

Do Communities in Developer Interaction Networks align with Subsystem Developer Teams? An Empirical Study of Open Source Systems

Usman Ashraf
Christoph Mayr-Dorn
Atif Mashkoor
Alexander Egyed

Johannes Kepler University, Austria
firstname.lastname@jku.at

Sebastiano Panichella
Zurich University of Applied Science, Switzerland
sebastiano.panichella@zhaw.ch

Abstract—Studies over the past decade demonstrated that developers contributing to open source software systems tend to self-organize in “emerging” communities. This latent community structure has a significant impact on software quality. While several approaches address the analysis of developer interaction networks, the question of whether these emerging communities align with the developer teams working on various subsystems remains unanswered. Work on socio-technical congruence implies that people that work on the same task or artifact need to coordinate and thus communicate, potentially forming stronger interaction ties. Our empirical study of 10 open source projects revealed that developer communities change considerably across a project’s lifetime (hence implying that relevant relations between developers change) and that their alignment with subsystem developer teams is mostly low. However, subsystems teams tend to remain more stable. These insights are useful for practitioners and researchers to better understand developer interaction structure of open source systems.

Index Terms—developer interaction network, system modularity, subsystem coordination, developer communities.

I. INTRODUCTION

Global software development is carried out by developers located in various parts of the globe. They face a number of challenges in relation to communication and coordination due to the distances involved in three dimensions – geographical, temporal, and socio-cultural [1]. Open source software development is considered as a successful example of large scale global software development [2]. In open source development environment, software systems do not have a pre-assigned organizational structure as developers can contribute to any part of the system. However, developer communities are self-organized within development teams [3]. Studies have shown that the developer communities organization has a significant impact on software quality [4]–[7].

Motivated by the aforementioned insights, researchers have built automated tools analyzing developers interaction data from various sources, such as version control systems, mailing lists, or issue trackers [3], [8]–[12], to investigate how “*emerging development teams*” are formed in open source projects.

However, a few researchers observed how the emerging “*community structure*” aligns with “*subsystem developer teams*”. Bird et al. [3], for example, hypothesized that files edited by developers of the same community are placed “closer together” than the files edited by randomly picked developers. Results, however, were inconclusive.

We argue that observing the alignment of developer communities with the system architecture can be more accurately studied at the level of coarse-granular subsystems, i.e., the components of a system, rather than analyzing these aspects at the level of fine-grained artifacts. In this regard, the work of Lenarduzzi et al. [13], for example, is important as they point out the potential negative impact of team independence at the subsystem (microservice) level. Additionally, Nagappan et al. [4] have shown that organizational metrics are the top bug proneness predictor in an industrial setting. Given the lack of explicit organizational structures in open source systems, we suggest to utilize implicit structures (i.e., communities)¹.

Take the excerpt from a fictive project in Fig. 1 as an example. The *developer interaction network* on the top exhibits edges based on being involved in the same issue via activities such as reviewing or commenting. A community detection algorithm assigns developers to communities (blue borders) based on the developers’ interaction intensity (indicated by line thickness). Developers who interact often tend to form communities. The lower part of Fig. 1 depicts files and issues belonging to three subsystems A, B, and C. Subsystem team members are then those developers that changed an artifact belonging to that subsystem, respectively, are active in an issue referring to this subsystem. Note that subsystem member sets may overlap as in open source projects key developers are often involved in multiple subsystems. In our example, developer Dave is a member of all three subsystems, Carol is a member of subsystem A only. In this example, when comparing the members of subsystem B (i.e., Dave and John),

¹We define a community as the set of developers that share responsibility or interest such as working on the same subsystem(s) [12].

we notice that they are members of different communities and there is no edge between them. This could be a warning signal about poor communication and coordination between the two.

To shed light into how communities and subsystems are aligned, we empirically investigate in this paper, in the context of 10 open source projects, how communities emerge and change over time – e.g., how developers join and leave sub-communities – and the extent to which these community patterns match the subsystems evolution. We find that developer communities change considerably across a project’s lifetime (hence implying that relevant relations between developers change) while subsystem developer teams (SDTs) remain comparatively stable. Overall, the community alignment with SDTs is often low, which implies that developers maintain significant communication ties with developers outside their (subsystem) work scope. We hypothesise that such an interaction network independent from subsystems emerges from the need to remain robust against the disruption of leaving developers and quick onboarding of new members.

The primary contribution of this paper is an empirical study investigating the evolution of developer communities and their alignment with the subsystems developer teams (Section IV). The secondary contributions of this paper are: i) a technique for measuring alignment (i.e., overlap) between subsystems and developer communities (Section IV(C)), and ii) a technique for determining developer communities evolution (Section IV(D)).

The paper is structured as follows. Section II introduces the research questions. Section III describes the applied data gathering method and the resulting evaluation data set. Section IV explains our community-SDT alignment and evolution measurement technique with Section V presenting the results. We discuss the obtained results and their implications in Section VI. Section VII compares this study to state of the art approaches before Section VIII, which concludes the paper with a summary and an outlook on the future work.

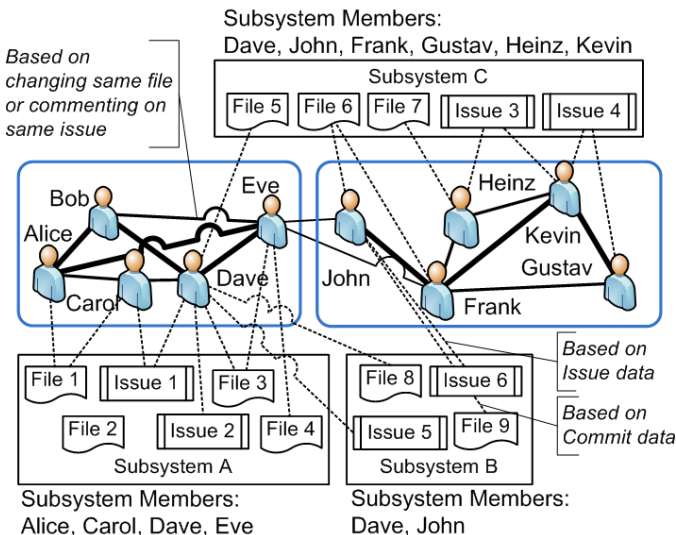


Fig. 1. Example developer community to subsystem alignment.

II. STUDY DESIGN

Developers that work on the same task, artifact, or subsystem need to coordinate and thus communicate in order to avoid incomplete change propagation, rework, or duplicate work. We therefore would expect them to form stronger interaction ties [14] than developers working in different subsystems. We hypothesize that these developer interactions give rise to a community structure. Therefore, the *goal* of this study is to obtain first insights into the alignment of sub communities within open source projects with the systems’ structure.

The *motivation* behind this study is to determine the extent to which emerging open source developer communities reflect the subsystem structure, proposing mechanisms for studying such communities aspects. The *perspective* is of practitioners and researchers that could leverage such mechanisms to identify subsystems where developers coordinate insufficiently.

The *context* of this study comprises commits, issues, source code (folder) structure, and the conceptual and structural links among the data. The selected 10 open source projects are part of our publicly available and published dataset [15].

Research Questions:

In this work, we investigate the following research questions:

RQ1: To what extent can we identify well defined developer communities across the investigated projects’ lifetime? Our goal is to propose a systematic approach able to detect emerging communities by analyzing the *developer interaction network* generated from issue involvement and commits linked to those issues (see Fig. 1). We are then interested in knowing how often and how many well defined communities we detect.

RQ2: Do the developers active in the same subsystem emerge in the same development communities? In other words, measuring the overlap of the SDTs with the detected communities allows us to identify subsystems where the communication among developers occurs mostly within the subsystem.

RQ3: How stable are the detected communities across time compared to the SDTs? While RQ1 and RQ2 investigate *developer interaction network* and its alignment with SDTs at separate snapshots in time, here, we inspect the stability of the detected communities and the SDTs across time. We analyse whether a pair of developers that belong to the same community in one time window also belong to it in the subsequent time window. Likewise, we are interested to know whether a pair of SDTs that emerge in the same developer community does so in the subsequent time window.

Overall, answering these questions provides insights into whether subsystems in open source projects represent decoupled work scopes that result only in limited coordination overhead compared to work coordination within a subsystem (as measured by interaction ties).

III. DATA GATHERING AND PREPROCESSING

A. Open Source Project Selection

The project selection (Fig. 2 (1)) is the first step of our study approach. We select the candidate project according to the following requirements:




Fig. 2. Steps of the study approach.

Subsystem Structure We manually selected projects which exhibit a non-trivial (i.e., at least 10 subsystems) and clear subsystem structure to avoid introducing potential bias by splitting the system into subsystems that do not reflect the real underlying decomposition.

We manually inspected the projects on Apache’s Jira server² that exhibited a significant number of issues with a *component* property set and where the top level source code folders (hosted on Github) closely (or identically) matched those component names. We interpreted those components/folders as subsystems.

We avoid selecting projects from a single mechanism for structuring subsystems by also selecting Github projects that manage code across multiple repositories (multi-repo), i.e., one repository per subsystem, as our selected Jira projects are mono-repo, i.e., one folder per subsystem.

Developers Further selection criteria included a minimum of 40 participating developers over the project’s life time to guarantee that even with heavy fluctuation of developers, that there are sufficiently many developers to form an interaction network with sub-communities. Note that for our selected projects the number of developers are well over 40.

Commits and Issues Furthermore, we filtered out projects with less than 1500 commits or 200 issues. This ensures that a relation between two developers based on being active in the same issue or committing an artifact update linked to the same issue may occur sufficiently often to signify a meaningful relation between the developers and not just an one-off occurrence.

This study focuses on 10 projects listed in Table I. The projects are limited in number and size as (i) there is a manual processing effort required in step 1 (Fig. 2) and (ii) non-negligible manual effort is necessary to investigate folder structure and confirm developer deduplication.

B. Data Extraction and Storing

Having identified 10 projects under investigation, we extracted commit and issue information (Fig. 2 (2)). We used *Perceval* [16] to extract commit and issue information from Github and the Jira python client (*Perceval* lacks to provide our desired Jira issue information) for extracting data from Apache Jira. Both these tools provide data in the JSON format.

Commits, the involved artifacts, and issue details make up the core information of our data set. In our datamodel, an issue may represent a Jira issue, a Github issue, or a Github pull request, or even a combination thereof whenever, for example, a pull request references a Jira issue.

A developer participates in a project in various ways: for example, from committing code changes, to commenting on issues, to reviewing pull requests. We introduce two *Involvement* types to harmonize activities across Jira and Github. *Contributing* is equal to committing artifacts (i.e., as indicated by a commit linked to an issue); *Informative* describes input to an issue such as having reported it, commented on it, or having reviewed artifacts. We encode these two *types* of actions as integers of value 3 and 2, respectively, to reflect the amount of effort behind the activities. Before continuing with these values we performed a value sensitivity analysis by assigning them values (4,2) and (1,1) respectively. We observed little change in the detected community structure with relatively less quality. More details of community structure and its quality metric are in the Section IV(A).

C. Developer De-Duplication

When recording a developer’s involvement across commits and issues, we need to take care of situations where a developer uses multiple accounts (email-ids) on Github. More importantly, a developer’s id on Jira does not match the Github account id of the same developer as Jira does not use email addresses as a part of a user id while Github does. This will result in inconsistencies as one developer will appear

²<https://issues.apache.org/jira/>

TABLE I
OVERVIEW OF THE TEN ANALYSED PROJECTS.

Project	Type	Programming Language(s)	Time Period [Months] (From-To)	Commits	Devs	Subsys	Arts	Issues	Arts Linked to Subsys [%]	Issues Linked To Subsys [%]	Issues Linked To Commits [%]	Commits Linked To Issues [%]
lagom	Multi-Repo	Scala, Java	44 (Mar 16-Nov 19)	6089	540	17	27490	3381	100	100	59	62
nameko	Multi-Repo	Python	85 (Sep 12-Nov 19)	3861	233	14	960	867	100	100	65	93
kumuluz	Multi-Repo	Java, JavaScript	54 (May 15-Nov 19)	2487	82	19	2083	274	100	100	45	22
jhipster	Multi-Repo	Java, JavaScript	71 (Nov 13-Nov 19)	10391	767	11	7146	3267	100	100	64	59
networknt	Multi-Repo	Java, JavaScript	38 (Sep 16-Nov 19)	5369	110	12	16127	2005	100	100	74	36
flume	Mono-Repo (Jira)	Java	115 (Jun 10- Jan 20)	2623	1199	16	3403	3638	51	57	29	86
stanbol	Mono-Repo (Jira)	Java	107 (Nov 10-Nov 19)	6953	149	21	17469	1490	54	80	47	62
falcon	Mono-Repo (Jira)	Java	87 (Nov 11- Mar 19)	2556	190	19	7362	2760	56	38	47	63
tika	Mono-Repo (Jira)	Java	154 (Mar 07-Jan 20)	5848	1327	21	6509	3335	71	65	46	67
openjpa	Mono-Repo (Jira)	Java	164 (May 06- Jan 20)	7215	831	27	10696	2849	54	73	57	68

as multiple individuals in the data set, hence the need for deduplication (Fig. 2 (3)). We used the *Dedupe* library [17], which uses machine learning to perform deduplication. It has an accuracy up to 95% [18]. More details of the training process are available in our dataset [15].

D. Extraction of Commit-To-Issue Relations

Github commit messages often refer to issues to identify the purpose of the changes. We manually inspected each project to identify what patterns developers tend to apply for referring to issue-id and applied that pattern when parsing the commit messages (Fig. 2 (4)). The following are two sample commit messages from our data referring to issue-id: FLUME-3311 and kumuluz#115 respectively.

- “FLUME-3311 Update User Guide In HDFS Sink”
- “Merge pull request #115 from Jamsek-m/master”

We applied the same mechanism to pull requests. Whenever a pull request identifies a Github issue, we merge the pull request and the Github issue into one dataset issue and link all commits that are part of the pull request to it. This is particularly helpful when commit messages lack a reference to an issue.

For the Jira-backed projects, we regularly find pull requests and commits without a reference to a Github issue but instead a reference to a Jira issue. Whenever a pull request identifies a Jira issue, we subsequently merged them into a single issue in our dataset (we never encountered multiple pull requests referencing the same issue) and link all involvements from the pull request and Jira issue to that merged Issue. We proceed likewise whenever we find a Github issue referencing (in the comments or title) a Jira issue.

E. Mapping Subsystems

For those projects that utilize Apache Jira for issue management, aside from mapping Jira to Github users, we also need to map a Jira project’s components to the corresponding folders on Github (Fig. 2 (5)). We manually mapped each component of the five Jira projects to its corresponding Github folder or multiple folders where necessary. We link all those folders, which do not map to a component to the main system (we store main system as a subsystem in our data model).

As briefly mentioned above, the mapping for Github multi-repo projects is straight forward: each repository in the project becomes a subsystem. Furthermore, we store the list of folders that make up a subsystem. This allows us to relate all artifacts to a unique subsystem.

F. Linking Issues to Subsystems

Linking issues to subsystems (Fig. 2 (6)) works similar to relating artifacts to subsystems. Jira issues exhibit a *Component* property that identifies all affected subsystems (potentially multiple). Github multi-repo projects provide separate issue lists for each repository, thus we link those to subsystems unambiguously.

G. Dataset Overview

Table I provides an overview of the 10 chosen open source projects. The columns *Type* and *Programming Language(s)* report the repository structure and languages of the project, respectively. The column *Time Period* describes the overall time in months in which we extracted commits and issues. All subsequent rows report values from this time period. The *Commits* column reports how many commits are made in total in a project. Recall that in multi-repo projects, commits per subsystem belong to one repository, thus the reported numbers are the sum of commits over all repositories in such a project. In case of mono-repo projects, all commits come from a single, main repository. The *Devs* column provides the total number of developers in each project. We consider any person a developer who is *Contributing* or *Informative* in at least one subsystems. *Subsys*, *Arts*, and *Issues* columns report the number of subsystems, total source code artifacts, and issues in a project, respectively.

The *Arts linked to Subsys* column shows the percentage of artifacts belonging to subsystems. In a Mono-Repo system each folder has its link to a subsystem or to the main system. So in this case, the column shows the percentage of artifact which are linked to subsystems and not to the main system. Each repository (subsystem) of a Multi-Repo system has its own Issues, thus, the *Issues Linked To Subsys* column shows their 100% link to subsystem. In a Mono-Repo system, we also include issues from Jira, where they have specified mapping to a subsystem. As mentioned in the data preprocessing section, commit(s) are made against an issue and issue(s) are linked against commits. The Columns *Issues Linked To Commits* and *Commits Linked To Issues* specify the percentage of such linkage, respectively. Note that as the linkage is not always present e.g., commit messages not always refer to issues to identify the purpose of the changes. Therefore, the linkage percentage ranges between 22 to 93.

IV. A TECHNIQUE FOR MEASURING COMMUNITY-SDT ALIGNMENT AND EVOLUTION

A. Developer Interaction Network Extraction

RQ1 requires building a developer interaction network. We built such a network from developers' involvement in issues such as contributing code or commenting. We sliced the project lifetime into time windows of 4 months, as done in close-related works [10], [12]. We presume that such a time window is sufficiently long to allow developers to repeatedly interact to produce a reliable interaction network while being short enough to capture changes of the interaction network as developers join and leave.

Developer Interaction Network Model We model the developer interaction network (Fig. 2 (7)) as an undirected graph with developers as nodes ($d \in V$) and edges ($e(d_i, d_k) \in E$) between those developers that interact. We add an edge between two developers when the two developers are involved in the same issue, for example, one developer adding a comment to an issue, the other developer committing an artifact update linked to the same issue. The edge's *issue involvement intensity* measures in how many issues the two developers are active in, weighted by their involvement type.

Developer Involvement Recall that the developer involvement types – *Contributing* and *Informative* – carry a weight of 3 and 2, respectively. We consider only the highest weighted involvement per developer and issue $inv_{max}(d, i)$. The intensity for a single issue i_h and pair of developers d_i and d_j (i.e., the per-issue score) is then the lowest of the two scores $\min(inv_{max}(d_i, i_h), inv_{max}(d_j, i_h))$. For example, the per-issue score between a contributing developer and an informative one is 2. The overall intensity for a pair of two developers across all their common issues (i.e., the *issue involvement intensity* edge property) is then the sum of per-issue scores ($\sum_k \min(inv_{max}(d_i, i_h), inv_{max}(d_j, i_h))$). The *issue involvement intensity* serves as an edge weight.

Community Detection The subsequent step after forming the developer interaction network is to discover communities in it. The communities are the natural divisions of network nodes into densely connected subgroups [19]. The developer interaction graph serves as the input to the step of community detection (Fig. 2 (8)). We first remove all developers who have no edges to any other developer. We consider those as not part of any community and retaining them in the graph would negatively influence the community detection results.

For the purpose of community detection, we use the order statistics local optimization method (OSLOM) [20]. This method is also used in close-related work [12]. *OSLOM* finds statistically significant communities and has distinct features such as to handle weighted graphs, to form overlapping communities, and to distinguish communities from pseudo-communities.

Community Verification and Validation The *OSLOM* algorithm identifies the best distribution of graph nodes as communities after calculating the statistical significance (the probability of finding the cluster in a random null model, i.e.

in a class of graphs without community structure) of each community. It uses the significance as a fitness measure to evaluate and include a community candidate in the resulting set of communities. We additionally checked the resulting quality of the community structure using conductance as follows.

Quantitative Community Quality Metric: We apply conductance [21], which measures the probability of having an edge leaving the community. A community with all edges connecting only member nodes has conductance 0 (an isolated community) while a community with no edges amongst its members has conductance of 1 (arguably not really a community at all).

Qualitative Detected Communities Inspection: In order to validate that the detected communities accurately reflect real-world developer collaboration, two of the co-authors manually inspected *OSLOM*'s detected communities of one of our projects, i.e., *kumuluz*. The project has 82 developers and 54 months of observed developers interactions. We inspected the detected communities for 5 time windows (20 months). We verified each detected community of every window by matching the strength of its collaborative relations compared to the other undetected communities. We also observed the correctness of developers division into communities based on edge weight.

The resulting communities consists of a subset of developers ($c_k \subset V$) where developers potentially are members of more than one community (overlapping communities). This is an important property as in developer interaction networks, the most active developers are often connected to two or more communities [22]. In the absence of overlapping communities, a developer will be placed into a single community which will skew the subsystem developer overlap calculations.

Observed Time Windows We apply the community detection algorithm for each time window separately. We skip time windows with less than ten developers as we do not expect to find meaningful communities among such a small number of developers. Manual inspection of sample windows with less than 10 developers confirmed these expectations. For all these windows the best community structures found by *OSLOM* exhibited high conductance values (> 0.8).

B. SDT Extraction

Answering RQ2 requires identifying for each subsystem ($s_y \in S$) who are the active developers: the SDT ($SDT(s) \subset V$) (Fig. 2 (9)). To this end, we simply consider a developer as active if one has at least two *Contributing* or five *Informative* involvements in that particular subsystem within the defined time window. This ensures we select only developers that have a longer-running interest in a subsystem and not just include anyone with minimal involvement. We continued with these values i.e., two and five after observing the trend of the number of *Contributing* and *Informative* involvements in the chosen projects.

In contrast to the interaction network community construction, a developer becomes member of an SDT purely based on one's contributions to source code and/or involvement in issues

regardless of interactions with other developers. For example, two developers changing the same artifact within the same time window that otherwise are not involved in any issue will end up in the same SDT.

Similar to developer overlapping communities, multiple SDTs may contain the same developer. Especially key developers are often involved in multiple subsystems. At the same time, subsystems may be too small to give raise to a dedicated SDT that focuses only on that single subsystem. Ultimately, for each developer we define the community membership $cm(d) \in [1, k]$ and SDT membership $SDTM(d) \forall S$.

C. Overlap Calculation

The second part for answering RQ2 is a metric for measuring how well a SDT matches a developer community (Fig. 2 (10)). In a project where the community structure represents subsystems, we would expect that members of the same SDT are also members of the same community (low membership heterogeneity). On the other side of the spectrum, we would find a SDT where every member belongs to a different community (high membership heterogeneity). We measure each SDT's membership heterogeneity using the normalized Shannon entropy [23] $mh(SDT(s)) = -\sum_k (p_k \ln p_k) / \ln k$ where p_k is the number of developers in $SDT(s)$ being member of community c_k and \ln is the natural logarithm. The normalized Shannon entropy provides a result in the interval $[0, 1]$: 1 when p_k is the identical for every k and 0 when $p_k = 0$ for all but one k . In other words, $mh(SDT(s)) = 1$ when all SDT members are exactly equally distributed across all communities, and $mh(SDT(s)) = 0$ when all SDT members are member of the same community. In our fictive example in Fig. 1 $mhSDT$ for Subsystem A is 0 and $mhSDT$ for Subsystem B is 1. Ideally, we observe minimal/low membership heterogeneity for all subsystems.

Having introduced the general concept of membership heterogeneity, we need to outline an adjustment to p_k . Without any adjustment each developer has equal impact. However, this does not properly reflect the different developer types in a typical open source project. There is a small number of key developers that get involved in issues and source code changes across the subsystems, subsequently being well connected and informed. Giving them equal weight as a developer who is focused only on one single subsystem would skew the measure. We thus redefine p_k as the weighted sum of developers in $SDT(s)$ being member of community c_k where the weight describes the developer's focus $f(d) \in [0, 1]$ across the subsystems (no other adjustment to $mh(SDT(s))$ needs to be made). The focus on a single subsystem is maximal when the developer is active in a single subsystem, and minimal when involved in every subsystem. To this end, developer focus is calculated as 1 minus the normalized Shannon entropy of the developer's involvement in each subsystem (see Subsection IV-B); hence $f(d) = 1 + \sum_y (p_y \ln p_y) / \ln y$ with $p_y = \sum_{i \in y} inv_{max}(d, i)$.

D. Membership Evolution

Regarding RQ3, we expect community structure as measured by conductance, and subsystem alignment as measured by SDT membership heterogeneity to change over time. Focusing on community structure, we want to be able to interpret community evolution not only through changes in conductance but also by looking at the external stability of developers (joining and leaving developers) and internal stability (shifting of developers from one community to another in two consecutive time windows w_t and w_{t+1}) (Fig. 2 (11)). Measuring external stability is simply a matter of tracking which developers appear as nodes in the interaction network graph in one time window but not in the next and vice versa. As the detected communities have no meaningful, intrinsic identity (we just assign them a number as identifier) internal stability requires tracking which developers are in the same community in one time window and remain in the same community (irrespective of the community identifier) in the next window. To this end, we counted how many pairs of developers were in the same community in w_t and are again in the same community in w_{t+1} ($\#sameC$) or were in different communities in w_t and are in different communities again in w_{t+1} ($\#diffC$). The internal stability metric is then $iStab = (\#sameC + \#diffC) / totalPairs$ and yields 1 for full stable communities and 0 when all communities become completely reshuffled.

Similar to developer community membership evolution, we are interested to know whether the SDTs that overlap with the same community in one time window will overlap with the same community in the subsequent time window (Fig. 2 (11)). We apply the same calculation as for internal stability above but instead of comparing a pair of developers being in the same or different community again, we measure in $sdtStab$ whether pairs of SDTs tend to overlap with the same community in the subsequent time window again or whether they overlap with different communities.

V. RESULTS

In this section, we present the detailed results for a single example project *flume* only due to page limitation. We thus provide aggregated numbers across all ten projects and refer for per-project details to our supporting online material (SOM) [24]. The SOM includes a database dump of the underlying dataset, the source code used for extracting the raw data, preprocessing, and metric calculation, as well as detailed figures for each project.

A. Answering RQ1

For answering [RQ1]: *To what extent can we identify well defined development communities [...]?* we analyse for each project how many communities we find for each time window and their conductance. Fig. 3 displays the count, size, and average conductance of communities found for the example *flume* project across the observed time windows (here twenty six windows (window-20 has less than 10 active developers therefore not included). The number of communities ranges between one and eight communities. Note that as communities

are latent – they do not have an explicit identifier but can be observed – community 4 in window 11 might be most similar to community 5 in window 12 in terms of overlapping members. The size of a community, i.e., number of developers in it, varies as well. As shown in Fig. 3, often communities exceed typical team sizes of ten members (as also previously observed [12], [25]) as communities also include rarely and non-contributing members. These rarely-contributing members are significant for estimating software quality (e.g., defect prediction) [25].




Fig. 3. Size of the detected communities and average community conductance for the *flume* project.




Fig. 4. Range of conductance per project. For each project, the x-axis reports also the number of analysed time windows with at least 10 active developers out of all time windows.




Fig. 5. Range of community count per project.

For comparing the investigated projects, Fig. 4 provides a boxplot for each project describing the range of conductance over all time windows that have at least 10 active developers.




Fig. 6. Range of average SDT membership heterogeneity per project.

We notice that the mean conductance for each project is below 0.4 except *networknt* and *falcon*, with no time window of any project exceeding a conductance of 0.6, a range similar to previous observations for such a developer count [12].

Fig. 5 reports the range of community count k . For all projects the number of community count ranges between one and eight.

Observation 1: There are far fewer communities that emerge from the developer interaction network than subsystems (e.g., *flume* has 16 subsystems).

Observation 2: Overall, communities tend to have medium to low conductance, thus the observed developer interaction networks exhibit a clear community structure.

B. Answering RQ2

For answering [RQ2]: *Do the developers active in the same subsystem emerge in the same development communities?* we calculate SDT membership heterogeneity for each subsystem in each observed project for all time windows with at least 10 active project developers and at least two communities (heterogeneity between a single community and SDT is 0).

Fig. 6 compares the range of average SDT membership heterogeneity across all intervals per project, i.e., we averaged the SDT membership heterogeneity metric over all subsystems where that metric was calculated. The box plots in Fig. 6 then describe for each project the range of that heterogeneity average when calculated for each time window.

We notice a wide range of heterogeneity behavior across the projects. While *nameko*, *networknt*, *flume*, and *falcon* display low average heterogeneity, *jhipster*, and *stanbol* yield rather high average heterogeneity. The remaining projects yield a medium range of mean average heterogeneity, yet have the occasional time window where average heterogeneity reaches high values, i.e., up to 0.7.

With significant spread of heterogeneity values across subsystems (in *flume*, for example, between 0.0 and 0.71), we investigate whether high heterogeneity might be correlated with large SDTs (i.e., the more members in a team, the more likely they come from different communities). The data show a weak correlation between SDT size and heterogeneity with a Pearson's correlation coefficient value of 0.46. We observe this phenomenon also in the remaining projects.




Fig. 7. Heatmap of SDT size for the project *flume* for all time windows with at least 2 clients.




Fig. 8. Heatmap of SDT conductance value for each subsystem (with at least 2 active developers) in the project *flume*.

We subsequently investigate whether SDTs form their own, small communities that OSLOM did not detect. Such communities may form when developers interact primarily with other developers working on the same subsystem and not so much with developers of other subsystems. To this end, we treat each SDT as the members of a virtual community and calculate its conductance.

Fig. 7 displays SDT size for each subsystem in the example *flume* project. For easier interpretation, conductance values are provided in Fig. 8 next to it. We observe that SDTs exhibit high conductance values (close to 1 and darker in color), thus they are tightly integrated in the developer interaction network and hence do not represent well defined communities. Recall that, a community with no edges amongst its members has conductance of 1 (arguably not really a community at all).

Observation 3: Projects are diverse with respect to how well SDTs typically overlap with a single community. All projects exhibit intervals in which SDTs experience high heterogeneity.

Observation 4: SDTs do not form subcommunities in the developer interaction network but rather maintain considerable interaction ties with developers not involved in the particular subsystem.

C. Answering RQ3

For answering [RQ3]: *How stable are the detected communities across time compared to the SDTs?* we need to measure internal stability of the communities between two subsequent intervals to understand how much the developer interaction network structure itself changes. We then compare the internal stability with the SDT stability.

Fig. 9 reports the evolution of the example project *flume* developer community membership as measured by external and internal stability (interval t reports the change from interval t to $t+1$). We notice a heavily fluctuating number of unchanged developers (between 0 to 27) while an even higher number of developers keep joining (between 5 and 58)




Fig. 9. Evolution of Developer Community Membership for example project *flume*: number of developers leaving (dotted line), joining (dash dotted line), and remaining (dashed line) as measured between subsequent intervals (right y-axis), and internal stability as well as SDT stability (left y-axis).

and leaving (between 9 and 51). This comparatively large set of fluctuating developers influences the community structure, hence few stable key members substantially shift between newly forming communities. This is reflected by the internal stability metric drifting between 0.27 and 0.65. The SDT stability remains above that between 0.65 and 0.91.

For every project we select the time windows with at least 10 active developers. We then determine the internal stability of developers across communities (Fig. 10). We also determine the stability of SDT overlap with communities (Fig. 11): i.e., whether two SDTs mostly overlapping with the same community in time window t do so also in time window $t+1$.

Observation 5: In general, we find that internal stability is not that strong, hence implying that communities change considerably across time. In contrast, we find considerably stronger SDT stability. This indicates that developers remain in the same SDT but switch among communities (i.e., pre-




Fig. 10. Range of internal stability per project.




Fig. 11. Range of SDT stability per project.

dominately interact with different members across intervals).

VI. DISCUSSION AND IMPLICATIONS

The findings of our study are that well defined communities emerge for time windows with ≥ 10 developers. We observe that the alignment between these communities and SDTs covers the complete spectrum from high to low overlap with no discernible influence of the SDT size. Hence, we draw the conclusion that there is no correlation between communities and SDTs: working on a subsystem does not cause interaction links to occur predominantly among SDT members.

We make the careful hypothesis that the low overlap is to mitigate the high developer fluctuation. Interaction links among members from multiple SDTs enable to better cope with situations when developers leave. Then the burden to take over or on-board new members is not placed only with members of that SDT. Cross-SDT interaction may thus be a mechanism to remain robust in the presence of high fluctuation. The low community stability can then be explained as the remaining developers rearrange (at least partially) their interactions and coordination around newly joining members (and the effect of leaving members).

The significant interactions across SDT boundaries also raise the question whether a microservices-centric approach may be suitable in the context of open source software development. The perceived advantages of microservices, among others, are a reduction of coordination needs across microservice-centric teams (i.e., SDTs) [13]. Assuming our hypothesis outlined above is correct, the open source community might benefit, even need, such cross-SDT interaction and hence might not be able to benefit fully from a microservice-centric approach. We are not expecting such high overlap in industrial settings – a subject of future studies.

Researchers and Practitioners Implications Our work provides a mechanism to measure development communities and their overlap with SDTs. We provide evidence of heterogeneity of development communities and where evolving sub-communities do not overlap with the subsystem compositions. This approach is potentially useful for practitioners, e.g., lead developers in open source projects, to assess the interaction structure among developers and potentially identify subsystems where teams are less well connected. Similarly, practitioners in industrial setting may find the approach useful

to identify subsystems where teams are highly interacting, thus perhaps identifying inadvertently strongly coupled subsystems.

A. Threats To Validity

Construct Validity We generated the developer interaction network from issues and commits linked to those issues to identify developers that coordinate their work. We did not contact developers to verify the detected community structure as we do not expect them to accurately recall differences in community structure for fine-granular time windows of 4 months for several years into the past.

Besides issue trackers and commits, we did not consider other possible developer communication channels, i.e., mailing lists, IRC, and conference calls, which could impact the links created among developers. We investigated that among our analysed projects only Jira-backed projects use mailing lists. Panichella et al. [11] investigated how developers collaboration links vary when data is gathered from different sources. They found, on the one hand, that communication links obtained from mailing lists have high overlap with links obtained from issue trackers and, on the other hand, that emails as the primary communication channel is increasingly replaced by chat and issue trackers. Other communication channels such as IRC and conference call are also used in practice. Four of the analysed projects use a chat platform (i.e., *Gitter*³ where conversations often happen in private threads which cannot be mined). However, research [11] has shown that mining links from chat is less reliable as 1) this tends to produce too many links, and 2) conversations are less easily associated with issues. We therefore believe that we miss only negligible information by leaving out mailing list and chats. We, however, like to point out that co-located developer interactions are not considered and hence restrict the applicability of our approach to projects with completely distributed developers.

Internal Validity In order to address the internal validity threat, we analyse data from multiple open source systems rather than conducting controlled experiments. The analysis focused on commits and issues in general and was not specifically tailored to Java projects. We sampled the chosen projects based on different subsystem structure mechanisms: mono-repo and multi-repo based to avoid a single mechanism influencing the SDT stability or evolution thereof.

³<https://gitter.im/>

Some of the chosen projects show a low percentage of linkage between commits and issues, thus a threat to the validity of the study. However, we observed similar community quality metric results of 10 additional projects with higher linkage values (the results are a part of the SOM [24]). These projects are not included in our study as they do not completely fulfill other inclusion criteria mentioned in the Section III.

We cannot be absolutely certain that the mapping between Jira issue components and Github folders is absolutely correct. A few Github folders for the five Jira-backed projects that could not be mapped to a component with high confidence were excluded to avoid skewing the results.

External Validity Generalizability of our observations is limited to open source projects that make significant use of issues (incl. pull requests) as their primary coordination and communication mechanism. We do not expect our insights to apply to commercial projects where developer teams are top-down defined rather bottom-up emerging. The developers in commercial projects are typically (at least partially) collocated. Even when a significant amount of communication happens online, a significant amount of coordination is expected to occur off-line, thus resulting in a less accurate dataset with respect to communication.

VII. RELATED WORK

Several researchers have analyzed developer interaction data from various sources, such as version control systems, mailing lists, or issue trackers, to investigate how emerging development teams are formed in open source projects.

Joblin et al. [12] provide an approach to identify the community structure of a software project. They capture a view on developer coordination, based on commit information and source-code structure. We applied the same algorithm for identifying communities, however, we applied developer interactions based on their involvement in issues along with commit information. Moreover, they capture a macro-level view (once for the overall project) of coordination compared to our approach which focuses on a more fine granular level (developer interaction network alignment with SDTs at separate snapshots in time).

Researchers also examined the effect of interaction between developers on software quality. Tamburri et al. [26] [27] explored the relation between community smells and code smells in open source environments. They base the detection of community smells on the micro-granular structural differences (i.e., non/existence of edges between individual developers) between the collaboration network generated from commits and the network derived from developer interactions in issues or mailing list. They also propose YOSHI [28], a tool to monitor key community traits in open-source projects. In contrast, our approach observes the community level but not patterns of edges around individual developers.

Bird et al. [25] examined the effect of different ownership measures on release failures in an industrial setting. They found that social network metrics are useful predictors.

Leibzon [29] studied the organization of software development teams and project communities at Github. Nzeko'o et al. [30] made a social network analysis and comparison of developers' and users' mailing lists of four open source software projects. Similarly, Bird et al. [3] analysed the latent social structure by forming a social network from the project's mailing list. They showed results that sub-communities arise within a project as the project evolves.

Several efforts aim to exploit the information embedded in the social structure. Canfora et al. [31] mined explicitly mentioned cross-system bug resolutions and correlated these activities with social attributes of developers who participated in discussions, e.g., developer mailing lists and commits in source code. Their study showed that cross-system bug fixing mainly involves developers who engage the most in mailing list interaction and developers who are among the top committers. Mockus et al. [32] used email archives of source code change history and issue reports to quantify aspects of developer participation, core team size, code ownership, and productivity for open source software projects.

Panichella et al. [33] studied how emerging teams evolve over time and that these teams tend to work on more structurally and semantically related set of files. Whether these files belong to or represent a particular subsystem was not part of the study. Hong et al. [34] also conducted a study to understand developer social network and its evolution. They observed that developers and their relationships change continually. Avelino et al. [35] studied how authorship-related measures evolve in open-source communities. While they focused only on one huge system, i.e., the Linux kernel, we in contrast focused on several comparatively smaller open source systems.

To the best of our knowledge, no approach studied how emerging communities align with subsystems, nor how stable SDTs are over time compared to developer communities.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we investigate the emergence of latent developer interaction communities and how they align with subsystem developer teams (SDT). We observed that developer community membership is less stable than subsystem developer teams. We noticed hardly any correlation between detected communities and SDTs. One cause behind the observed low overlap between communities and SDTs could be the need to remain robust against high developer fluctuation in an open source development environment.

As a future work, we are interested in comparing in more detail those projects that experience different levels of developer fluctuation, respectively what other factors might bring about the lack of correlation between communities and SDTs.

ACKNOWLEDGMENT

This work was supported by the Austrian Science Fund (FWF): P31989 and P29415-NBL, and the LIT Secure & Correct Systems Lab funded by the state of Upper Austria.

REFERENCES

- [1] Eoin Ó Conchúir, Pär J. Ågerfalk, Helena Holmström Olsson, and Brian Fitzgerald. Global software development: where are the benefits? *Commun. ACM*, 52(8):127–131, 2009.
- [2] Gary Gaughan, Brian Fitzgerald, and Maha Shaikh. An examination of the use of open source software processes as a global software development solution for commercial software engineering. In *35th Euromicro Conference on Software Engineering and Advanced Applications, SEAA 2009, Patras, Greece, August 27-29, 2009, Proceedings*, pages 20–27. IEEE Computer Society, 2009.
- [3] Christian Bird, David Pattison, Raissa D’Souza, Vladimir Filkov, and Premkumar Devanbu. Latent social structure in open source projects. In *ACM SIGSOFT International Symposium on Foundations of software engineering, SIGSOFT ’08/FSE-16*, pages 24–35. ACM, 2008.
- [4] Nachiappan Nagappan, Brendan Murphy, and Victor R. Basili. The influence of organizational structure on software quality: an empirical case study. In Wilhelm Schäfer, Matthew B. Dwyer, and Volker Gruhn, editors, *30th International Conference on Software Engineering (ICSE 2008), Leipzig, Germany, May 10-18, 2008*, pages 521–530. ACM, 2008.
- [5] Andrew Meneely, Laurie Williams, Will Snipes, and Jason A. Osborne. Predicting failures with developer networks and social network analysis. In Mary Jean Harrold and Gail C. Murphy, editors, *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2008, Atlanta, Georgia, USA, November 9-14, 2008*, pages 13–23. ACM, 2008.
- [6] Martin Pinzger, Nachiappan Nagappan, and Brendan Murphy. Can developer-module networks predict failures? In Mary Jean Harrold and Gail C. Murphy, editors, *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2008, Atlanta, Georgia, USA, November 9-14, 2008*, pages 2–12. ACM, 2008.
- [7] Marcelo Cataldo and James D. Herbsleb. Coordination breakdowns and their impact on development productivity and software failures. *IEEE Trans. Software Eng.*, 39(3):343–360, 2013.
- [8] Luis López-Fernández, Gregorio Robles, Jesús M. González-Barahona, and Israel Herraiz. Applying social network analysis techniques to community-driven libre software projects. *IJITWE*, 1(3):27–48, 2006.
- [9] Andrejs Jermakovics, Alberto Sillitti, and Giancarlo Succi. Mining and visualizing developer networks from version control systems. In *Proceedings of the 4th International Workshop on Cooperative and Human Aspects of Software Engineering, CHASE ’11*, pages 24–31, New York, NY, USA, 2011. ACM.
- [10] Andrew Meneely and Laurie Williams. Socio-technical developer networks: should we trust our measurements? In Richard N. Taylor, Harald C. Gall, and Nenad Medvidovic, editors, *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu, HI, USA, May 21-28, 2011*, pages 281–290. ACM, 2011.
- [11] Sebastiano Panichella, Gabriele Bavota, Massimiliano Di Penta, Gerardo Canfora, and Giuliano Antoniol. How developers’ collaborations identified from different sources tell us about code changes. In *International Conference on Software Maintenance and Evolution, 2014*, pages 251–260.
- [12] M. Joblin, W. Mauerer, S. Apel, J. Siegmund, and D. Riehle. From developer networks to verified communities: A fine-grained approach. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, volume 1, pages 563–573, May 2015.
- [13] Valentina Lenarduzzi and Outi Sievi-Korte. On the negative impact of team independence in microservices software development. In *Proceedings of the 19th International Conference on Agile Software Development: Companion, XP ’18*, New York, NY, USA, 2018. Association for Computing Machinery.
- [14] M. M. Mahbul Syeed and Imed Hammouda. Socio-technical congruence in OSS projects: Exploring conway’s law in freebsd. In *Open Source Software: Quality Verification - 9th IFIP WG 2.13 International Conference, OSS 2013, Koper-Capodistria, Slovenia, June 25-28, 2013, Proceedings*, volume 404 of *IFIP Advances in Information and Communication Technology*, pages 109–126. Springer, 2013.
- [15] Usman Ashraf, Christoph Mayr-Dorn, Alexander Egyed, and Sebastiano Panichella. A mixed graph-relational dataset of socio-technical interactions in open source systems. In *MSR ’20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020*, pages 538–542. ACM, 2020.
- [16] Santiago Dueñas, Valerio Cosentino, Gregorio Robles, and Jesus M. Gonzalez-Barahona. Perceval: Software project data at your will. In *Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings, ICSE ’18*, pages 1–4, New York, NY, USA, 2018. ACM.
- [17] Forest Gregg and Derek Eder. Dedupe. <https://github.com/dedupeio/dedupe>.
- [18] Dedupe blog: Should i use dedupe.io or the dedupe python library? <https://dedupe.io/documentation/should-i-use-dedupeio-or-the-dedupe-python-library.html>.
- [19] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69(2), Feb 2004.
- [20] Andrea Lancichinetti, Filippo Radicchi, José J. Ramasco, and Santo Fortunato. Finding statistically significant communities in networks. *PLOS ONE*, 6(4):1–18, 04 2011.
- [21] Jure Leskovec, Kevin J. Lang, and Michael W. Mahoney. Empirical comparison of algorithms for network community detection. In Michael Rappa, Paul Jones, Juliana Freire, and Soumen Chakrabarti, editors, *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, pages 631–640. ACM, 2010.
- [22] Gerardo Canfora, Luigi Cerulo, Marta Cimitile, and Massimiliano Di Penta. Social interactions around cross-system bug fixings: the case of freebsd and opensbd. In Arie van Deursen, Tao Xie, and Thomas Zimmermann, editors, *International Working Conference on Mining Software Repositories, MSR 2011 (Co-located with ICSE)*, pages 143–152. ACM, 2011.
- [23] Claude E Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948.
- [24] Supporting online material. <https://doi.org/10.5281/zenodo.4591724>.
- [25] Christian Bird, Nachiappan Nagappan, Brendan Murphy, Harald C. Gall, and Premkumar T. Devanbu. Don’t touch my code!: examining the effects of ownership on software quality. In *SIGSOFT/FSE’11 19th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-19) and ESEC’11: 13th European Software Engineering Conference (ESEC-13), Szeged, Hungary, September 5-9, 2011*, pages 4–14, 2011.
- [26] D. A. A. Tamburri, F. Palomba, and R. Kazman. Exploring community smells in open-source: An automated approach. *IEEE Transactions on Software Engineering*, pages 1–1, 2019.
- [27] F. Palomba, D. A. A. Tamburri, F. Arcelli Fontana, R. Oliveto, A. Zaidman, and A. Serebrenik. Beyond technical aspects: How do community smells influence the intensity of code smells? *IEEE Transactions on Software Engineering*, pages 1–1, 2018.
- [28] Damian A. Tamburri, Fabio Palomba, Alexander Serebrenik, and Andy Zaidman. Discovering community patterns in open-source: a systematic approach and its evaluation. *Empir. Softw. Eng.*, 24(3):1369–1417, 2019.
- [29] William Leibzon. Social network of software development at github. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2016, San Francisco, CA, USA, August 18-21, 2016*, pages 1374–1376, 2016.
- [30] Armel Jacques Nzekon Nzeko’o, Matthieu Latapy, and Maurice Tchuente. Social network analysis of developers’ and users’ mailing lists of some free open source software. In *2015 IEEE International Congress on Big Data, New York City, NY, USA, June 27 - July 2, 2015*, pages 728–732, 2015.
- [31] Gerardo Canfora, Luigi Cerulo, Marta Cimitile, and Massimiliano Di Penta. Social interactions around cross-system bug fixings: The case of freebsd and opensbd. In *Proceedings of the 8th Working Conference on Mining Software Repositories, MSR ’11*, page 143–152, New York, NY, USA, 2011. Association for Computing Machinery.
- [32] Audris Mockus, Roy T Fielding, and James D Herbsleb. Two case studies of open source software development: Apache and mozilla. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(3):309–346, 2002.
- [33] Sebastiano Panichella, Gerardo Canfora, Massimiliano Di Penta, and Rocco Oliveto. How the evolution of emerging collaborations relates to code changes: an empirical study. In *International Conference on Program Comprehension, ICPC 2014*, pages 177–188, 2014.
- [34] Qiaona Hong, Sunghun Kim, Shing Chi Cheung, and Christian Bird. Understanding a developer social network and its evolution. In *2011 27th IEEE international conference on software maintenance (ICSM)*, pages 323–332. IEEE, 2011.
- [35] Guilherme Avelino, Leonardo Passos, Andre Hora, and Marco Tulio Valente. Assessing code authorship: The case of the linux kernel. In *Open Source Systems: Towards Robust Practices*, pages 151–163, Cham, 2017. Springer International Publishing.