

Community Structure, Individual Participation and the Social Construction of Merit

Matthias Studer

University of Geneva, department of econometrics, 40 Bd du Pont d'Arve,
1211 Geneva 4, Switzerland, matthias.studer@metri.unige.ch

Abstract. FLOSS communities are often described as meritocracies. We consider merit as a social construction that structures the community as a whole by allocating prestige to its participants on the basis of what they do. It implies a hierarchy of the different activities (web maintenance, writing code, bug report...) within the project. We present a study based on the merging of two datasets. We analyze the archive of KDE mailing lists using a social network. We also use responses to a questionnaire of KDE participants. Results bring empirical evidences showing that this hierarchy structures the community of KDE by allocating more central position to participants with more prestigious activities. We also show that this hierarchy structures individuals participation by giving greater "membership esteem" to members involved in more prestigious activities.

Key words: Collective Self-esteem, Community of Practice, Meritocracy, Open Source, Social Network Analysis, Social Structure.

Introduction

It is often proposed that the distinctive social structure of FLOSS communities could be one of the key reasons of its success [1]. This organisation is often described as meritocratic [2] or at least, as having “an ideological commitment to meritocracy” [3]. In this article, we propose to discuss the concept of meritocracy and to describe how it structures the community. We intend to bring empirical

Please use the following format when citing this chapter:

Studer, M., 2007, in IFIP International Federation for Information Processing, Volume 234, Open Source Development, Adoption and Innovation, eds. J. Feller, Fitzgerald, B., Scacchi, W., Sillitti, A., (Boston: Springer), pp. 161–172.

evidences to support our matter by taking KDE¹ as a case study. Our study is based on two data sources: e-mails archives and a questionnaire of KDE participants.

Meritocracies are social systems in which the social position is determined by merit. In other words, the social position is determined by the social valorisation of the activities done in the community. Thus, some activities lead to more influential position than others according to the merit linked to each activity. This distribution of the merit between different activities structures the community as a whole. According to our point of view, the definition of the merit is not objective, but results from a social construction that we need to understand better. Indeed, it seems to form the base of power relationships in FLOSS communities.

According to the theory of communities of practice [4], this allocation of power is described using the concept of mutual engagement. What a person does allows to “categorize him as”, i.e. to assign him to a peculiar social position. This classification (which can be either positive or negative) is carried out by the other members of the community in an unconscious way through the returns (comments) made on each contribution. But what we do also makes it possible to “categorize us ourselves as” through the returns carried out by the other members of the community. Thus, the process of “power allocation” is done at the community level by the allocation of power to people who deserve it. It is also an individual process since each one “categorize himself as” compared to other members of the community.

The socially constructed definition of merit structures community on two levels. Firstly, it structures the community as a whole by allocating more influential positions to members carrying out more valorised activities. Secondly, it structures individual participation by giving a conscience of oneself position in the community. Thus, the centre of our analysis is the “person-in-the-social-world” [5], i.e. we seek to clarify the relationships between social norms and individual participation. According to the theory of communities of practice [4], we should look at the relationships between these two levels in order to better understand the social structure. The social structure acts through its internalization by contributors.

Following these two points of view, we selected two data sources. The first one is the archive of all e-mails archived by MARC since the beginning of KDE. This data source will enable us to approach the social structure as a whole through a social network analysis. The second data source comes from a questionnaire of KDE participants. It will enable us to approach the meaning given to each activity and some more subjective elements of our assumptions like membership esteem of contributors.

This article is organized in the following way. We present the data sources and the methods we used to extract relevant information from it. Then, we will analyze the hierarchical organization of activities. After that, we will look at the internalisation done by KDE contributors before concluding our article.

¹ See <http://www.kde.org>

Presentation of data sources

As we said, we have two data sources: e-mails archive and a questionnaire of KDE contributors. We used e-mails sent to KDE mailing-lists and archived by MARC². We used the data from beginning of January 2006 until end of June 2006 covering a six month period. These e-mails come from the lists of discussions within each project and sub-project. We also used information from the “kde-bug-dists” mailing-list, a list gathering automatic notifications for all changes made to Bugzilla. The use of e-mails archive enables us to bring all these sources together on the basis of names and e-mail addresses. We used this information to build a social network analysis of all participants to KDE mailing-lists using Pajek [6]. In this kind of analyses, two definitions are essential: inclusion and relationship.

Regarding inclusion, two problems quickly arise: neither the e-mails addresses nor the names can be considered unique. Consequently, we used an in-depth search algorithm to put together the couples of “name-email”. This algorithm was used to propose possible merges to a human. Since all regroupings were human-supervised, we were forced to use a selection criterion. We thus regrouped and selected all person having sent at least ten messages over a period exceeding one month. One can argue that we introduced a systematic bias in our analysis by using this selection criterion. However, we think that this bias will not be very influential. We are interested in relationships with people who are important for the community. It is not abusive to think that these peoples sent at least ten e-mails.

One usually defines the relation using “point-to-point” information such as the “In-Reply-To” header of emails. However, this information was not available in KDE archives. Thus, we used the definition of “thread” from MARC to constitute our network. We have then defined the relationship between two persons as:

The relationship between a person A and B is equal to the sum of all messages sent by A in “threads” where B also sent at least one message.

The relationship has a direction (from A to B) and the value is different according to this direction. However, and this rises from the given definition, if A has a relationship with B, then B has a relationship with A. It will not be automatically the same value. The graph obtained is directed and valued. Our measurement also contains a scale about the “force” of the relationship. Thus, taking part in a discussion with a lot of different participants implies more “relations” than taking part in a small discussion. In other words, each message is not equal in our construction of the network. This corresponds to some logic. By taking part in a large discussion (which has more chance to be considered as important), one acquires a greater visibility than in a small discussion implying only two people.

Consequently, how to interpret the values of relationships? We suggest thinking in term of contacts. One “unit of relation” corresponds to one contact initiated if one thinks in terms of outgoing relationships, or with one contact received in the case of incoming relationship. The construction of our network makes difficult to compare values of incoming and outgoing relationships since the same message does not have

² Mailing list ARChives (<http://lists.kde.org>)

the same weight according to the direction: a message received corresponds to a contact whereas a message sent can correspond to several contacts. Being given that most messages are “functional”, we think that we should interpret our network as some kind of cooperation network. In our sense, it would be abusive to speak of friendly networks. In fact, the use of IRC discussions archive would be probably closer to such an interpretation of the network.

We computed several indicators from the social network analysis in order to test our hypothesis, namely the degree, the sum of incoming lines, the sum of outgoing lines and the maximum value of all arcs. The degree is simply the number of arcs connected to a given vertex (contributor) in our social network. According to our definition of the network, the sum of incoming lines corresponds to the number of messages received. The sum of outgoing lines can be interpreted as an indicator of influence in our network. Finally, the maximum values of all arcs (incoming or outgoing) should reflect the strength of the connection with other members.

All these indicators are local. So, we also computed “proximity prestige” [8]. This indicator is equal to the number of vertex that a given vertex can reach according to the arcs present in the network divided by the mean length of path to reach all these vertices. It is an indicator of the position in the global network.

We also used a questionnaire of KDE participants. The answers were collected online. Thus, the answerers were voluntary. KDE contributors were informed about the existence of the questionnaire through several messages sent on KDE mailing-lists. We took care to include all kind of mailing-lists such as users-oriented, translators and developers mailing-lists. We also took care to address our questionnaire to all kind of KDE contributors (including bug-reporter, translators, developers, etc.). However, we did not send the advertisement to all mailing-lists – we did not want to be considered as a “spammer” by KDE contributors.

We collected 131 answers. This low number of answers leads us to think that generalization of our results should be done with caution. Because we are here within a framework of observation and not of experimentation, the variations could be caused by factors for which we do not control. Answering the questionnaire was not especially long (approximately 15 minutes). However, 25 peoples did not answer the whole of the ten pages presented. The first question made it possible to establish the link between the questionnaire and the e-mails archives. This question was optional since some answerers may prefer to remain anonymous on Internet. Only 95 people gave an answer to this question.

The questionnaire was made up of questions about involvement in KDE, collective identity, demographic information and individual motivations. In this article, we will use the answers to three questions which were related to:

- The activities carried out within the community on a 6 item scale of frequency ranging from “Never or almost never” to “Every day or almost every day”.
- The prestige granted to each activity on a scale ranging from 1 (No prestige) to 10 (very prestigious).
- Two questions related to the importance that one think one have in a given social group. These questions were “I am a worthy member of my KDE group(s)” (positive formulation) and “I feel I don't have much to offer to my KDE group(s)” (negative formulation) on 7 item scale ranging from “Strongly disagree” to “Strongly agree”. These questions were taken from the “membership

esteem” subscale of the collective self-esteem scale proposed by Luhtanen and Crocker [7]. We transformed the formulation of the scale by replacing references to “social groups” with “KDE groups”. The “membership esteem” subscale is equal to 8 plus answer to the first question minus answer to the second one.

We can use the data from the social network analysis to get an idea of the representativeness of our answerers. Unsurprisingly, our population is not representative but far more involved in the community. The mean degree of all network members is 10.49 against 54.64 for our answerers. This difference is statically significant and independent of the network indicator used. To get a better idea of the involvement of our answerers, we present on the figure below the maximum frequency between all activities carried out in KDE. As we can see, almost 75% of our respondents contribute at least “once or twice a week”. Hence, most of our respondents can be considered as frequent contributors.

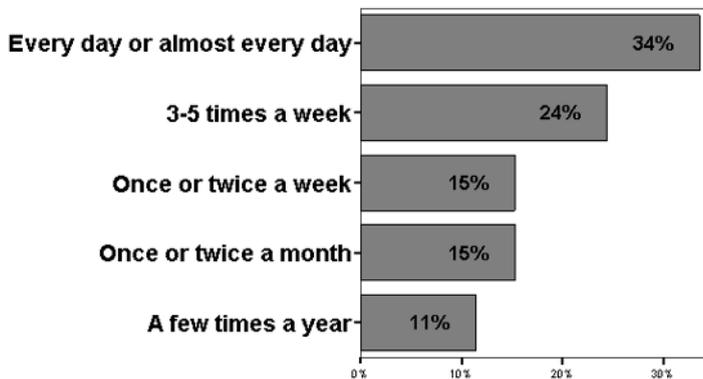


Fig. 1. Maximum frequency of involvement of answerers

In this article, we are interested in the social structure of the community. According to our point of view, the social structure comprises frequent as well as infrequent contributors. Thus, in following analysis, we included all answerers³.

After having presented our data source, we will look at the social valorisation of activities and how these valorisation' schemes structure the community.

Hierarchical classification of activities

We argued that the social construction of merit organize the activities into a hierarchy. We propose to describe this organization before showing how it organizes the community as a whole. After that, we'll show how this social construction structures the individual participation of contributors.

³ However, interpretations (and main results) do not change if we include only respondents contributing at least “once or twice a week”.

The table below presents the mean prestige accorded to each activity by the answerers on a one (no prestige) to ten (very prestigious) scale. Seven answerers refused to answer this question by giving ten to all activities. The classification does not fundamentally change if we use a transformation such as rank. Changes in the hierarchy may appear if prestige averages are close⁴. The third column shows a rank for each activity. A difference of rank indicates a statically significant difference in the prestige distributions⁵. On the contrary, if the ranks are the same, we can not assume that one of the distributions is significantly higher.

Table 1. Hierarchical classification of the activities according to prestige scale

Activity	Average Prestige	Rank
Code	8.46	1
Coordination	7.44	2
Discussion about future development	6.98	3
Art	6.11	4
Bug Management	5.43	5
Help	5.42	5
Documentation	5.19	5
Translation	5.10	5
Packaging	5.03	5
Web	4.55	6
Bug Reports	4.29	6

According to our hypothesis, this ranking of activities is a social construction. It does not mean that some activities are less valuable or less important than others, but that the social interpretation and meaning given to these activities are different. This construction is not arbitrary but corresponds to some logic that we seek to understand.

The activities turned towards technology seem to be the most prestigious. Thus, the writing of code occupies the first position and the discussion of future developments the third. The presence of the activity “Art” in fourth position shows us that the creative and productive activities are largely valorised. Finally, let us note that the coordination of the community is seen as a prestigious activity. This is not surprising since our societies (as a general rule) tend to valorise this kind of activities.

The activities which come after are more difficult to distinguish, because the averages are close and the ranks are the same. It shows us that we can not assume a clear hierarchy between these activities. In this group, we find other form of contribution to the community such as bug management, user assistance, documentation, translation and packaging. It will be noticed that the user assistance appears relatively valorised for an activity which is not productive (or whose result is

⁴ We always check all possible transformation in the reasoning presented below.

⁵ We computed the Wilcoxon signed ranks test for all pair of activity scales. We then set the activity ranks according to statically significant (at the 5% level) difference in scale’s distribution.

not visible in the final product). It seems to occupy a similar position to much more productive activities such as translation or documentation.

Finally, we find a last group with much less valued activities with an average below five which would correspond to the mathematical average: website maintenance and bug reports. In our sense, the rank of website maintenance is quite surprising since it's a productive activity quite important for the community. Finally bug reporting occupies the lowest position in the hierarchy.

Translation is not classified in a uniform way between the answerers. Translators (N=54) tend to classify this activity significantly higher (Fisher's test: $F=10.59$; $df=1$; $p=0.002$) by giving an average of prestige of 5.85 whereas the "non-translators" (N=55) gives an average of 4.38. The translation passes thus from the fourth place (for translators) to the bottom of the hierarchy (for other contributors). This difference is even more significant if one takes into account the people making translation at least "one or twice a week" (N=32). The average of prestige is then 6.5 for translators against 4.52 for the others. The differences in distributions (Wilcoxon test) are statistically significant ($Z=-3.291$; $p=0.001$) in the first case and the second ($Z=-3.841$; $p<0.0001$).

This difference is important. It means that non-translators will evaluate less prestigious "translation" than translators. Thus, translators will gain much more prestige inside their own sub-community than in the community as whole. It shows us that translators seem to form a sub-community with their own definition of merit. However, this definition is not totally different. The top of the hierarchy is not disputed. Translation does not imply what translators consider as the most prestigious activities such as coding. Thus, translators are in a dynamic where their own activity (and their sub-communities) remains necessarily peripheral. We can also give a second interpretation of this difference. There is a social necessity for translators to valorise their own activities in order to maintain a regular involvement. Indeed, translation seems to be less valued by the community as whole.

The activities performed inside KDE are not evaluated in a uniform way. Some activities are more prestigious than others. Therefore, there is a social construction of merit. In this social construction, technical and creative activities seem to be the most prestigious alongside with coordination. After having presented the valorisation's scheme of the activities, we will show how these differences of prestige structure the community as whole.

Activity Prestige and Community structure

According to our hypothesis, the social construction of merit should structure the community as a whole. So, we computed the correlation between the frequency of each activity and indicators computed from the network analysis. The correlations presented in the table below are Kendall's τ -b correlations⁶ and are all significant at 1% level. In the table below, non-significant correlations are ignored.

⁶ We always computed the Kendall's τ -b correlation rather than Pearson or Spearman correlation. Kendall's τ -b correlation is known to better handle ties values (which are typically frequent with ordinal measure) and make no assumption on the distribution of variable or on the form of the relationship between both variable (such as linear

Table 2. Correlation (Kendall’s τ -b) between network indices and frequencies of activities

	Degree	Sum of incoming lines	Sum of outgoing lines	Maximum values of all arcs	Proximity prestige
Help	0.228	0.233	0.217	0.223	0.207
Code	0.400	0.421	0.416	0.428	0.377
Discussion	0.436	0.473	0.456	0.484	0.422
Translation					
Doc.	0.238	0.237	0.224	0.227	0.231
Art					
Web	0.337	0.332	0.324	0.304	0.318
Coordination	0.255	0.275	0.262	0.285	0.245
Bug Management	0.404	0.414	0.414	0.439	0.390
Bug reports	0.244	0.220	0.249	0.226	0.261
Packaging					

We expect all correlations to be significant and positive. Whatever we do, if we do it more frequently, we should have more relationships in the network and our relationships should be stronger⁷. By looking at the table below, we can distinguish three groups of activities:

- Code, discussion and bug management show strong correlations. We should notice that the first two correspond to the most prestigious activities.
- Coordination, Help, Documentation, web and bug reports show correlations around 0.25. This set of activities is in the middle of our hierarchy. The only exception is coordination but we had only few answerers who stated doing it.
- Art and Packaging activities show no significant correlation. But this is mostly because of the small number of answerers who stated doing it. This is not the case of translation which does not show any significant correlation.

The correlations seem, generally speaking, to follow the prestige accorded to each activity. However, “bug management” and “web” show much stronger correlations than expected. One possible explanation is that these activities come alongside with influential position. In other words, it is possible that influential positions imply responsibilities and activities that are not necessarily prestigious or “fun”. In these cases, the influential positions would not be the result of such activities, but from the others performed alongside. Hence, “bug management” is highly correlated to “code” (τ -b = 0.47) and “web” is mostly performed by long time contributors.

relationship for instance). This coefficient is known to be more conservative than the other and values are typically lower than for Pearson correlation. See Arndt et al. [9] for a full discussion.

⁷ Some network indicators show stronger correlations with the frequency of each activity than others. The “sum of incoming lines” and the “maximum values of all arcs” show the strongest correlations. These indicators take into account the values of the arcs. Hence, it’s not only the number of relations (degree) but also the values of the arcs that are important. Correlation between activities and proximity prestige are weaker than with other indicators.

The absence of correlation between network indicators and “translation” is interesting. Translation had a peculiar position in the hierarchy of the activities: non-translators were evaluating this activity as less prestigious than translators. This absence of correlation means that performing more often translation does not lead to more connections in the community. In other words, translation seems to be a peripheral activity. One can argue that social power is poorly linked with social network indicators. In our sense, we should interpret these indicators as a necessary condition. It is necessary to have different and strong link in order to exercise some sort of social power.

The socially constructed definition of merit seems to structure the community by allocating more central position to people performing more prestigious activities. According to our hypothesis, this construction of merit should also be internalized by contributors. This is what we will try to show now.

Activity Prestige and Individual Participation

For recall, membership esteem refers to the importance that one thinks one has within social groups to which one is identified. Membership esteem is strongly correlated with the indicators from the network (τ -b = 0.38 with the sum of incoming line; $p < 0.0001$). Hence, members more strongly involved have more chance to feel important for the community. It is not necessarily the number of relations which counts more, but also the force of these bonds. Thus, one observes a positive correlation with the value of the strongest relation (τ -b=0.40; $p < 0.0001$).

Table 3. Correlation (Kendall's t-b) between activities and membership esteem

	Membership esteem	Positive formulation	Negative formulation
Help	0.18*		-0.19 *
Code	0.38 **	0.27 **	-0.38 **
Discussion	0.39 **	0.35 **	-0.31 **
Translation			
Doc.	0.21 *		-0.19 *
Art			
Web	0.19 *		-0.23 **
Coordination	0.22 **		-0.26 **
Bug Management	0.23 **	0.17 *	-0.23 **
Bug reports			
Packaging			

** Significant at 1% level; * significant at 5% level.

Membership esteem is not only linked to network indicators but also with the activities done by contributors. Indeed, if we look at correlations on the table above, one can identify three groups of activities:

- The activities “codes” and “discussions” show the most important correlations with a value around 0.4.
- The other activities come then with correlations of about 0.2. These correlations are mainly the fact of the rejection of the negative formulation rather than of strongest acceptance of the positive formulation. Thus, we should conclude that

this is the expected correlation between frequencies and membership esteem. The only exception is “coordination”. Indeed, we had only few answerers that state doing it.

- Finally, “translation” and “bug report” do not show any significant correlations; the same applies to “packaging” and “art”. But these last ones do not have enough answerers to enable us to deduce something from it.

The first group of correlation shows us that the most prestigious activity comes together with stronger membership esteem. The second group of activities shows the expected relationships between activity and membership esteem. The presence of the last group is interesting. According to our assumptions, we should expect positive correlations. More one makes, in a given activity, better the membership esteem should be. This absence of correlation indicates us that this hypothesis is not verified regarding “translation” and “bug report”. In other words, doing these activities more frequently does not lead to greater membership esteem. These activities were also less valued in the global classification. Therefore, we can conclude that membership esteem is linked to valorisation’s scheme of the activity.

We should clarify some points. Our results do not show that “translation” or other activities are devaluated within FLOSS communities. In fact, membership esteem was usually high and we should remember that the frequency of interaction (maximum frequency found between all activities) is highly correlated with membership esteem ($\tau\text{-}b=0.39$). Therefore, all activities lead to greater membership esteem. Hence, our conclusion is that doing less prestigious activities contributes less to membership esteem than very prestigious activities.

We showed that the social construction of merit structures individual participation. Doing more prestigious activities contributes more to membership esteem than doing less valued activities. The social construction of merit does not only structure the community as a whole. It is also internalised by contributors.

Conclusion

By using two distinctive sources of information, namely a social network analysis and a questionnaire of KDE contributors, we brought a new insight of FLOSS communities’ structuration since we were able to locate answerers inside the social structure. These two sources allowed us to think on two levels: the individual participation and the community. It also enabled us to think the relationship between these two levels of analysis which are usually measured separately.

We showed that there is a social construction of merit that implies a hierarchy of the activities performed in the community. This construction valorises the activities turned toward technological development such as coding and discussion about future development. Creative and coordination activities are also valorised. This does not mean, by any way, that some activities are less valuable or less time-consuming. We presented the social interpretation of activities not an evaluation.

We showed that the whole community is structured according to the activities performed and their social valorisations. The most prestigious activities seem to lead to more central position in the social network whereas we did not find such link for

less prestigious activities. We noticed that some less prestigious activities (bug management and web maintenance) seem to come alongside influential position. We showed that different activities relate to different social positions as measured with social network indicators. Therefore, we can conclude that the social power also comes from the activity performed and not only from the frequency of interaction with the community. In other words, aside from frequency, the kind of activity performed is also a key dimension of social position.

This structuration dynamics is not only observable at the community level. We showed that it seems to be internalized by contributors. Esteem of its own importance within the community is linked with the kind of activity performed. Specifically, some activities seem to be more linked with membership esteem than others. The classification of the activities is internalized and not only an external factor of individual participation.

We showed that the social construction of merit structures the community as whole as well as individual participation of contributors. The KDE community can be described as meritocratic. However, we did not explain the process in detail. A more in depth or ethnological analysis is needed in order to precisely describe the social construction of merit. More precisely, we need to understand how the different contributions (within a given activity) are evaluated. Our analysis showed us that the social construction of merit defines which activities are linked with more influential social positions.

Beyond the structuration of KDE according to the social definition of merit, our analysis showed us that individual participation to a FLOSS project should be understood in relation with the social structure of the community. Contributors internalise the social structure of the community and the social structure influences their own participation. From a theoretical and methodological perspective, we should think the relationships between individual participation and social structure. Activities done inside a FLOSS community are not individualistic but a form of participation.

References

1. K. Crowston and J. Howison, The social structure of free and open source software development, *First Monday* **10**(2), (2005).
2. E. S. Raymond, Homesteading the noosphere, *First Monday* **3**(10), (1998).
3. J. Howison, K. Inoue and K. Crowston, Social dynamics of free and open source team communications, in *proceedings* The Second International Conference on Open Source Systems (2006).
4. E. Wenger, *Communities of Practices: Learning, Meaning, and Identity* (Cambridge University Press, Cambridge, 1998).
5. J. Lave and E. Wenger, *Situated learning: legitimate peripheral participation* (Cambridge University Press, Cambridge, 1991).
6. V. Batagelj and A. Mrvar, Pajek 1.14 – Program for Large Network Analysis' <http://vlado.fmf.uni-lj.si/pub/networks/pajek/>
7. R. Luhtanen and J. Crocker, A Collective Self-Esteem Scale: Self-Evaluation of One's Social Identity, *Personality and Social Psychology Bulletin* **18**(3), 302-318 (1992).

8. W. de Nooy, V. Batagelj and A. Mrvar, *Exploratory Social Network Analysis with Pajek* (Cambridge University Press, Cambridge, 2005).
9. S. Arndt, C. Turvey and N. C. Andreasen, Correlating and predicting psychiatric symptom ratings: Spearmans r versus Kendalls tau correlation, *Journal of psychiatric research* **33**(2), 97-104 (1999).