

The Secret Life of Software Communities: What we know and What we Don't know

Gemma Catolino¹, Fabio Palomba², Damian A. Tamburri³
Delft University of Technology¹, University of Salerno²,
Eindhoven University of Technology & Jheronimus Academy of Data Science (JADS)³
g.catolino@tudelft.nl, fpalomba@unisa.it, d.a.tamburri@tue.nl

Abstract—Communities of software practice are increasingly playing a central role in the development, operation, maintenance, and evolution of good-quality software, as well as DevOps pipelines, lean Organizations and Global Software Development. However, it is still unknown the structures and characteristics behind such communities. For this reason, in this paper, we tried to explore the organizational secret of communities, trying to offer a few practical extracts of (1) what we know and is known, (2) what we know to be unknown, and (3) what we know to be tentatively discoverable in the near future from an empirical research point of view. Moreover, the paper provides a number of recommendations for practitioners to help and be helped in their community endeavors.

Index Terms—Community Types; DevOps; Road Map.

I. INTRODUCTION

Communities are the foundation and backbone of organized societies and just as much as our software-driven civil societies are dense with communities, the software makers are also organized into communities of practice (*e.g.*, in open-source), of intent (*e.g.*, in standardization groups), of purpose (*e.g.*, as part of agile movements). Recent studies showed how the community's health can reflect the quality of the software produced [5], [6], for this reason research community tried to deeper analyze and characterize the structure of software communities (*i.e.*, defined as social units of size, dense strongly-typed and diverse interactions across diverse roles and fluid characteristics) [12].

In this paper, we discuss about a few basic facts and observations about the secret life of software communities which we were able to empirically establish over previous research studies as well as our own practice as part of software communities themselves.

We outline the roundup to encourage further research into understanding and harnessing the (still) very much secret life of open-source communities, as well as their structural, socio-technical, and health characteristics. Furthermore, we aim to call for help from the practitioners themselves in supporting their own software community activity and therefore, we provide practitioners with practical recommendations or calls for help defined with the intent of engaging their interest around the matter and/or allowing further understanding their needs from a research perspective, such that we as a research community can aid them in a better fashion. Moreover, we provide a brief road map for further research along this intriguing emerging topic.

II. PRACTICAL RECOMMENDATION

A. *FACT 1. Software Communities Have Types*

Several community types have been studied in organizational structures and social networks research - Figure 1 offers an overview of these types extracted from a systematic studies on the matter [11], [12]. In Layman's terms, a community type is a series of characteristics, a pattern, which remain constantly true across the social network underlying that community. What is still understood fairly little in the state of the art of software engineering research is that, on the one hand, software communities not surprisingly exhibit the same types already known in literature, but, on the other hand, the role of community types and characteristics for the benefit (or fallacy) of software code qualities as well as software processes remains mostly unknown.

What we don't know: (a) the influence of community types over software qualities as well as the qualities of software processes; (b) algorithms and measurements to precisely pinpoint type shifts across the community's life-cycle; (c) design patterns for community structures which are contiguous with design patterns in underlying software architectures.

What should practitioners do: (1) follow community tracking and measurement initiatives such as OPENHUB¹ or BITERGIA² to gather insights over their own community and act upon it - this ensures that progress in community management, measurement, and steering practices are harnessed; (2) engage in community quality and health initiatives such as CHAOSS³ - this ensures that initiatives struggling to crack the code of sustainable software communities are well fed with engaged practitioners; (3) manifest their own organizational and socio-technical issues as much as the code-quality issues - this ensures that social software engineering [13] researchers have both quantities and qualities of the right data and evidence to study.

B. *FACT 2. Types Influence Software Qualities*

We conducted several experiments to exploratively figure out the boundaries around the influence that software community types may be playing over the qualities for software production and its maintenance. Figure 2, for example, shows the influence of several types over the rate of re-opened

¹ <http://openhub.net/> ² <https://bitergia.com/> ³ <http://chaoss.community/>

Name	CoP	IN	FN	IC - Informal	NoP	WG	PT	FG	PSC	Problem-	LC	SN	KC	SC - Strategic
Features	Community of Practice	Informal Network	Formal Network	Community	Network of Practice	Working Group	Project Team	Formal Group	Solving Community	Learning Community	Social Network Site	Knowledge Community	Community	Community
Community Structure	high	high	high	high	high	high	high	high	high	high	high	high	high	high
Situatedness	high	low	low	low	low	high	high	high	low	low	low	low	low	low
Dispersion	low	high	high	high	high	low	low	low	high	low	low	high	high	high
Problem-Solving	low	low	low	low	low	high	high	high	high	low	low	low	low	high
Informality	high	high	low	low	high	high	high	high	high	high	low	high	high	low
Formality	low	low	high	low	low	high	high	high	low	low	low	low	high	high
Engagement	high	high	low	high	low	high	high	low	high	high	low	low	low	low
Cohesion	low	low	high	low	low	high	high	low	low	low	low	low	low	low
Duration	high	high	high	high	high	high	low	high	high	high	low	low	low	high
ROI-Tracking	low	low	low	low	low	low	high	low	low	low	low	low	low	high
Governance	low	low	high	low	low	high	high	high	low	low	low	low	low	high
Culture	low	low	low	low	low	high	low	low	high	high	low	low	low	low
Visibility	low	low	low	low	low	low	low	low	low	low	low	high	high	low

Fig. 1. An Overview of Community Types from the State of the Arts [12]

issues investigated over 25 open-source communities sampled according to community participants age, size, gender diversity, community programming language, type of product, geographical dispersion, and activity, the last one measured as the mean activity stemming from BITERGIA and OPENHUB. The Box-plot in Figure 2 was generated by operationalizing




Fig. 2. The influence of community types over re-opened software issues

a series of metrics to measure essential characteristics from Figure 1 and match them with corresponding types as well as their mean proneness to re-opening issues, that is, the mean ratio of re-opened issues over 6 months worth of releases in the projects lifetime. The plot shows that something is in fact going on - in this case, the data shows that Formal-Networks (second bar from the left) may be exhibiting an organizational behavior which tends to re-open issues more often than other types. Conversely, the opposite seems true for Informal-Communities (third bar from the left).

What we don't know: (a) a general quality model of what software qualities are influenced by which community structure qualities; (b) which community structure characteristics play a role for software communities; (c) how to quantify what is the additional cost or socio-technical debt connected to sub-optimal characteristics; (d) how to measure a type and use those measurements as devices for improved governance or organizational structure agility.

What should practitioners do: (1) open issues on issue-trackers not only for software code but also for perceived community structure issues - they are as much impactful as they are dangerous and can even lead to breaking the internet (see the NPM incident⁴); (2) report software community accidents on your version control system as much as you do on STACKOVERFLOW - researchers need to study both to come up with proper empirically established ground truths as well as practical outputs to support your work.

C. FACT 3. Types narrow with Practitioners' Experience

Similarly to results in Figure 2, we report that the number of types intermixed in the same community types, that is, the number of attributes diversity across a community reduces as the practitioners' age (or their experience, skills, community participation). In the same study that led us to prepare the plot in Figure 2, we reported an inverse correlation of ~ -0.40 (p -value $\ll 0.05$) between the mean developers experience with the community (i.e., the total time they spent working in the community) across our projects sample and the amount of community types and characteristics that manifest explicitly across the community.

What we don't know: (a) how to quantify mentorship and experience as assets across software communities and how to reward both in a proper fashion; (b) how to infer community

⁴ <https://tinyurl.com/yafertj3b>

evangelists and use them as thought leaders for their software communities; (c) what skills and characteristics matter more than others in software communities and which skills grow with age and which others do not, as well as how to foster the growth and nourishment of skills that cannot grow with time (e.g., knowledge communication)

What should practitioners do: (1) prepare codes of conducts [1] which reflect also the mentorship, governance structure, experience and reward management as well as the code contribution policies or good-behavior across the community; (2) make the software code reflect the code of conduct, namely, use comments in software code to reprimand or exercise social sanctioning [9], [11] wherefore codes of conducts are violated.

D. FACT 4. Communities with more formal characteristics may result more often into abandon-ware

We observed that 82% of the abandonware communities in our sample exhibited a formal type in the last 6 months of their life. This is in line with previous research and observations from Crowston *et al.* [14] where informality is signaled as an essential software community health parameter.

What we don't know: (1) whether it is in fact true if formal characteristics and types lead to abandon ware — empirical software engineering needs to be instrumented to actually assess and establish this link; (2) what 'formal' means, in the software engineering world and in terms of socio-technical and organizational relations — typically in software engineering a formal formulation involves proofs, foundational theories, but from the realm of social-networks analysis and organizations research, the word and concept of "formality" assumes a sensibly different meaning; (3) the measured role of "informal" as opposed to formal — how can one foster informal? Is Informal beneficial as formal is not? These and similar research questions are still out there for the taking;

What should practitioners do: we don't know, yet.

E. FACT 5. Communities can be healthy and sustainable

Many studies in the literature identify healthy values for several community characteristics, *e.g.*, socio-technical congruence [2] or community structure verification [3]. These indicators, stemming from top studies, lead to argue that software communities, like other thriving communities in our societal structures, can be turned healthy and made sustainable with appropriate and dedicated efforts — sites such as BITERGIA, OPENHUB, as well as initiatives and research projects such as OSSMETER⁵ are fundamental assets to drive the societal challenge of making software communities aware and sensible to their health, as sustainable communities should be. In this respect, the state of the art in organizations research, social networks analysis as well as emerging disciplines such as sustainable community development [4], [14] can aid but even typical, structured approaches used in software engineering such as formalization [7], *e.g.*, to better understand and measure the socio-organizational dynamics playing a role across software communities. And for those skeptical enough

⁵ <http://ossmeter.org/>

to think that this endeavor is pointless or less important than, say, software testing or maintenance, please read your literature, *e.g.*, a couple of hints are in Smite *et al.* [10] or [8].

What we don't know: (a) a systematically-generated, general quality model for software communities with supporting evaluation and analytic tools and automatons; (b) a practical implementation of sustainability in software communities which are currently managed in a rather trial-and-error fashion, with due exceptions (*e.g.*, the Apache Software Foundation); (c) the dimensions, qualities, and metrics of software community sustainability, to be used jointly with a quality model for performing, long-lived, high-quality software communities.

What should practitioners do: we don't know, yet.

III. CONCLUSIONS

As it turns out, there is still information that need to be further investigate . As a matter of fact, our non-exhaustive list cannot convey enough that there is much more we don't know with respect to how much we do know; most especially, we don't know a lot about what should practitioners do to cater for their communities, striving for healthy types matching their organizational requirements over time and in a sustainable fashion. Our conclusion is not only that more work is needed to explore the aspects we non-exhaustively pointed at — rather, the most dire observation is that, in order to better support the (secret) life of software practitioners in their communities, practitioners themselves may be well encouraged to treat the community as a software-influencing artifact itself, thus, for example, opening issues if a community issue or smell does in fact manifest. An increased practitioner community consciousness will increase awareness over the problem, making it more explicit, measurable, and hence improvable by practitioners as well as for researchers. In the future it is our ultimate intention to further pursue the road map that emerges from the previous identified shortcomings, and, at the same time, work to improve software forges, collaborative coding environments, IDEs or other software equipment that practitioners may use to build, maintain, or work upon their code as part of a lively and healthy community and with more appropriate software engineering 'ergonomics', intended as the disciplines of software engineering which engage into designing or arranging software commons, communities, workplaces, work-products, and working systems so that they fit the people and goals around them [14].

REFERENCES

- [1] C. M. ACM. Acm code of ethics and professional conduct. *Code of Ethics*, 1992.
- [2] M. Cataldo, J. D. Herbsleb, and K. M. Carley. Socio-technical congruence: a framework for assessing the impact of technical and work dependencies on software development productivity. In *Proceedings of the Second ACM-IEEE international symposium on Empirical software engineering and measurement*, pages 2–11. ACM, 2008.
- [3] M. Joblin, W. Mauerer, S. Apel, J. Siegmund, and D. Riehle. From developer networks to verified communities: a fine-grained approach. In *Proceedings of the 37th International Conference on Software Engineering-Volume 1*, pages 563–573. IEEE Press, 2015.

- [4] J. Lucena, J. Schneider, and J. A. Leydens. Engineering and sustainable community development. *Synthesis Lectures on Engineers, Technology, and Society*, 5(1):1–230, 2010.
- [5] F. Palomba, D. A. Tamburri, A. Serebrenik, A. Zaidman, F. A. Fontana, and R. Oliveto. How do community smells influence code smells? In *ICSE (Companion Volume)*, pages 240–241. ACM, 2018.
- [6] F. Palomba, D. A. A. Tamburri, F. A. 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*, 2018.
- [7] D. Rombach and F. Seelisch. Formalisms in software engineering: Myths versus empirical facts. In *IFIP Central and East European Conference on Software Engineering Techniques*, pages 13–25. Springer, 2007.
- [8] J. Rost and R. L. Glass. *The dark side of software engineering: evil on computing projects*. John Wiley & Sons, 2011.
- [9] K. Sigmund, C. Hauert, A. Traulsen, and H. De Silva. Social control and the social contract: the emergence of sanctioning systems for collective action. *Dynamic Games and Applications*, 1(1):149–171, 2011.
- [10] D. Šmite and Z. Galviņa. Socio-technical congruence sabotaged by a hidden onshore outsourcing relationship: lessons learned from an empirical study. In *International Conference on Product Focused Software Process Improvement*, pages 190–202. Springer, 2012.
- [11] D. A. Tamburri, P. Lago, and H. Van Vliet. Uncovering latent social communities in software development. *IEEE software*, 30(1):29–36, 2012.
- [12] D. A. Tamburri, P. Lago, and H. v. Vliet. Organizational social structures for software engineering. *ACM Computing Surveys (CSUR)*, 46(1):3, 2013.
- [13] W. Tracz. Lord of the files: essays on the social aspects of software engineering by russel ovans. *ACM SIGSOFT Software Engineering Notes*, 36(6):31–31, 2011.
- [14] R. T. Watson, M.-C. Boudreau, and A. J. Chen. Information systems and environmentally sustainable development: Energy informatics and new directions for the is community. *MIS quarterly*, 34(1), 2010.