since during these two years, there are more records in the repository as we counted that participants are more active; also, the activities are representative, both the Android platforms and their developer groups are larger and more diverse during the latest two years and it was also more relevant to modern Android handsets. The data we used of these two years covers 14,432 out of 20,169 total bug records and 46,806 out of 67,730 total bug comments from the whole dataset. Related to these bug and comment records, there are 30,969 people who have either reported a bug or made comments on a bug.

The bug and comment records are grouped into 30-day windows sliding by 1 day. We extracted 673 windows in total from the bug reports during year 2010 and 2011.

#### 3.2 Windowing Bug Reports and Extracting Social Networks Methodology

We windowed the data and constructed networks that indicated the relations among the participants within each specific time window. For each window, we calculated the betweenness centrality of each participant and we plotted the centrality values per participant in a visualization. The steps of our methodology are explained as following:

**Step 1: Pruning the data.** We pruned the records of the reporters and commenters into a pure name format, which are originally recorded in semi-anonymous email formats in the XML repository dump. For example, given the original email address which is represented by “mathias....@gmail.com”, we truncate the string starting from “...” and keep the front part “mathias” at the beginning as the name of the reporter or commenter. This strategy could lead to name aliasing problem, especially for common names or email addresses starting at just a simple letter like “e....@gmail.com”. Although algorithms have been provided to reduce the extent of the problem, [20], [2], [21], it is difficult or even impossible to eliminate the influence from this data quality issue. When applied to other repositories that do not anonymize this would be less of a problem. Hence, we focus on participants whose names are less common and less ambiguous in our study.

**Step 2: Windowing the records.** We windowed the data into periods of 30 days with a 29-day overlap. 30 days was chosen as a window size because it is smaller than the periods between a major and minor release, it is similar to a month of work, but long enough to contain the resolution of multiple bugs. We have compared sliding by 1 day with our previous result of sliding by 7 days, 1 day sliding produces gradual and smoother transitions of centrality.

**Step 3: Establishing the network.** We made a tool to perform the SNA with sliding windows. The tool is implemented in Java and built on top of the JUNG Graph Framework, that converted bug reports and bug comment records within a window to a social network graph.

The nodes of these networks represent participants who have either reported some bugs or made comments on bugs. The edges represent connections between two nodes. All the edges are weighted. For a bug within a selected time window, whenever a person makes a comment on this bug, the edge between the bug commenter and the bug reporter would get weight plus one, as well as the edges linking to the participants who previously made comments on this bug. Bug reports or comments in different windows would have separate network graphs depending on the activity of their reporters or commenters. An example in Figure 1 indicates how the weighted network graph is built.

**Step 4: Calculating the centrality.** We calculated the betweenness centrality using JUNG, and normalized the centrality with the number of node pairs, as in Equation (2). We then get a list of all the bug participants and their betweenness centrality values for the total 673 overlapping windows.

**Step 5: Removing irrelevant participants.** We removed the participants with betweenness centrality value 0, who might have either reported a bug/bugs with no comments, or made the only comment on a bug so that no other participants are related. Afterwards, we get 1654 participants with betweenness centrality value larger than 0, out of the 30969 in total.

**Step 6: Generating the analysis graph.** The activity of each bug participant is represented by a 673 dimensional vector representing their betweenness per window. Each element of the vector indicates the betweenness centrality value extracted from the graph, which is generated from the window for that specific time period (in our case, the specific time period is 30 days starting from the date of the window start point). Then we clustered all the vectors by using K-means ( $k = 100$ ) with cosine distance

to 100 clusters. Finally, we plotted the results, as shown in Figure 2 to visualize the clustered data so that we could easily analyze our results.

### 3.3 Validation using the Android Release History and the Git

In addition to the methodology of mining the Android bug repository, we made use of the git version control repository and inspected the release history highlights to validate the purpose behind the clusters and patterns we observed. We looked into the participants who contributed to the git repository in order to find their areas of expertise and validate our analysis conclusion about how the community participants act in accordance with the project development.

The types of files modified and the corresponding projects are highly correlated with the specialization of those who commit changes. For instance, if a developer always submits kernel related code files, he is more likely to be specialized in kernel techniques. Types of files include document files, test files, source files, etc; dictionary paths of files usually indicate what projects the files belong to. We manually identified the participants' areas of expertise by observing the *project* and the *target* for all of their commits (such as source code or documentation). To give a specific example, if there were commits from a developer, Mr. Guilfoyle, on the target file `media/java/android/media/Ringtone.java` under the project `platform_frameworks_base`; then, we would suggest that Mr. Guilfoyle likely has some specialized knowledge about the platform's ringtone. Thus this is how we derive participant expertise [3].

Also, we could further relate their expertise to their centrality patterns. The Android release history could, on the other hand, help to relate the release highlights to participants central behaviour during that release. Further validation is discussed in Section 5.

## 4 Results and Analysis

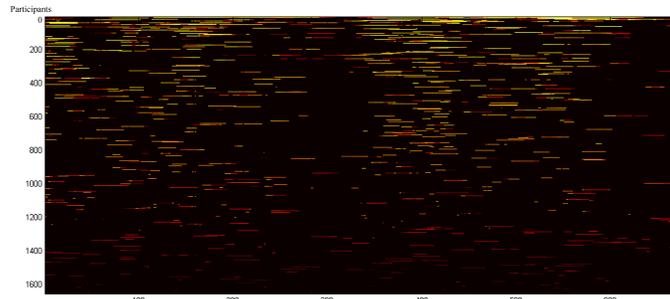
We study the results shown in Figure 2. Each horizontal line represents the 673 betweenness centrality values for the selected bug participant during year 2010 and 2011. In total, we have 1654 bug participants. By studying these results, we answered the following questions:

### 4.1 Global Analysis

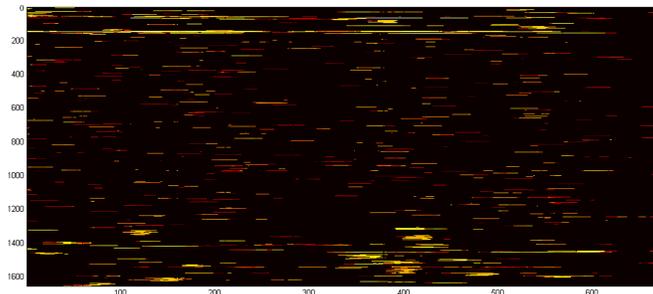
#### **RQ1. How does the number of active bug participants change over time? Why?**

To give an overview, we compared the interaction of bug participants between January, 2010 and December, 2011, and found that the interaction among participants in the Android bug community in 2011 was similar to the interaction of participants in 2010 but more frequent, as we can see in Figure 2. One "gap" occurs around window 300, which we will explain in the next Local Analysis subsection.

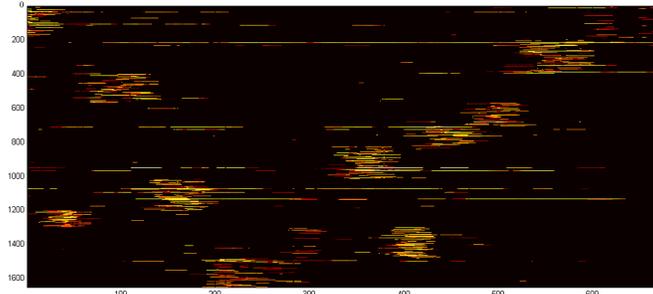
Correspondingly in Figure 4, that we counted the number of participants with betweenness centrality value larger than 0 within each window, the number of active



(1) Participants ordered by betweenness value, numbered from 1 to 1654 in decreasing order.



(2) Participants clustered into 100 clusters using Euclidean distance.



(3) Participants clustered into 10 clusters using cosine distance.

Figure 3: Betweenness centrality along time line. Participants on the y-axes are ordered differently by betweenness values or various clusterings.

participants during 2011 is slightly larger than that of 2010. Figure 5 shows the sum of betweenness values along the two years' time line, we can see that the trend is very similar to that of the number of active participants in Figure 4. This also suggests that the betweenness centrality reflects the interaction among participants.

Moreover, a possible reason for the changes of the number of active bug participants and the betweenness centrality values is that around major or minor releases of SDKs, API fixes or improvements, participants seem to become more active in bug reporting, discussing and fixing activities. Also, during these time periods, bugs are more likely to be discovered and reported. Perhaps the pressure of the release is causing developers



Figure 4: Number of active participants across time.

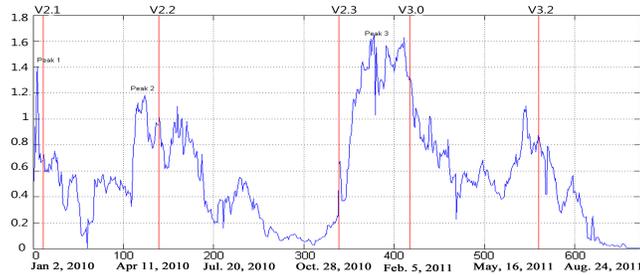


Figure 5: Sum of betweenness centrality of participants across time.

to address outstanding bugs more than usual. After a release, users also take part in the activity of discovering the bugs and problems so that in this case both users and developers would like to discuss the bugs.

**RQ2. How does the betweenness centrality of a participant change over time? What are the reasons when they have a certain activity pattern?**

Observing the continuity of betweenness centrality in Figure 2, some participants have kept active during the entire two years, and correspondingly they have a very continuous and bright line. For participants of this type, there are a few possible explanations. First, our conjecture is that these participants are professional developers who belong to the core development team so that what they reported are more important issues which attract more participants to discuss and fix them. Their identities of being professional developers will be discussed in Section 5.1.

Second, some of these participants are of high community status or expertise, and they might supervise and guide the development of the project. For example, when we validated, we did find one developer, romanguy, who has experiences on almost every component relevant to platforms so that he can be considered to be an expert. Developers related to these continuous lines are listed in Table 1, and we will further discuss and validate on them in Section 5.1.

However, in most cases, participants' betweenness values are highly variant, as observed in Figure 2. To investigate the variation in betweenness values over time, we decided to count the number of times that a user experienced a range of consecutive

windows in which the user had non-zero betweenness.

Participants with a count of distinct ranges greater than 1 would be *phasers* who periodically participate within the Android bug community. Here phasers are those who phase into centrality and later out of it. These randomly phasing participants (phasers) are very likely to acquire less expertise or have lower community status in their community, than those with continuous high centrality. Phasers might be interested in limited topics and only central and active during the appearance of bugs relevant to those topics. Participants who only had 1 distinct range of betweenness are considered to be participants who only appeared once, and are probably users. We validate the roles these participants play in Section 5.1.

To summarize, among the 1654 participants with betweenness values larger than 0, we analyzed their centrality patterns and divide them into three categories: 1) participants appeared only once with a betweenness greater than 0 (71 out of 1654 participants), 2) participants recurred periodically (1575 participants) and 3) participants who are central along the entire project history (8 participants).

## 4.2 Local Analysis

### **RQ3. Are there special time ranges during which participants are more/less active or central than normal? Why?**

By inspecting the Android release history highlights, we found that the v2.1 SDK was released on 12 January 2010, which corresponds to the first peak value in Figure 5. Android v2.2 SDK was released on 20 May 2010 and this corresponds to peak 2. From Dec. 2010 to the beginning of Mar. 2011, several minor updates were released and on 22 Feb. 2011, one major update v3.0 SDK was released. These releases explain the summit, i.e., peak 3, in Figure 5. This is correlated with more participation at the same time.

In addition, during the first obvious “gap”, which covers the time from October 2010 to the end of 2010 (around window 300), the social network during this time period is almost inactive and even “quiet”. There were fewer releases during the “gap”.

The other low value showing up in the end of Figure 5 results from the fact that there are no bug reports recorded (right tail censoring) in the given dataset. This piece of data is still meaningful because it contains comments belonging to bug reports several weeks or months before. The betweenness value is thus simply calculated by the comments here.

### **RQ4. What are the possible scenarios for a very sharp change of the participants' centrality? Why?**

Considering individual participants, almost all of them has experienced centrality oscillations. In addition, some participants tend to become active and core members during the same time period and then they fade away together.

We suspect that the phasers tend to be interested in one or several categories of problems so that they appear only along with the occurrence of these issues. They take part in activities related to the bugs or technical issues and become inactive after the problems are solved. Or in the case when they are working on a project, they would become inactive when the projects are finished. As showed in Figure 2, the participants' tend to get clustered together around important releases, which supports

that the phasers are working along with projects or related issues. Meanwhile, by observing the clustered participants of their activity patterns in Figure 2, we suspect that the phasers that show up densely together could be interested in similar categories of topics. This assumption is validated in Section 5.2.

## 5 Validation

We made use of the git repository and inspected the release history to validate our answers to the research questions in the previous section. For RQ1, it could only get answered based on assumption and the number of active participants across time as we counted in Figure 4, but not thoroughly validated. RQ3 is intuitively answered when we match the betweenness distribution with the release history by time, and no further validation is needed. For RQ 2 and RQ4, we have made a detailed validation in this section.

### 5.1 Activity pattern validation

From the mining results, among the 1654 participants with betweenness values larger than 0, we notice that there is a small group of participants who have been central for most of our analysis period (8 participants out of the 1654); another relatively larger group appear without any recurrence (71 participants out of the 1654); the majority would periodically become central in their community (1575 participants out of the 1654). Based on the three activity patterns proposed in RQ2, we confirmed many of our previous suspicions:

#### 5.1.1 Participants that appeared only once tend to be pure users

We look into the git repository to find the files submitted by the 71 participants who have appeared only once in the bug community. We found that only 7 of them have ever committed a change, which means that these 7 are developers rather than pure Android users. The rest do not have commits in the version control system. This verifies our assumption that participants appeared only once in the bug community would more likely to be pure users, as introduced in RQ2.

#### 5.1.2 Participants showed up periodically should be a combination of users and developers

Periodically appearing participants are the majority and we call them phasers. Based on the methodology in Section 3, we looked into the commit history in the git repository in order to verify the expertise of phasers. With as many as 1575 participants, we sampled 156 participants. 21.8% of the sampled participants were developer phasers, who have submitted changes. We studied the expertise of the developer phasers from this sample. All except two of them have worked on specialized tasks that implied some specific kind of expertise or specialization. The rest 78.2% have never submitted files to the development community. They are probably users of Android. Thus, phasers

Table 1: 5 continuously central participants who have submitted changes to the git.

Participant	#Submitted-changes	Related Project
fadden	1259	device_samsung_crespo, platform(bionic, build, dalvik, etc.)
xav	3501	platform(frameworks_base, build, external_bouncycastle, etc.), device_sample,
mbligh	80	kernel(common, experimental, linux-2.6, msm, omap, qemu, samsung, tegra)
ralf (Ralf.-Hildebrandt)	665	kernel(common, experimental, linux-2.6, dalvik, external_libpng, sdk,system_core, etc.)
romainguy	1455	device_htc_passion, device_samsung_crespo, platform(build, cts, dalvik, development, external_bouncycastle, libcore, ndk, apps(AccountsAndSyncSettings, AlarmClock, Bluetooth, Browser, Calculator, etc.), input-methods(LatinIME, iOpenWnn, PinyinIME, CalendarProvider), providers(DownloadProvider, GoogleSubscribedFeedsProvider), wallpapers(Basic, LivePicker, MagicSmoke, MusicVisualization), prebuilt, sdk, system_core)

consist of both users and developers. This answers to our assumption of the phasers' role in RQ2.

### 5.1.3 Participants who were continuously central for a long time period could have multiple areas of expertise

5 out of the 8 participants in this group have submissions in the git. We extracted the projects these 5 participants have submitted changes to, as listed in Table 1. (On the forth row, `ralf` and `Ralf.Hildebrandt` are email alias of the same person, as we observed that the `author_name` attributes are the same for the two email alias.)

Firstly, considering the number of changes they made, all of them except `mbligh` have more than 500 commits within the git, which means that they are quite active in Android development community. This supports that they are experts or advanced developers since more submissions indicates a broader range knowledge about the related techniques.

Moreover, `fadden`, `xav`, and `romainguy` are all working on the platform layer, which includes `build`, `dalvik`, `development`, `framework base`, `libcore`, `sdk`, etc. All of their areas of expertise are related to the platform layer or system core layer.

The participant `romainguy` has experiences modifying almost every component relevant to platforms, including both the apps and the core, and hence should be considered as Android platform development leader.

Furthermore, when investigating these continuous lines we found some partici-

Table 2: 5 clusters we have chosen, out of a total number of 21.

Cluster	Time
1	May 16, 2010 - Jun. 24, 2010
2	Jun. 2, 2010 - Jul. 24, 2010
3	Jan. 13, 2011 - Mar. 3, 2011
4	Dec. 3, 2010 - Jan. 31, 2010
5	Feb. 4, 2011 - May 1, 2011

Table 3: Participants and their areas of expertise in cluster No. 4

ID	Name	Areas Of Expertise
1	charles	kernel - sound, kernel_linux-2.6
2	jasta00	ringtone, media
3	kristoff	driver(net, video, serial, input)
4	rik(rik.bobbaers)	kernel_linux-2.6(mlock)
5	rik(rikard.p.olsson)	kernel_linux-2.6(arm)
6	rik(riku.voipio)	kernel_linux-2.6(arm), driver
7	snp	platform sdk(eclipse plugin)

Table 4: Participants' common areas of expertise of each cluster. Participants number is counted as the number of participants within each cluster who has ever submitted a change and appeared in the git, ie., developers.

Cluster	Participants Number	Areas Of Expertise
1	5	netfilter, driver(video), tests, MIPS
2	13	driver(usb, wireless, mouse), sound, net, i386, performance(tools), input methods
3	9	sound, driver, frameworks_base, tests, platform, kernel
4	7	sound, media, kernel_linux-2.6, driver, platform sdk, kernel video/serial
5	63	net(bluehooth, net driver, ipvx, kernel_linux-2.6), driver(dvd, media, usb, gpu, net), ia64, sound, tests

pants were Google employees, for example, two developers with alias `mbligh` and `romainguy`. Their email account recorded in the git repository is from the “google.com” domain, and moreover, when we googled them, they are indeed introduced as software engineers at Google.

To summarize, this subsection demonstrates that three different centrality patterns correspond to participants of three categories, which supports our analysis hypothesis about activity patterns in Section 4.

Table 5: Highlights of identified clusters from Figure 2

Release	Time	Highlights	Related cluster
v2.2	May 20, 2010	camera and gallery, portable wifi, multiple keyword language, performance(general, browser), media framework, Bluetooth, kernel upgrade, APIs(media, camera, graphis, data backup, device administrator, UI framework)	1
v2.2.1	Jan. 18, 2011	bug fixes(one is about root and unroot), security updates, performance improvements	3
v2.2.2	Jan. 22, 2011	fixed minor bugs, including SMS routing issues	3
v2.3	Dec. 6, 2010	UI refinements, faster text input, power management, NFC, multiple cameras, download management, new multimedia, new developer features(gaming, communication, multimedia, garbage collector, event distribution, video driver, input, native access-audio, graphics, storage, development), linux kernel upgrade to 2.6.36, Dalvik runtime, mixable audio effects	4
v2.3.3	Feb. 9, 2011	NFC, Bluetooth, Graphics, media, framework, speech recognition, voice search, API(identifier, build-in app, locales), emulator skins	5
v3.0	Feb. 22, 2011	UI design for tables, redesigned keyboard, improved text selection, copy and pase, connectivity options(USB, WIFI, media, keyboard, bluetooth), apps update, browser, camera and gallery, contacts, email, development support	5

## 5.2 Cluster validation

As we have discussed above, participants are more active around important releases. Moreover, we can observe from Figure 2 that participants' centrality distributions tend to form into groups or clusters, that often are found around the releases. Participants belonging to the same group become central during the same time periods and then fade away together.

We labeled 21 visible clusters from Figure 2 and looked into five of them which are located more around releases. The five clusters we chose are listed in Table 2.

We extract changes submitted by the members of each cluster from the Android git. (For those who do not have records in the git, we regard them as pure users and do not consider them in this case). After inspecting their submissions, we would get an idea about what kind of tasks they have been mostly working on. Based on release history and the commit logs we found that these clusters tend to be coherent efforts undertaken by multiple kinds of participants.

### 5.2.1 Participants clustered together share similar areas of expertise and tasks

Our analysis in Section 4 shows that the phasers that show up densely together could be interested in similar categories of topics or working on tasks related to the same area.

As described in Section 3, we extract the targets and project names from the git for each member appeared within the cluster. The areas of expertise could be inferred by the contents of the targets and the topics of the projects. We summarized the areas of expertise of participant clusters (from Figure 2) in Table 4.

Inspecting the areas of expertise, we find that each cluster has their own topics, which are relatively different from each other. Also, the topics of each cluster are concentrated to specific layers of Android's architecture. For example, cluster No.1 covers techniques about net filters, drivers, tests, and MIPS, while cluster No.2 is about drivers for connecting devices (usb, wireless, and mouse), net, processor, and input. It is easy to tell that participants of these two clusters are working on different tasks. The other clusters could lead to the same conclusion. Thus we conclude that clusters often exist around a topic.

Take cluster No.4 as an example. There are 7 developers contained in this cluster, as listed in Table 3. It can be observed that work of participants in this cluster could be generally divided into two groups: one is about the Linux 2.6 based kernel, another is related to multimedia. Charles, rik.bobbaers, rikard.p.olsson, and riku.voipio (the pruned bug reporter alias rik is related to three developers in the git and we look them all; this issue would be discussed in Section 6) are all modifying the Linux 2.6 kernel. Charles, jasta00, and kristoff are working on multimedia topic, which includes sound, video drivers, and ringtone.

When we look into other clusters, we get similar conclusions. Thus, from the observation and analysis above, we can conclude that participants with similar centrality patterns often share similar areas of expertise and tasks. This validates our assumption about the phasers being clustered on specific techniques in RQ4.

### 5.2.2 Clusters' working areas of expertise are in accordance with the release contents along the time line

When observing the Android release history, we concluded that the overall betweenness centrality becomes higher around releases, and more active participants appear around important releases, at least according to Figure 4 and Figure 5.

In addition, when taking participants' areas of expertise into consideration, we find that the release highlights are in accordance with the areas of expertise for members of each cluster. Table 5 lists releases and their corresponding clusters together with the highlighted release contents.

Comparing the release contents and the cluster areas of expertise, these two subjects are mostly matched on release topics and cluster's working contents. For example, cluster No.4 covers from December 3, 2010 to January 31, 2011, which occurs before release v2.3. Participants in cluster No.4 have areas of expertise relevant to sound, media, and kernel-video, which match the release contents of new multimedia, APIs for native audio, and mixable audio effects in v2.3; We can also find that 4 out of 7 developers in cluster No.4 have worked on the kernel when the linux kernel was

upgraded to 2.6.36 in Android v2.3.

Cluster No.3 was centered around the releases of v2.2.1 and v2.2.2 (January 18, 2011 and January 22, 2011 respectively). Release 2.2.1 contained security updates and performance improvement; participants in cluster No.3 are specialized mostly on kernels or platforms. This occurs in cluster No.1 and its corresponding release v2.2 as well.

Our conclusion is that participants' work is relevant to areas of expertise associated with clusters, and at the same time, the clusters and participation tends to be correlated with releases. This further validates our answer to RQ4 that developers tend to work as groups on specific projects or issues they are specialized, and their centrality patterns are related to the occurrences of projects or issues.

## 6 Limitations

In this study we explicitly trust that the same account of email addresses, i.e., the part before "@", belongs to the same bug participant. With the given semi-anonymous email addresses in Android bug repository, we pruned the part starting from "..." and kept the front part as the names of bug participants. However, it is possible that some common names share the same start string. For example, "Benjamin Franzke", "Benjamin Tissoires" and "Benjamin Romer" have the same first name. We cannot distinguish these names with the email address "Benjamin@XXX". Besides, some of the email addresses start with a simple letter which is ambiguous identifying a person, while we analyze the results without excluding such data.

We validate our analysis based on the assumption that the types and projects of submitted files reflect the areas of expertise that the developers are specialized in. Hence, we tagged the participants with the techniques according to their submitted files in the Android git. However, there could be inconsistency between the techniques and the submitted files.

Our manual inspection increased the validity of the results, but it still relied on the authors judgment, interpretation and potential bias.

## 7 Conclusion and Future Work

In this paper, we mined the Android bug repository and studied the data of 2010 and 2011. We combined overlapping time windows with social network analysis in order to analyze the participants interactions within the Android bug repository, as part of the Android open source community.

We conducted a thorough case study of the bug reporter activity within the Android bug repository with our method. We analyzed the temporal evolution of the Android bug reporting community both globally and locally. We found that most minor or major releases lead to high betweenness centrality in general. We found and explained sharp changes of participants' betweenness values and we inspected three activity patterns for the participants. Also, we found out that participants tend to get clustered into groups. Then, we validated these results by manually inspecting the Android version control

system (git) and the Android release history highlights. We validated the three activity patterns of bug participants as well as their corresponding reasons. For participants who were clustered in same groups in our plots, they showed interest in a set of similar topics as we inspected in our validation.

Thus we conclude that by combining the SNA with sliding windows, we were able to find many local interactions that would be lost in a global analysis. The sliding windows make these local collaborations more visible, instead of drowning them out in a global analysis. In this case, we can get a more accurate knowledge about participants' working patterns as well as their group working. Furthermore, we validated our findings by inspecting other repositories to confirm that the local behaviour occurred and was of relevance. This work could be used by managers and researchers to produce project dashboards, and automated project status reports.

Future work includes applying the approach in this paper to other open source projects' repositories in order to improve its generality. We want to further validate if our overlapping time windowing SNA plots are trustworthy enough to depict the actual develop processes of various projects.

## References

- [1] K. F. S. Wasserman, *Social network analysis : methods and applications*. Cambridge ; New York : Cambridge University Press, 1994.
- [2] C. Bird, A. Gourley, P. Devanbu, M. Gertz, and A. Swaminathan, "Mining email social networks," in *Proceedings of the 2006 international workshop on Mining software repositories, MSR '06*, (New York, NY, USA), pp. 137–143, ACM, 2006.
- [3] A. Meneely and L. Williams, "Socio-technical developer networks: should we trust our measurements?," in *Proceeding of the 33rd international conference on Software engineering, ICSE '11*, (New York, NY, USA), pp. 281–290, ACM, 2011.
- [4] C. Bird, D. Pattison, R. D'Souza, V. Filkov, and P. Devanbu, "Latent social structure in open source projects," in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering, SIGSOFT '08/FSE-16*, (New York, NY, USA), pp. 24–35, ACM, 2008.
- [5] G. Madey, V. Freeh, and R. Tynan, "The open source software development phenomenon: An analysis based on social network theory," in *Proceedings of the Americas Conference on Information Systems (AMCIS 2002)*, (Dallas, Texas), pp. 1806–1813, 2002.
- [6] 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* (A. Bahulkar, K. Kesavasamy, T. V. Prabhakar, and G. Shroff, eds.), pp. 195–204, ACM, 2011.
- [7] Y. Kamei, S. Matsumoto, H. Maeshima, Y. Onishi, M. Ohira, and K. ichi Matsumoto, "Analysis of coordination between developers and users in the apache community.," in *OSS* (B. Russo, E. Damiani, S. A. Hissam, B. Lundell, and G. Succi, eds.), vol. 275 of *IFIP*, pp. 81–92, Springer, 2008.
- [8] M. L. Bernardi, G. Canfora, G. A. D. Lucca, M. D. Penta, and D. Distanto, "Do developers introduce bugs when they do not communicate? the case of eclipse and mozilla.," in *CSMR* (T. Mens, A. Cleve, and R. Ferenc, eds.), pp. 139–148, IEEE, 2012.
- [9] R. Abreu and R. Premraj, "How developer communication frequency relates to bug introducing changes," in *Proceedings of the joint international and annual ERCIM workshops on Principles of software evolution (IWPSE) and software evolution (Evol) workshops, IWPSE-Evol '09*, (New York, NY, USA), pp. 153–158, ACM, 2009.
- [10] M. Pinzger, N. Nagappan, and B. Murphy, "Can developer-module networks predict failures?," in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering, SIGSOFT '08/FSE-16*, (New York, NY, USA), pp. 2–12, ACM, 2008.

- [11] N. Bettenburg and A. E. Hassan, “Studying the impact of social structures on software quality,” in *Proceedings of the 2010 IEEE 18th International Conference on Program Comprehension, ICPC '10*, (Washington, DC, USA), pp. 124–133, IEEE Computer Society, 2010.
- [12] Y. K. E. Shihab and P. Bhattacharya, “Mining Challenge 2012: The Android Platform,” 2012.
- [13] U. Brandes and T. Erlebach, *Network Analysis: Methodological Foundations*. Springer, 2005.
- [14] A. Hindle, M. W. Godfrey, and R. C. Holt, “What’s hot and what’s not: Windowed developer topic analysis.,” in *ICSM*, pp. 339–348, IEEE, 2009.
- [15] L. C. Freeman, “A set of measures of centrality based on betweenness,” *Sociometry*, vol. 40, pp. 35–41, March 1977.
- [16] R. A. Hanneman and M. Riddle, *Introduction to social network methods*. University of California, Riverside, Riverside, CA, 2005.
- [17] Y. H. Kidane and P. A. Gloor, “Correlating temporal communication patterns of the eclipse open source community with performance and creativity,” *Comput. Math. Organ. Theory*, vol. 13, pp. 17–27, Mar. 2007.
- [18] B. P. T. Iba, K. Nemotob and P. Gloor, “Analyzing the Creative Editing Behavior of Wikipedia Editors: Through Dynamic Social Network Analysis,” vol. 2, no. 4, pp. volume = 13, number = 1, month = mar, year = 2007, issn = 1381–298X, pages = 6441–6456, year = 2010.
- [19] J. B. MacQueen, “Some methods for classification and analysis of multivariate observations,” in *Proc. of the fifth Berkeley Symposium on Mathematical Statistics and Probability* (L. M. L. Cam and J. Neyman, eds.), vol. 1, pp. 281–297, University of California Press, 1967.
- [20] G. Robles and J. M. Gonzalez-Barahona, “Developer identification methods for integrated data from various sources,” in *MSR '05: Proceedings of the 2005 international workshop on Mining software repositories*, 2005.
- [21] M. Goeminne and T. Mens, “A comparison of identity merge algorithms for software repositories,” *Science of Computer Programming*, Dec. 2011.