# *"The Canary in the Coal Mine..."* A Cautionary Tale from the Decline of SourceForge

Damian Andrew Tamburri<sup>\*1</sup> | Kelly Blincoe<sup>2</sup> | Fabio Palomba<sup>3</sup> | Rick Kazman<sup>4</sup>

<sup>1</sup> Jheronimus Academy of Data Science, The Netherlands

- <sup>2</sup>University of Auckland, New Zealand
- <sup>3</sup>University of Salerno, Italy <sup>4</sup>University of Hawaii, United States

Correspondence

\*Corresponding author. Email: d.a.tamburri@tue.nl

#### Summary

Forges are online collaborative platforms to support the development of distributed opensource software. While once mighty keepers of open-source vitality, software forges are rapidly becoming less and less relevant. For example, of the top 10 forges in 2011, only one survives today—SOURCEFORGE—the biggest of them all, but its numbers are dropping and its community is tenuous at best. Through mixed-methods research, this manuscript chronicles and analyze the software practice and experiences of the project's history—in particular its architectural and community/organizational decisions. We discovered a number of sub-optimal social and architectural decisions and circumstances that, may have led to SOURCEFORGE's demise. In addition, we found evidence suggesting that the impact of such decisions could have been monitored, reduced, and possibly avoided altogether. The use of socio-technical insights needs to become a basic set of design and software/organization monitoring principles that tell a cautionary tale on what to measure and what not to do in the context of large-scale software forge and community design and management.

#### **KEYWORDS:**

Software Failure, Forge Design, Automated Architecture Analysis, Community Analysis.

# 1 | INTRODUCTION

Over the years the open source software movement has contributed to a dramatic reduction of software costs and release times, while increasing general quality [18]. At the same time, the way in which open source software is built has changed radically, from the domain of just a few, to an enormous economic force. Now entire open source software ecosystems have become subject of flourishing research and practice [34] [32]. However, one of the phenomena that played a role in this change is the failure of software forges. A software forge provides a platform for hosting software projects and usually offer code hosting and bug tracking. While many software forges existed in 2011, only one seems to still be in existence: SOURCEFORGE [1] SOURCEFORGE was the biggest of all forges by count of projects, commits, and participants, but these numbers have dramatically declined in recent years.

In this study, we aim to understand the shape and indicators on the *Canary* in the *Coal Mine*, namely, we examine potential early indicators of trouble in SOURCEFORGE using mixed-methods empirical software engineering research [54]. To do this, we adopted a mixed-method research approach, by (1) conducting interviews to learn about team members' perceptions on the project and community health during the duration of the project and (2) analysing nine years of commit activity and four years of issue reports and mailing lists. Our analysis subjects are SOURCEFORGE itself, and its supporting software infrastructure.

Members of the original team perceived various social and technical issues that contributed to the decline of SOURCEFORGE. Paired to that, in the quantitative analysis, we also observed a lack of organisational stability in SOURCEFORGE. This instability occured at the same time as an

<sup>1</sup>Link: http://sourceforge.net/ Note that other more timely platforms like GITHUB provide additional collaboration tools and enable various social network analyses. They are, therefore, more advanced than classical forges; this is the reason why we do not consider them as forges.

increasing number of changes to the code base that resulted in huge increases of technical debt whose introduction may be due to the lack of communications/coordination among the contributors. Thus, a variety of social and technical factors contributed to the downfall of SOURCEFORGE and its supporting infrastructure.

**Structure of the paper**. Section 2 describes the research setting and overviews the research methodology adopted to address the posed research questions. Section 3 reports our cautionary tale on SOURCEFORGE and ALLURA its supporting infrastructure. In Section 4 we discuss the key findings and lessons learned from our study, while Section 5 summarizes the possible threats to the validity of the study and how we mitigated them. Section 6 discusses the related literature on the topic; finally, Section 7 concludes the paper.

# 2 | METHODOLOGY

## 2.1 | Research Setting

This study investigates SOURCEFORGE and its supporting software infrastructure, the APACHE ALLURA SOURCEFORGE was created by VA Software and was first launched in 1999. It was one of the first to offer free code hosting for open source software projects, a revolutionary service at the time. It offered free access to Concurrent Versioning System (CVS), a bug tracker, and mailing lists to open source projects. For many years, it was the biggest open source software development and collaboration website 56.14.

In 2009, SOURCEFORGE initiated the APACHE ALLURA platform. ALLURA is "an Open Source, extensible, web-based platform that provides integrated software tools for collaborative software development [5]". ALLURA provides an integrated issue tracker, built-in discussion forums, a code repository, and more. It was submitted to the Apache Incubator in 2012 and became a Top-Level Project in 2014 [5]. Originally, the project stems from the codebase that sustained developer tools for SOURCEFORGE (in PHP) and was redesigned using leaner programming languages and frameworks such as Python. Although it is best known as being the platform behind SOURCEFORGE, ALLURA also powers several software-intensive platforms such as the Open Source Projects Europe<sup>3</sup> the DLR German Aerospace Center<sup>4</sup> and DARPA's VehicleForge<sup>5</sup>

## 2.2 | Research Problems and Questions

In this study, we aim to understand how socio-technical factors contributed to the decline and downfall of SOURCEFORGE. We do this by considering the team member perceptions around the decline of SOURCEFORGE and examining the social and technical structures of the APACHE ALLURA software community, which is responsible for the development of SOURCEFORGE. We formulate three research questions:

- RQ1. What were the perceptions of the SOURCEFORGE team on the project and its community health?
- RQ2. Is there evidence of problems in the community and technical structures of SOURCEFORGE?
- RQ3. Is there evidence of architectural problems in SOURCEFORGE?

To answer **RQ1**, we interviewed members of the SOURCEFORGE team to learn about their perceptions on the project and community health during the duration of the project<sup>6</sup> To answer **RQ2**, we analyzed nine years of commit activity of the SOURCEFORGE/ALLURA project to identify the evolution of the source code and its structure. We also analyzed the last 4 years worth of data describing its community and organisational structure. To answer **RQ3** we analyzed the architecure of eight versions of APACHE ALLURA, focusing on its Decoupling Level (DL) and its architectural flaws.

## 2.3 | Interview Data

We interviewed four members of the SOURCEFORGE team. Our participants covered 25 years of SOURCEFORGE history. The participants were SOURCEFORGE developers, managers and analysts. Participants included junior (<2 yrs experience), senior (3-5 yrs experience), and expert (6+ yrs experience) members, in terms of SOURCEFORGE expertise. Our participants had varied backgrounds including business informatics, software operations, software engineering, management, and digital IT marketing.

<sup>&</sup>lt;sup>2</sup>https://allura.apache.org/

<sup>&</sup>lt;sup>3</sup>https://opensourceprojects.eu/

<sup>&</sup>lt;sup>4</sup>http://www.dlr.de/dlr/en/desktopdefault.aspx/tabid-10002/

<sup>&</sup>lt;sup>5</sup>https://cps-vo.org/group/avm/vehicleforge

<sup>&</sup>lt;sup>6</sup>interviewees kindly asked to share interviews selectively and therefore, interview transcripts are available upon written request.



**FIGURE 1** High-level representation of the Developer Social Network (DSN) structure we used for community smell detection showing both communication and collaboration graphs.

## 2.4 | Contribution and Communication Data

The APACHE ALLURA software repository was mined in its entirety by using the CODEFACE tool [27]. We elicited all 4 years worth of data for ALLURA about: (a) file change quantities and commit statistics; (b) total quantity of changes over time; (c) commit-message sizes and page-ranking of contents. A total of 8,902 commits were analyzed. Finally, in terms of communication data, we scraped ALLURA issue-tracking and mailing-lists, computing 3-month community structure snapshots from both sources and for all the 4 years of data currently available for ALLURA [7] A total of 32 community sociograms [23] were generated, visualized, and analyzed using time-series analysis.

A discussion of the analyses conducted on the above data is contained in the respective results sections for each research question (see Sections 3.1 3.2 and 3.3).

# 2.5 | Operationalisation and Data Mining

To attain our results, we exploited previous work in community analysis and socio-technical measurement of software development networks [44] [39]. Both works report respectively on the operationalisation and re-application of community smells and other relevant socio-technical metrics (see Table 4 from Tamburri et al. [39]) for the qualities of software processes and products. Furthermore, both works reflect extensions to the wellknown CodeFace tool, a Siemens tool for Application Lifecycle Intelligence (ALI) originally introduced by Joblin et al [27]. More specifically, the following operationalisation is used in the scope of this work (tailored from Palomba et al. [39][44]).

Starting from the developer networks built by CODEFACE, we detect instances of smells according to the formalisation below. For all of them, a premise is needed:

 $\label{eq:premise.} \text{ Premise. Let } G_m = (V_m, E_m) \text{ be the communication graph of a project and } G_c = (V_c, E_c) \text{ its the collaboration graph.}$ 

More precisely, for *communication* we mean the relation by which two or more developers communicate with each other through any channel: for example, a communication link between two developers is established in case they reply to the same discussion within a mailing list or they comment on the same issue in the issue tracker. As for *collaboration*, we mean the relation for which two or more developers have worked on the same source code elements. This is established by considering the change history of a project, looking for cases where two or more developers have modified the same code entities.

<sup>7</sup>As an example, ALLURA mailinglists are available here: http://mail-archives.apache.org/mod\_mbox/allura-dev/



FIGURE 2 Organisational Silo Effect Community Smell identification pattern.

#### 2.5.1 | Organisational Silo Effect

The Organisational Silo effect occurs when the developers break into isolated sub-communities with little or no coordination between the subcommunities [48, 46]. That is, there are two sub-teams that cannot properly communicate with each other. In the communication graph this manifests itself as two (relatively isolated) sub-graphs with just one or two people connecting them. With the occurrence of Organisational Silo effects, social debt manifests as decaying communication across sub-communities and consequent negative effects on developers' situational awareness [25] as well as degradation of projects' socio-technical congruence [48, 40]. Also, according to recent findings [48], the Organisational Silo effect may lead to tunnel-vision, since participants may focus their cooperation and communication solely on other members of their narrow sub-community rather than on the broader community. Finally, community members belonging to an *Organisational Silo* may exhibit egotistical behaviour leading to unsanctioned architectural decisions [49] as well as defiance of the decisions of others [48].

Based on this definition, we define the set of Organizational Silo pairs S as follows:

$$\{(v_1,v_2)|v_1,v_2\in V_c,(v_1,v_2)\not\in E_m^*\}$$

where  $E_m^*$  is the transitive closure of  $E_m$ . With *transitive closure* we indicate the transitive closure of a graph. More specifically, given a directed graph, the operation finds out if a vertex j is reachable from another vertex i for all vertex pairs (i, j) in the given graph. With *reachable* we mean that there is a path from vertex i to j. The reach-ability matrix is called *transitive closure* of a graph. For the sake of precision, we capture the Organisational Silo Effect at the finest grain possible, i.e., that of collaboration dyads: pairs of co-committing developers. An example is shown in Figure 2 Here the Organisational Silo effect is reflected on developer "1", who does not communicate with developer "2" even though "1" is collaborating with "2". Conversely, developer "2" is communicating with (at least) one other developer, "3", who belongs to a sub-community other than "1". Considering the example proposed in Figure 1 an operationalisation of the identification pattern for the Organisational Silo effect has two steps. In the first step, the identification mechanism compares the collaboration network (bottom half of Figure 1) with its communication counterpart (top half of Figure 1). Then it verifies that the developer identified by the letter A is present in the collaboration network, i.e., A commits to files co-committed by others, but is not present in the communication DSN reflecting those files.

#### 2.5.2 | Lone Wolf Effect

The *Lone Wolf* community smell reflects circumstances in which communication may indeed be present but insufficiently addressing project needs [48] 46]. The result is developer free-riding and unsanctioned architectural decisions that cause nasty ripple effects such as code duplication and churn [25]. Thus, we define the set of *Lone Wolf* pairs L as the set of collaborators that do not directly or indirectly communicate with each others. More formally:

$$\{(v_1,v_2)|v_1,v_2\in V_c,(v_1,v_2)\in E_c,(v_1,v_2)\not\in E_m^*\}. \text{ By definition } L\subseteq S.$$

The identification pattern for the *Lone Wolf* smell is based on the detection of development collaborations between two community members that have intermittent communication counterparts or feature communication by means of an external "intruder", i.e., not involved in the collaboration. A simple example is given in Figure <sup>3</sup> In this example two developers, "1" and "2", are collaborating on some code, but they are not connected by any communication link other than developer "3", who is not co-committing on a shared file. In this case, either developer "1" or developer "2" (or both) can develop a *Lone Wolf* community smell.

This smell reflects the presence of possible side effects generated by the Organisational Silo such as communication decay or negative influence on developer awareness and heavy socio-technical congruence degradation. Our conjecture is that the occurrence of the Organisational Silo effect is



FIGURE 3 Lone Wolf Community Smell identification pattern.

not negative per se. But when that occurrence is compounded by the occurrence of *Lone Wolves*, extra attention must be paid to avoid negative consequences such as delays and unmanageable social debt. The *Lone Wolf* smell reflects dyads of co-committing (collaborating) software developers who exhibit uncooperative behaviour and mistrust by not appropriately communicating.

#### 2.5.3 | Black-cloud Effect

The Black Cloud Effect leads to negative social interactions within a software development community featuring: (a) community members' inability to cover knowledge and experience gaps between two different software products developed within the same software community [48]; (b) lack of periodic and official opportunities to share and exchange knowledge between all community members [48]46]. Whenever these two circumstances occur together, they can generate a "black-cloud" of misinformation (e.g., confusing, delayed, or unnecessary communication that generates communication overload) within the community. The main consequence of the *Black Cloud* effect is to obfuscate project vision, compromising progress [48]46]. The occurrence of this community smell can be generated or worsened by several socio-technical triggers:

- absence of ad-hoc protocols for knowledge sharing;
- lack of boundary spanners;
- presence of inefficient communication filtering protocols.

Moreover, the *Black Cloud* Effect smell is associated with several other side effects such as: lowering of trust between developers, information obfuscation, as well as inception of the *Organisational Silo* Effect, due to the rise of egoistic behaviour. The identification pattern for the *Black Cloud* Effect smell reflects sub-communities that in subsequent release windows do not communicate, with the exception of two community members (i.e., boundary spanners in social-network jargon [53]), one belonging to each sub-community. The detection of the *Black Cloud* instances starts with the identification of vertex clusters as already implemented in CODEFACE. More specifically, let  $P = \{p_1, \dots, p_k\}$  be a mutually exclusive and completely exhaustive partition of V<sub>m</sub> induced by the clustering algorithm. From the partition, *Black Cloud* is the set of pairs of developers C that connect otherwise isolated sub-communities, more formally:

$$\{(v_1, v_2) | v_1, v_2 \in V_m, (v_1, v_2) \in E_m, \forall i, j(((v_1 \in p_i \land v_2 \in p_j) \Rightarrow i \neq j) \land \forall v_x, v_y((v_x \in p_i \land v_y \in p_j \land (v_x, v_y) \in E_m) \Rightarrow v_x = v_1 \land v_y = v_2))\}$$

The smell manifests if the above condition holds for at least two consecutive organisational time-windows (fixed to 3-month intervals, in the case of CODEFACE4SMELLS). An example is presented in Figure 4 Here the occurrence of *Black Clouds* reflects two developers, "3" and "4", who are the lone boundary spanners across two different sub-communities and over time—at least two subsequent analysis windows (3 months, in our case).

Detecting black clouds requires eliciting the communication network and applying known community detection algorithms 36 to identify subcommunity structures and boundary spanners across them. For example see Figure 1 where two sub-communities (previously specified) can be detected by considering the density of communication links.



FIGURE 4 Black-cloud Effect Community Smell, an identification pattern.

#### 2.5.4 | Bottleneck Effect

The *Bottleneck* community smell is characterised by the occurrence of the following sub-optimal characteristics within a software development community: (a) proposed changes within every software development phase require an extraordinary quantity of time to be implemented [48]; (b) time waste [46]; (c) hidden or counterintuitive information (and broker) locations [48]; (d) highly formal or complex organisational structure [48]; (e) highly regularized procedures [48]; 46].

The fundamental side-effect generated by this community smell is a massive delay that characterises key organisational processes within the community such as decision-making, due to personnel unavailability or communication overload. The identification pattern of this smell is based on the detection of unique knowledge and information brokers in different sub-communities.

In our attempt to define an automatic identification pattern for this community smell we focused on the analysis of project communication networks. We considered the six key factors around *Bottleneck* as reflecting the presence, within a project organisational structure, of a unique boundary spanner across several different sub-communities (i.e., more than 2). The social-network analysis concept of unique boundary spanner [53] has, in fact, a remarkable similarity to *Bottleneck*. A unique boundary spanner interposes him/herself into every formal interaction across two or more sub-communities and if the organisational structure of the project is complex and characterised by highly formal procedures, it will not be possible to incept parallel information channels between other members of the sub-communities. From a formal perspective, we define the set of *Bottlenecks* as:

$$\begin{aligned} \{v | v \in V_m, \exists i (v \in p_i \land \forall v_x (v_x \in p_i \Rightarrow v = v_x))\} \cup \{v | v \in V_m, \exists v_x, i, j (v \in p_i \land v_x \in p_j \land (v, v_x) \in E_m \land \forall v_y, v_z ((v_y \in p_i \land v_z \in p_j \land (v_y, v_z) \in E_m) \Rightarrow v_y = v)\} \end{aligned}$$

To further elaborate on the definition of this community smell, let consider the example proposed in Figure [5] As shown, detecting *Bottleneck* requires the identification of community members who are the <u>only</u> members of their sub-community that communicate with (at least) two other sub-communities. Therefore, assuming a communication link was present between developer "A" and "B", then developer "B" is the pivot of *Bottleneck*.

# 2.6 | CodeFace4Smells Extension

From the perspective of the CODEFACE tool, a *community* is operationalised as a densely connected set of nodes within the community group (i.e., the members that make up a development community) which is sparsely connected to all other nodes in the network. To identify and properly characterise the community structure, the CODEFACE tool enacts two community detection strategies, defined as follows, paraphrasing from Joblin et al. [27]:

Function — To recover a community structure, CODEFACE uses a heuristic for identifying when two developers are engaged in a coordinated
effort using a fine-grained heuristic based on code structure, where developers are considered to be coordinated when they actually contribute code to a common function block. Furthermore, CODEFACE uses the commits' timestamp for identifying the appropriate directions
of the edges in the recovered community structure.



FIGURE 5 Bottleneck Community Smell.

- Committer-Author In this method, the tool uses tags to identify relationships between all people that contributed to a common commit, including authors, reviewers, and testers. For example, sign-off tags are self-reported acknowledgments of participation on a commit, therefore the tag-based networks undoubtedly capture real-world collaboration [27].
- 3. Community-Verification to verify the recovered community structure, CODEFACE uses a random null-model to compute the probability of observing the identified community in an equivalent class of null-model graphs that lack a community structure. The tool generates the null-model using a standard approach called the configuration model for random graphs, where nodes are joined uniformly at random under the constraint that the degree distribution is identical to the observed graph [21].

To the above heuristics, we add a systematic implementation of the operationalisation provided in the previous section, to allow for automated detection of community smells at the same time as CODEFACE operates community structure recovery. The output of the tool is represented by a CSV file containing the community smell instances identified over a social structure representation known as a Developer Social Network (DSN), a notation previously used for bug prediction and error-proneness [41,4].

# 3 | ANALYSIS OF THE RESULTS

## 3.1 | Perceptions of Team Members

To address RQ1, we interviewed four key forge managers, designers, developers, and operators who, together, cover the entire 25 year history of the forge. We supplemented this data through the analysis of blogs and news articles that documented the history of SourgeForge. This supplementary data was used to obtain specific details of events described by the participants (e.g. to identify the details of a change in ownership). In cases where details were obtained from news articles or blogs, they are referenced.

#### 3.1.1 | Research Methods

We performed semi-structured interviews. The goal was to understand their perspectives on project and community health. Each interview lasted about one hour. The interviews were transcribed by an independent third-party. Thematic content analysis was used to analyze the interview transcripts [6].

#### 3.1.2 | Results for RQ1

*Early Success.* We asked the interviewees what factors they perceived to have led to the early success of SOURCEFORGE. The interviewees described three main factors:

- Emerging Need. SOURCEFORGE addressed an emerging need. Offering free code hosting when SourgeForge launched in 1999 was revolutionary. One participant said "we wanted like 1,000 projects in the first year, and we had 1,000 projects in the first month." By 2007, there were more than 150,000 users and over 1.5 million projects [52]. Thus, SOURCEFORGE grew much larger and faster than expected.
- First To Move. Similarly, SOURCEFORGE was the first to offer free and versioned code-hosting. As the demand for this service grew, SOURCEFORGE's popularity exploded. This can largely be attributed to the fact that they were the only ones offering this service in the beginning.

• Skunkwork Team. Interviewees also suggested that the Skunkworks nature of SOURCEFORGE contributed to its aggressive expansion of popularity. Skunkwork teams are a small group of people who work on a project in an unconventional way; the group's purpose is to develop solutions quickly with minimal management constraints. In the case of SOURCEFORGE, there was no direction from management on what should be developed. The project itself happened by happenstance. The team of four developers were given the goal of generating online traffic. The developers came up with the idea of creating what would eventually become SOURCEFORGE because of their passion for open source software and their belief that it would attract traffic. New features originated within the team. One participant said "we did the development of SOURCEFORGE on SOURCEFORGE. So as we needed a feature ourselves we would write it." Thus, the team had significant freedom in deciding what to develop.

The early success of SOURCEFORGE was not seen to be related to the technical competitiveness of the project, but rather being the first to fill an emerging need.

*Eventual Downfall.* We also asked the interviewees about the project and team health throughout the project and what (if anything) they would have done differently. We identified several factors that interviewees believe eventually led to the downfall of SOURCEFORGE:

- Not Considering ROI. Due to the Skunkwork nature of the project, the team paid little attention to how much money was being spent or how much money was coming in. One participant said "it was never part of the plan to make money." However, after the company went public, interviewees described an emerging need to track the return-on-investment (ROI) that the forge was producing (if any). For example, one of the interviewees mentioned that "[...] a lot of hardware was coming in and nothing was coming out, so people in the high places started to ask questions". Another interviewee said "once we were public and now we had responsibilities to shareholders and things like that, people obviously started wanting to know where all this money was being dumped into, and wanted to know how we were going to return on investment and things like that. And so six, eight months later is when we started actually getting pressure from executives to figure out how we were going to make money."
- Deceit. Due to the push to increase ROI, in July 2013, SOURCEFORGE introduced a new program, called DevShare. This program bundled third-party software with project downloads, following the model of the widely known CNET download network<sup>8</sup>

DevShare was conceived as a way for open source software projects to monetize their efforts while still keeping the software open source and free. The ad revenue would be shared with the projects [20]. The developers thought that enabling ways for the OSS projects to make money was a good goal, but the way DevShare was implemented was too dishonest, putting ROI over trustworthiness. One participant said, "SOURCEFORGE has been trying to help projects to sustain what they were doing with dubious initiative like Devshare. Might have been a good thing if that was run the right way. As a matter of fact today projects still need money and they don't have the solve for that."

While DevShare was an opt-in service, some projects complained that SOURCEFORGE bundled third-party adware in their downloads without their consent [43]. Ads were added to project download pages with fake download buttons to trick users into clicking on the ad. Often, clicking on these ads resulted in the download of adware. At this time, many projects announced they were abandoning SOURCEFORGE, citing DevShare as one of the main reasons. The GNU Image Processor (GIMP) was the first big project to announce it was leaving in November 2013 [43].

• Two-Masters Syndrome. Part of the reason DevShare was implemented was driven by some of the bigger OSS projects using SOURCEFORGE at the time. Many large projects called SOURCEFORGE home, including VLC Media Player and GIMP. At this point, there were plenty of competitors providing source-code hosting, and SOURCEFORGE wanted to keep the projects. One participant said "We had a lot of bigger projects on the site, and then we had lots and lots of little projects right. So the bigger projects, they wanted some kind of revenue sharing. ... we started catering to the really big projects and trying to implement enough of the little features that the smaller projects wanted to stay as well."

The two-masters syndrome 24 is a partially unknown organisational effect 13 where the team needs to work to satisfy two "masters" while making money out of both, but both end up having a conflicting agenda that creates an impasse.

• Organizational instability and disconnect between management / developers. SOURCEFORGE changed ownership many times. One participant said "every 18 months we'll have a new owner and a new set of managers. At one point in time we all got together for beers and wrote them down, and it was almost every 18 months on the dot." With these changes in ownership, there were also changes in direction and a disconnect between management and the development team. A participant complained that "by 2008 we had already actually changed hands of who owned us and what we had been doing two or three times, that the new people that had come in and acquired us, they didn't care about the open source ethos, they didn't care about anything except for making money back."

DevShare was introduced after one of the changes of ownership without much consideration for the opinions of the developers, causing many to leave the team. One participant said "Because the company was sold one of the many times it was. But it was sold at that time, so changing hands, the new owners decided to try to give a spin using the Devshare program among other things. So this is not a team decision, it was of course a company decision. So the idea was, let's try to do this and see if that might help, but if you talk about the developer team, some of them decided to move on and went to join other initiatives. So not everyone decided to stay at SOURCEFORGE at that time."

- Skeleton Crew. There were very few core contributors to the SOURCEFORGE project. One interviewee said "there was three main developers and the manager type that was also a developer but working part time. We needed more people, and we just couldn't get people. We needed more hands." The unexpected popularity and growth of SOURCEFORGE coupled with the need to serve two masters with conflicting agendas, meant that the small skeleton crew was insufficient.
- Blinded by Technical Debt. The small skeleton crew was further complicated by the grassroots start of the project. At the time when GitHub was released, the SOURCEFORGE team was in the midst of a complete refactoring and redesigning of their code base. One participant said, "A little bit before that, 2008 or whatever, we had started rewriting the entire site, trying to pay down the almost decade of technical debt that we'd accumulated." This effort prevented the team from fully noticing the disruptive change that was occurring in the software landscape with the advent of collaborative development tools.
- Missed Paradigm Shift. Meanwhile, new collaborative, highly-distributed software development, hosting, and versioning tools were disrupting the market. Interviewees stated they believed these were fads and did not consider modernizing SOURCEFORGE to keep up to date with this paradigm shift. This was likely the result of having an overworked, skeleton crew who did not have time to really step back and look at the changing landscape. One interviewee said, "I don't think we quite got the importance of, you know, of the social element ... we were still in very much a sort of dot com 1.0 framework of content producers and content consumers being very distinct populations."

The quality of service went down and projects started moving away. The exodus of projects was facilitated by the advent of GITHUB, which was launched in 2008. GITHUB, a competing code hosting website, differed from SOURCEFORGE as it was built on top of git, a distributed version control system. GITHUB also offered many collaboration features. In June 2011, ReadWriteWeb reported that GITHUB had surpassed SOURCEFORGE in total number of commits for the period January to May 2011.

There were many factors that seem to have contributed to the downfall of SOURCEFORGE. Yet, all of the factors seem to be connected in various ways. Fig. 6 offers a chronicle of the SOURCEFORGE story. The figure plots participation (number of users shown by the continuous line) and size (total lines of code (LoC) added shown by the dotted line) over time. The participation and size numbers were obtained from the SOURCEFORGE Research Data Archive (SRDA) [52]. The figure also highlights the major external events that pertained to SOURCEFORGE (top part) as well as the internal re-organizations (bottom part) where we highlight the re-organisation start and finish (forked-line arrow from the bottom). The figure clearly illustrates a *decline* in the population and size starting in 2011. In 2013, after the introduction of DevShare, we witness the *downfall* of SOURCEFORGE. By downfall, we mean an abrupt, steady, and continuing loss of people and projects.

**Summary for RQ1.** The results coming from the semi-structured interviews confirm that the decline of SOURCEFORGE has been facilitated by both social and technical aspects around the software development community.

## 3.2 Community and Technical Structures of SOURCEFORGE

To address RQ2, we studied the evolution of APACHE ALLURA from both social and technical perspectives.

#### 3.2.1 | Finding Evidence of Social and Technical Debt

From a socio-technical perspective, we sought out the most established indicators of social debt [49]48], namely, (1) socio-technical incongruences, or breaks in socio-technical congruence [3] as well as (2) the presence of sub-optimal patterns in the community structure, or *community smells* [45]. Specifically, in this study we considered three of the community smells defined by Tamburri et al. [48]: Organisational Silo Effect, Lone-wolf Effect, and Bottleneck or "Radio-silence" Effect, which were previously defined.

Furthermore, in addition to the previously defined community smells, we considered two well-established measurements for socio-technical issues, namely:

1. Smelly-Quitters. This ratio reflects the number of people who were part of a community smell for two subsequent time windows and left the community for the remaining time windows in the available range of data 1.



FIGURE 6 A chronicle of SOURCEFORGE using participation (#people, continuous line) and size (total #LoC added, dotted line) over time. Indications of a declining Forge are observable earlier than the DevShare incident and coincided with many developers leaving; downfall (abrupt drops of both people and projects) manifests 2 years later.

2. Socio-Technical Congruence. Paraphrased from previous work [51] as "the state in which a software development organisation harbors sufficient coordination capabilities to meet the coordination demands of the technical products under development" and operationalised in this study as the number of development collaborations that do communicate over the total number of collaboration links present in the collaboration network.

Along with community-related information, we also collected technical data. In the first instance, we studied how the quantity and sizes of file changes applied by developers vary during the evolution of ALLURA. Secondly, we extracted data related to technical debt. In particular, we considered two different sources of information:

- Code smells. These represent sub-optimal design or implementation solutions applied by contributors during the evolution of the project [19]. While the most-known code smell detectors only work for Java programs [17], we needed an automatic tool able to identify design issues in Python. For this reason we employed the detector proposed by Chen et al. [10], which can identify 10 different smell types including both traditional code smells (e.g., Large Class and Long Method [19]) and Python-specific ones (e.g., Long Ternary Conditional Expression [10]). We relied on the original implementation of the tool made available by the authors. It is important to note that this detector is the only one currently supporting the detection of Python smells; its accuracy has been reported to be high (average precision of 98%), being therefore suitable for our purpose.
- Self-admitted technical debt (SATD). Potdar and Shihab [42] defined SATD as a practice that developers use to admit the existence of temporary design solutions (at any level, from requirement to code debt) that should be fixed. To identify SATD, we exploited regular expressions to match inside comments the 62 patterns defined by Potdar and Shihab [42]. This list includes a set of keywords that are likely to indicate the presence of a SATD (e.g., fixme, this is a hack, etc.), and has been defined by manually analyzing more that 100,000 code comments [42].

#### 3.2.2 | Results for RQ2

Results of our socio-technical data synthesis and analyses are reported in Figures 7 to 10. In the following, we first report on the evolution of number of developers as well as community smells in the considered period; furthermore, we decribe how additional socio-technical factors such as the socio-technical congruence and the smelly-quitter ratio evolved in the same period.

First, Figure 7 outlines the numbers pertaining to core (i.e., contributing to both communication and committing) and periphery contributors (i.e., contributing to communications only) to the APACHE ALLURA project. The figure shows a heavy fluctuation of all contributor types across our data sample, with a standard deviation of 4.6 developers per every 3-month snapshot, meaning that the community acquired and lost an average of 4-5 developers around every considered snapshot.



FIGURE 7 Numbers of developers for APACHE ALLURA - highlighted is the the period we denoted with downfall.



FIGURE 8 Numbers of community smells for APACHE ALLURA - highlighted is the the period we denoted with downfall.

Furthermore, Fig. Boutlines the progression of the reported numbers of community smells in project APACHE ALLURA — the figure also highlights in the grey area to the left-hand side, the period we denoted with downfall on Fig. The figure highlights a coincidental drop of community smells (deviation of -60%) all along the entire period of downfall. The growth of smells begins again between late 2014 and mid 2015, precisely around the corporate acquisition of SOURCEFORGE/ALLURA by yet another external company, the well known videogame giant GameStop, in this case. The figure also shows an almost identical drop (standard deviation – is 0,5) of the number of community smells coincidentally to a final acquisition by SOURCEFORGE/ALLURA between the end of 2015 and mid 2016 by its current owner, private company BIZX.

Finally, Fig. 10 plots results for the smelly-quitter ratio across APACHE ALLURA. The figure shows a recurrent pattern clearly appearing twice (see the greyed-out areas on the figure, degree of isomorphism >80%) in the project's community structure and, not surprisingly, once again coinciding with the two subsequent organisational rewirings in the respective periods.

**Summary for RQ2, Part a) - Socio-Organizational Perspective.** Our evidence shows recurrent organizational turmoil patterns that coincide with organisational rewiring scenarios around the forge. This evidence suggests that a lack of organisational stability in SOURCE-FORGE/ALLURA.

Furthermore, while organisational change is inevitable, most organisations do not track it. Tracking emerging metrics such as the number of community smells or the ratio of "smelly" quitters offers a reliable and fine-grained (e.g., see Fig. 10) perspective over the effects and manifestations of organisational rewiring. The recurrence of a drop in smelly-quitters likely indicates a recurring turmoil coinciding with the organizational rewiring taking place around the forge.

The recurrence of the above patterns is reflected partly on socio-technical congruence (see Fig. 9) and even more clearly in the smelly-quitters ratio (see Fig. 10). In particular, Fig. 9 shows steadily increasing congruence values throughout 2014 and late 2015, when a blatant drop in socio-technical congruence happens around APACHE ALLURA, at the time of the acquisition by BIZX. This suggests a misalignment between the

**Socio-Technical Congruence** 



FIGURE 9 Socio-Technical Congruence for APACHE ALLURA - highlighted is the the period we denoted with downfall.



FIGURE 10 Smelly-Quitter Ratio for APACHE ALLURA - highlighted is the the period we denoted with downfall.

organizational structure of developers and operators with respect to the acquisition by BIZX; the drop likely indicates the turnover surrounding this acquisition.

From a technical perspective, we analyzed the (i) quantity and sizes for code changes applied and (ii) code smells and self-admitted technical debt detected on APACHE ALLURA over the entire analysis time-window running from early 2014 until late 2016, extracting essential code quality metrics to control code quality as well. Figures 11 to 14 outline the achieved findings.

In particular, we observed that the socio-technical observations made in Sec. 3.2.2 recur. For example, in coincidence to the same organizational rewiring scenarios, code-changes are localized (e.g., the top-right plot from Fig. 11 highlights an series of changes focused on 5-7 files), the difference-in-size range is almost constant (bottom-left plot on Fig. 11 range fluctuates around 100 lines) and the code-commit tags disappear after the first 2 periods of our analysis while the size of commit messages is almost constant and quite considerable, between 50 and 100 words – this evidence seems to denote a constantly changing code-base around an organization that disconcerts the use of typical Apache projects' coordination practices such as commit-tagging.

Similarly, Fig. 12 provides another confirmation of our organizational turmoil observations as reflected on the technical artifacts. In particular, all highest plot "peaks" in the top-most and second plot on Fig. 12 reflect time-ranges coinciding with the afore-mentioned organizational turmoil.



FIGURE 11 APACHE ALLURA Change sizes (Y-Axis) per commit (X-Axis).

Even more importantly, in coincidence with the last organizational acquisition, a massive campaign of code-changes was initiated, as if the forge system needed again to undergo major redesign.

Furthermore, Figures 13 and 14 are perfectly in line with the results discussed so far. Indeed, we observe that during the time period referred as downfall in Fig. 1 both code smells and self-admitted technical debt tended to decrease, while their numbers increase once the organizational structure of ALLURA became unstable. It is important to remark that, despite the absolute number of code smells and technical debt is not that high (e.g., at most, we observed 23 Large Class instances), it may still substantially contribute to the drift of the design of a software system, possibly compromising the overall quality—as demonstrated by previous work in the field [2] 38. The observations above are especially true when considering the joint fluctuation between community smells and three particular code smell types, i.e., Large Class, Long Method, and Long Parameter List. This close relation is confirmed from a statistical point of view: specifically, we applied the Granger causality test [22] to determine whether one time series (i.e., the introduction/removal of community smells) is useful in forecasting another (i.e., the introduction/removal of code smells). In other words, we tested whether the presence of a community smell c<sub>i</sub> can be used to "predict" the presence of a code smell cs<sub>i</sub>. Note that we used the Granger test instead of association rule discovery [2] because we are interested in assessing the statistical significance of the temporal relation between community and code smells, rather than of just their co-occurrences. As a result, we found that relation between community smells and their co-occurrences. As a result, we found that relation between community smells and the  $\rho$ -values are < 0.001).

Summary for RQ2, Part b) - Technical Perspective. Our evidence shows recurrent organizational turmoil patterns influencing technical stability. The combined technical and socio-organisational evidence denotes an organisational and technical structure in downfall, featuring major restructurings and negative manifestations across both perspectives.

# 3.3 | Architectural Decisions of SOURCEFORGE

To address **RQ3** we analysed the software architecture of eight releases of the Apache Allura project covering a four year time-span, from early 2014 to early 2018. These were releases: 1.1.0, 1.2.0, 1.3.0, 1.4.0, 1.5.0, 1.6.0, 1.7.0, and 1.8.1.



15





**FIGURE 13** Numbers of code smells for APACHE ALLURA; an evident peak reflects the frantic refactoring connected to the period in which the decline started, a clear indication of technical turmoil and haphazard technical debt generation - highlighted is the the period we denoted with downfall.

# 3.3.1 | Research Methods

We analyzed the software architecture of each of these releases. Or, more precisely, we analyzed the module structure of each of these releases using a commercial version of the Titan tool suite 33 called DV8<sup>9</sup> This tool suite allowed us to reverse engineer each release of Allura, and perform two analyses:

- measure the Decoupling Level (DL) of the module structure of each release [35];
- measure the architectural flaws in the module and package structure of each release [35];

Each of these measures has been shown to be highly correlated with maintenance effort and bugginess [35].



FIGURE 14 Numbers of SATD items for APACHE ALLURA; the conflict between the SATD and the smells reported indicates an extent of failure reticence, people cannot admit the frantic refactoring is not helping but only making things worse - highlighted is the the period we denoted with downfall.

Release	Date	DL Score
1.1.0	1/2014	32%
1.2.0	12/2014	34.6%
1.3.0	6/2015	34.6%
1.4.0	4/2016	33.8%
1.5.0	8/2016	33.7%
1.6.0	12/2016	33.5%
1.7.0	3/2017	33.5%
1.8.1	3/2018	33.5%

#### **TABLE 1** DL Scores from Eight Releases of APACHE ALLURA.

#### 3.3.2 | Results for RQ3

The DL scores of all eight releases are shown in Table 1 There are two points to note about these scores:

- The scores are uniformly low. These DL scores place APACHE ALLURA in the bottom 10th percentile of all projects analyzed in [35], where 129 projects, both commercial and open source, were measured.
- The DL scores are relatively stable over time, decreasing slightly from release 1.2.0 to 1.8.1. This suggests that the team was not attempting to address any of the architecture debt [55] that had accumulated.

DL is a measure of how (de-)coupled the system's source code files are. A low measure suggests a highly coupled system, one that is likely to be difficult to understand, debug, and modify. But this is just a single number. To understand these low DL scores more precisely we analyzed the architectural flaws (also known as "hotspots") in each of the APACHE ALLURA releases. These flaws have been shown, in several studies ([35][16]), to be very strongly correlated with bugs, changes, and churn across a wide variety of projects.

The rationale behind performing this analysis is that we wanted to measure the flaws over time. This would help us understand whether the developers were attempting to pay down the architecture debt, by removing at least some of the architectural flaws.

As for the DL analysis, the results are summarized in Table 2 In this table we report on the number of instances of four flaw types: Unhealthy Inheritance (UnhInt), Cliques (Clique), Modularity Violations (ModVio), and Package Cycles (PkgCyc). These flaw types are described in 35. As can be seen from the table, number of flaws in this project is large and it remains stable or even increases slightly from release to release. This suggests

Release	UnhInt	Clique	ModVio	PkgCyc
1.1.0	11	5	186	17
1.2.0	13	6	232	17
1.3.0	13	6	238	18
1.4.0	13	4	254	17
1.5.0	13	4	254	18
1.6.0	13	4	256	18
1.7.0	13	4	253	17
1.8.1	13	4	250	18
		•		

TABLE 2 Architectural Flaws in Eight Releases of APACHE ALLURA.



FIGURE 15 APACHE ALLURA activity, in commits per month, 2014-2018 – picture was extracted via OpenHub.

that the architecture had a heavy load of debt and that this debt was not being addressed. As such it was increasing slightly with each subsequent release, making the project harder to maintain and evolve.

This observation is consistent with the recorded activity, in terms of commits, on the Allura project over the studied time-frame. Over this time-frame commit activity dropped precipitously, as shown in Figure 15 taken from openhub.com. In January 2014 there were 140 commits per month. By January 2018 there were just 28 commits per month. While we can not claim a causal relation here, this change in activity is consistent with our expectations.

Summary for RQ3. The project had a heavy and continuously increasing load of architectural debt and this debt was never repaid.

# 4 | DISCUSSION AND IMPLICATIONS

We identified a variety of both social and factors that contributed to the decline and downfall of SOURCEFORGE. These factors were both internal and external. SOURCEFORGE's early success was because they were the first to fill a need. At the time, their skunkwork team was a benefit and enabled them to identify and fill a gap in the market quickly. However, as the popularity of SOURCEFORGE grew, the skunkworks team was no longer advantageous. The team tried to serve two masters (small projects and large projects) with different and, at times, conflicting needs. The skeleton crew introduced more and more technical debt as they tried to meet the needs of all of their stakeholders quickly. While they focused on dealing with technical debt and continued to try to satisfy the needs of their stakeholders, they missed a large paradigm shift that was occurring - the movement to distributed version control systems and more collaborative development tools. GitHub's growth became exponential while the SOURCEFORGE team largely dismissed it as a fad.

Meanwhile, SOURCEFORGE changed ownership several times, and a disconnect between the development and management started to cause major problems. Management focused on ROI; SOURCEFORGE was expensive to run and did not have a plan to bring in revenue. This led to the introduction of DevShare, but since management did not understand the open source ethos and the development team was not included in management decisions, DevShare was a major failure. It prioritized ROI over trust and bundled adware with project downloads. Many projects started leaving SOURCEFORGE, citing DevShare as a main reason.

Through analysis of the social and technical structures of the project, we also observed a lack of organizational stability and technical and architectural debt. There is substantial indication of community smells presence in the organizational structure around SOURCEFORGE, indicating intense organizational turmoil and lack of stability [31]; at the same time the architecture of Allura is reportedly highly coupled, replete with flaws. The team around the forge never attempted to pay down this architecture debt, presumably too focused on paying back the more basic, code-level, and self-admitted technical debt. There are several lessons that can be learned from this cautionary tale. Here we highlight lesson for other software projects and directions for future research.

- Advice: Determine business plan early. One of the key factors that contributed to the major exodus of projects was DevShare. While the deceitful nature of its implementation is largely to blame, it also illustrates how important it is for a software project to identify its ROI strategy early. DevShare was an effort to increase ROI since the original project did not have a plan.
- Advice: Management and development teams must be aligned. Another contributing factor to the downfall was the misalignment between management and the development team. The developers clearly indicated that they were not involved in key decisions, despite having more expertise in the domain. This disconnect caused an already small team to lose valuable talent, but also caused poor decisions that did not align with the ethos of the community their product served. The aforementioned condition is a known phenomenon in organisations and social networks research often referred to as "two-masters syndrome" [29,15], which could itself be a *community smell* which was not previously manifested nor formalised in software engineering research and practice and may deserve further attention at least to support collaborative and computer supported cooperative work in the scope of software produciton and operation platforms such as GitHub.
- Advice: Be careful of changing paradigms. There were many reasons for the missed paradigm shift, particularly they had an overworked, skeleton crew who were blinded by technical debt. Yet, one important aspect to discuss is the expertise of the team. The developers of SOURCEFORGE were designing and developing software for a domain in their own area of expertise software development. They were very comfortable with centralized version control systems and did not appreciate the importance of incorporating collaboration into software development tools. This high level of domain expertise, in a way, hindered their ability to appreciate the change that was happening since they were very likely considering their own needs as a software developer. Given the small size of their own team, collaborative tools may not have seemed important.

Software projects should be careful to continue to assess the landscape in which they operate. Diverse teams can help ensure different perspectives about changing landscapes are considered.

Future research: More comprehensive project and community health metrics are needed. Our evidence showed that SOURCEFORGE suffered from a variety of social and technical problems. We cannot claim that additional project and community health tracking would have prevented the downfall of SOURCEFORGE, but we do observe that using some of the community smells and technical and architectural debt metrics we examined in this study would have revealed problems. Potentially if some of the problems had been revealed earlier, the fate of SOURCEFORGE would have been different. Future research should develop tools that enable the tracking of more fine-grained metrics related to the community, the code, and the architecture.

Turnover (e.g., as measured by its earlier manifestations such as the Bus-Factor [12]) and smelly-quitters are the only two quantities that we investigated which exhibit variability both at the micro-structure (single or small-world network interactions) and at the macro-structure level. The only apparent exception to the turnover and smelly-quitters variability in the organizational structure of the forge under study is in the very beginning of our sample of observations, namely, the period between early 2014 and mid 2015, where the standard deviation for both drops by half a point — this period coincides almost identically to the *downfall* time range from the timeline reported in Fig. [6] The coincidence could indicate that Forge designers and managers detected the negative trend, trying to pick up the disaster scenario caused by DevShare and sought external collaboration as well as increasing the number of paid maintainers. On the other hand, the indication in question could have been aided by automated-tracking of turnover and smelly-quitters [39] emerging from organizational tracking. *Further research should be invested in these factors to correctly establish the feasibility of the aforementioned metrics in this context.* 

There is also an intrinsic relationship between the social, technical, and architectural factors. Thus, additional comprehensive metrics should be developed beyond the existing measures like socio-technical congruence. While some initial work on this has been carried out [39], more research is needed on the relation between social and technical debt as well as on the methods for assisting developers in reducing the joint issues coming from such relation. Some communities are already forming to develop additional health metrics, for example, the CHAOSS initiative<sup>10</sup> or the SECO-health FNRS international project<sup>11</sup> Future efforts should develop tools that enable comprehensive analysis of project and community health that consider the relationships between the social, technical, and architectural components of a project.

- Future research: Analysis of failures and Retrospectives. While there has been some recent work that studies the failure of software projects (described in Section, additional investigations are needed to understand more deeply the reasons for failure of declines of software projects. Additional case studies could be performed to add to the evidence collected here.
  - <sup>10</sup>https://chaoss.community/

<sup>&</sup>lt;sup>11</sup>https://secohealth.github.io/

# 5 | THREATS TO VALIDITY

Like any study of comparable magnitude and scale, this study is affected by several threats to validity. In what follows we outline the major ones in our study design and execution.

Internal and Sampling Validity. Internal validity refers to the internal consistency and structural integrity of the empirical research design. We focused our study on APACHE ALLURA and SOURCEFORGE individually, operating a mixed-methods research approach also adopting several observer, data, and sample triangulation strategies (e.g., an additional set of our study of the Forge features an interview dataset whose coding was executed twice and  $K_{\alpha}$  evaluated) being adopted. This notwithstanding, there are up to 90 factors from the state of the art in organisations research [47] that may still be affecting our findings and results. Also, the quantities and effect sizes of the factors themselves were not addressed in this study. Stemming from this limitation, we are planning further study of our target subjects in follow-up quantitative and qualitative research over the projects hosted on SOURCEFORGE or further interested parties involved with the forge (e.g., developers for those projects) that may confirm the validity of this work.

**External Validity.** External validity is the degree to which results from a study may be generalized to other contexts. One threat to external validity is the small size of our interviewee pool —just four project members were interviewed. While four is a small number of interviewees, they represent a large percentage of the core developer team, and so we believe that their opinions are likely to be accurate representations of the project. Further, our analysis considered only one organisation, SOURCEFORGE, so may not generalize to other projects. Future studies should investigate if the social and technical factors we identified also contribute to the downfall of other software projects.

**Conclusion Validity.** Conclusion validity represents the degree to which conclusions about the relationship among variables are reasonable. In the scope of the discussions of our results we made sure to minimise possible interpretations, designing the study with reference to known hypotheses. Also, our conclusions were drawn from statistical analysis of our dataset and analysis of the source code of the eight ALLURA releases we could get grips onto. Further data triangulation might improve conclusion validity.

## 6 | RELATED WORK

While the majority of research on OSS projects has focused on their successes, a number of papers have investigated reasons for their failure. For example, Coelho and Valente [11] report on the results of a survey of 104 developers of failed (deprecated) GITHUB projects. As a result of this survey they provide 9 reasons for failure, such as: "usurped by competitor" (27 projects), "project is obsolete" (20 projects), or "lack of time of the main contributor" (18 projects). Khondhu et al. [28] analyzed a set of SOURCEFORGE projects and classified them into "active", "dormant", and "inactive". They then analyzed the "maintainability index" (MI) [37] of the code form a sample of each of those sets of projects. Their results, however, were inconclusive; there was not clear trend in the MI between the different sets of projects, and the sample sizes were not large. Yee et al. [30] conducted a study where they identified "five determinants of OSS success and the relationships among them through a literature review of previous IS success models". These factors were: Software quality, OSS use, Community service quality, User satisfaction, and Individual net benefits. Their key finding was that "usage of OSS is predominantly determined by user satisfaction and software quality".

There has also been a substantial literature on software project failures: while this research did not specifically focus on OSS projects, a complete overview is available in [26]. A number of previous works have investigated factors likely to influence the success/failures of software projects. For instance, Capiluppi et al. [7] investigated four dimensions, i.e., (i) community of developers, (ii) community of users, (iii) modularity and documentation, and (iv) software evolution, of 406 projects coming from a deprecated open source repository called FreshMeat. They found that most of the projects (57%) have one or two developers and that only a few of them (15%) can be considered active. Tourani et al. [50] investigate the impact of codes of conduct in open source projects, finding that they aim at providing a safe and inclusive community to avoid community-related issues. Ye and Kishida [57] conducted a study on the motivations leading developers to engage open source development, finding that learning is the major reason that motivates people to start contributing to open source projects. However, when this need is not satisfied, contributors tend to quit the project, possibly creating critical issues for its success. As a matter of fact, Avelino et al. [3] found that the survival of popular GITHUB projects heavily depend on one or two developers.

# 7 | CONCLUSIONS

This paper tells the cautionary tale of the downfall of ALLURA/SOURCEFORGE. We identified both internal and external causes, both social and technical in nature. We also found evidence of socio-technical and architectural problems in project archives. Project and community health metrics could have been used to predict and understand the downfall by tracking its early manifestations. The team reported to be aware of the technical

debt in the project, yet this was insufficient to maintain the health of the project. Community managers should also measure and manage the mutual impacts of technical and social debt. Furthermore, our analysis suggests that a combination of managerial and technical flaws doomed this project: architectural flaws, power-distance, and measurable management mishaps eventually led to the forge's demise. Practitioners and maintainers of other software projects can benefit from this cautionary tale, e.g., as indications of **what not to do**, in the scope of their community and architecture management infrastructure.

# ACKNOWLEDGMENTS

This research is partially supported by the European Commission grant no. 825480 (H2020), SODALITE, and by the grant no. 825040 (H2020), RADON; finally, by grant ANPCyT PICT-1725-2017. Palomba gratefully acknowledges the support of the Swiss National Science Foundation through the SNF Project No. PZ00P2\_186090 (TED).

#### References

- [1] Magnoni, Simone. An approach to measure community smells in software development communities. 2016.
- [2] Rakesh Agrawal and Ramakrishnan Srikant. Mining sequential patterns. In <u>Data Engineering</u>, <u>1995</u>. Proceedings of the Eleventh International Conference on, pages 3–14. IEEE, 1995.
- [3] Guilherme Avelino, Leonardo Passos, Andre Hora, and Marco Tulio Valente. A novel approach for estimating truck factors. In <u>Program</u> Comprehension (ICPC), 2016 IEEE 24th International Conference on, pages 1–10. IEEE, 2016.
- [4] Christian Bird, Nachiappan Nagappan, Harald C. Gall, Brendan Murphy, and Premkumar T. Devanbu. Putting it all together: Using sociotechnical networks to predict failures. In <u>ISSRE</u>, pages 109–119. IEEE Computer Society, 2009.
- [5] The Apache Software Foundation Blog. The apache software foundation announces apache allura as a top-level project. <a href="https://tinyurl.com/gjq8489">https://tinyurl.com/gjq8489</a>, 2014.
- [6] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. Qualitative research in psychology, 3(2):77-101, 2006.
- [7] Andrea Capiluppi, Patricia Lago, and Maurizio Morisio. Characteristics of open source projects. In <u>Software Maintenance and Reengineering</u>, 2003. Proceedings. Seventh European Conference on, pages 317–327. IEEE, 2003.
- [8] Marcelo Cataldo, James D. Herbsleb, and Kathleen M. Carley. Socio-technical congruence: a framework for assessing the impact of technical and work dependencies on software development productivity. In H. Dieter Rombach, Sebastian G. Elbaum, and Jï£;rgen Mï£;nch, editors, ESEM, pages 2–11. ACM, 2008.
- [9] Alexander Chatzigeorgiou and Anastasios Manakos. Investigating the evolution of bad smells in object-oriented code. In <u>2010 Seventh</u> International Conference on the Quality of Information and Communications Technology, pages 106–115. IEEE, 2010.
- [10] Zhifei Chen, Lin Chen, Wanwangying Ma, and Baowen Xu. Detecting code smells in python programs. In <u>Software Analysis, Testing and</u> Evolution (SATE), International Conference on, pages 18–23. IEEE, 2016.
- [11] Jailton Coelho and Marco Tulio Valente. Why modern open source projects fail. In Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, pages 186–196. ACM, 2017.
- [12] Valerio Cosentino, Javier Luis CÃanovas Izquierdo, and Jordi Cabot. Assessing the bus factor of git repositories. In <u>SANER</u>, pages 499–503. IEEE Computer Society, 2015.
- [13] David Crandall, Dan Cosley, Daniel Huttenlocher, Jon Kleinberg, and Siddharth Suri. Feedback effects between similarity and social influence in online communities. In <u>Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining</u>, KDD '08, pages 160–168, New York, NY, USA, 2008. ACM.
- [14] Paul A David and Francesco Rullani. Dynamics of innovation in an open source collaboration environment: lurking, laboring, and launching floss projects on sourceforge. <u>Industrial and Corporate Change</u>, 17(4):647–710, 2008.

- [15] Rohan Dutta, David K. Levine, and Salvatore Modica. Damned if you do and damned if you don't: Two masters. J. Economic Theory, 177:101– 125, 2018.
- [16] Qiong Feng, Rick Kazman, Yuanfang Cai, Ran Mo, and Lu Xiao. An architecture-centric approach to security analysis. In <u>Proceedings of the</u> 15th International Conference on Software Architecture, May 2016.
- [17] Eduardo Fernandes, Johnatan Oliveira, Gustavo Vale, Thanis Paiva, and Eduardo Figueiredo. A review-based comparative study of bad smell detection tools. In <u>nternational Conference on Evaluation and Assessment in Software Engineering</u>, EASE '16, pages 18:1–18:12, New York, NY, USA, 2016. ACM.
- [18] Brian Fitzgerald and Joseph Feller. A further investigation of open source software: community, co-ordination, code quality and security issues. Information Systems Journal, 12(1):3–6, 2002.
- [19] Martin Fowler. Refactoring: Improving the Design of Existing Code. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [20] Roberto Galoppini. Today we offer devshare (beta), a sustainable way to fund open source software. https://sourceforge.net/blog/today-weoffer-devshare-beta-a-sustainable-way-to-fund-open-source-software/, 2013.
- [21] Christos Gkantsidis, Milena Mihail, and Ellen W. Zegura. The markov chain simulation method for generating connected power law random graphs. In Richard E. Ladner, editor, ALENEX, pages 16–25. SIAM, 2003.
- [22] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. Econometrica: Journal of the Econometric Society, pages 424–438, 1969.
- [23] Bernie Hogan, Juan Antonio Carrasco, and Barry Wellman. Visualizing personal networks: Working with participant-aided sociograms. <u>Field</u> Methods, 19(2):116–144, 2007.
- [24] Florian Hoos, William Messier, Jason Smith, and Paulette Tandy. The effects of serving two masters and using the internal audit function as a management training ground on internal auditors' objectivity. 2014.
- [25] Qian Huang, Hefu Liu, and Xuepan Zhong. The impact of transactive memory systems on team performance. <u>IT & People</u>, 26(2):191–212, 2013.
- [26] Watts S Humphrey. Why big software projects fail: The 12 key questions, 2005.
- [27] Mitchell Joblin, Wolfgang Mauerer, Sven Apel, Janet Siegmund, and Dirk Riehle. From developer networks to verified communities: A finegrained approach. In Antonia Bertolino, Gerardo Canfora, and Sebastian G. Elbaum, editors, <u>ICSE (1)</u>, pages 563–573. IEEE Computer Society, 2015.
- [28] Jymit Khondhu, Andrea Capiluppi, and Klaas-Jan Stol. Is it all lost? a study of inactive open source projects. In <u>IFIP International Conference</u> on Open Source Systems, pages 61–79. Springer, 2013.
- [29] Ned Kock, David E. Avison, Richard L. Baskerville, Michael D. Myers, and A. T. Wood-Harper. Is action research: can we serve two masters? (panel session). In Prabuddha De and Janice I. DeGross, editors, <u>ICIS</u>, pages 582–585. Association for Information Systems, 1999.
- [30] Sang-Yong Tom Lee, Hee-Woong Kim, and Sumeet Gupta. Measuring open source software success. Omega, 37(2):426-438, 2009.
- [31] Meir M. Lehman, Dewayne E. Perry, and Juan F. Ramil. On evidence supporting the feast hypothesis and the laws of software evolution. IEEE Press, 1998.
- [32] Ying Liu, Eleni Stroulia, and Hakan Erdogmus. Understanding the open-source software development process: a case study with CVSChecker. In Proc. Intl'I Conf. on Open Source Systems, pages 154–161. NRC 47453, 2005.
- [33] Yuanfang Cai Lu Xiao and Rick Kazman. Design rule spaces: A new form of architecture insight. In <u>Proceedings of the 38th International</u> Conference on Software Engineering, 2014.
- [34] Tom Mens, Maëlick Claes, and Philippe Grosjean. Ecos: Ecological studies of open source software ecosystems. In <u>2014 Software Evolution</u> Week-IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE), pages 403–406. IEEE, 2014.

- [35] Ran Mo, Yuanfang Cai, Rick Kazman, and Lu Xiao. Hotspot patterns: The formal definition and automatic detection of architecture smells. In Proceedings of the 15th International Conference on Software Architecture, May 2015.
- [36] M.E.J. Newman. Fast algorithm for detecting community structure in networks. Physical Review E, 69, September 2003.
- [37] Paul Oman and Jack Hagemeister. Metrics for assessing a software system's maintainability. In <u>Software Maintenance, 1992. Proceerdings.</u>, Conference on, pages 337–344. IEEE, 1992.
- [38] Fabio Palomba, Gabriele Bavota, Massimiliano Di Penta, Fausto Fasano, Rocco Oliveto, and Andrea De Lucia. On the diffuseness and the impact on maintainability of code smells: a large scale empirical investigation. Empirical Software Engineering, pages 1–34, 2017.
- [39] Fabio Palomba, Damian Andrew Tamburri, Francesca Arcelli Fontana, Rocco Oliveto, Andy Zaidman, and A Serebrenik. Beyond technical aspects: How do community smells influence the intensity of code smells? <u>IEEE Transactions on Software Engineering</u>, 2018.
- [40] Fabio Palomba, Marco Zanoni, Francesca Arcelli Fontana, Andrea De Lucia, and Rocco Oliveto. Smells like teen spirit: Improving bug prediction performance using the intensity of code smells. In ICSME, pages 244–255. IEEE Computer Society, 2016.
- [41] Martin Pinzger, Nachiappan Nagappan, and Brendan Murphy. Can developer social networks predict failures? In <u>FSE '08: Proceedings of the</u> 16th ACM SIGSOFT international symposium on Foundations of software engineering, 2008.
- [42] Aniket Potdar and Emad Shihab. An exploratory study on self-admitted technical debt. In <u>Software Maintenance and Evolution (ICSME), 2014</u> IEEE International Conference on, pages 91–100. IEEE, 2014.
- [43] Alexandre Prokoudine.
   Anatomy of sourceforge/gimp controversy.
   http://libregraphicsworld.org/blog/entry/

   anatomy-of-sourceforge-gimp-controversy, 2015.
   http://libregraphicsworld.org/blog/entry/
- [44] D. A. A. Tamburri, F. Palomba, and R. Kazman. Exploring community smells in open-source: An automated approach. <u>IEEE Transactions on</u> Software Engineering, pages 1–1, 2019.
- [45] Damian A. Tamburri, Rick Kazman, and Hamed Fahimi. The architect's role in community shepherding. IEEE Software, 33(6):70-79, 2016.
- [46] Damian A. Tamburri, Patricia Lago, and Hans van Vliet. Uncovering latent social communities in software development. <u>IEEE Software</u>, 30(1):29 – 36, jan.-feb. 2013.
- [47] Damian A. Tamburri, Patricia Lago, and Hans van Vliet. Organizational social structures for software engineering. <u>ACM Comput. Surv.</u>, 46(1):3:1–3:35, July 2013.
- [48] Damian Andrew Tamburri, Philippe Kruchten, Patricia Lago, and Hans van Vliet. Social debt in software engineering: insights from industry. <u>J.</u> Internet Services and Applications, 6(1):10:1–10:17, 2015.
- [49] Damian Andrew Tamburri and Elisabetta Di Nitto. When software architecture leads to social debt. In Len Bass, Patricia Lago, and Philippe Kruchten, editors, WICSA, pages 61–64. IEEE Computer Society, 2015.
- [50] Parastou Tourani, Bram Adams, and Alexander Serebrenik. Code of conduct in open source projects. In <u>Software Analysis, Evolution and</u> Reengineering (SANER), 2017 IEEE 24th International Conference on, pages 24–33. IEEE, 2017.
- [51] Giuseppe Valetto, Mary Helander, Kate Ehrlich, Sunita Chulani, Mark Wegman, and Clay Williams. Using software repositories to investigate socio-technical congruence in development projects. <u>Mining Software Repositories, International Workshop on</u>, 0:25, 2007. doi = http://doi.ieeecomputersociety.org/10.1109/MSR.2007.33, publisher = IEEE Computer Society, address = Los Alamitos, CA, USA.
- [52] Matthew Van Antwerp and Greg Madey. Advances in the sourceforge research data archive. In <u>Workshop on Public Data about Software</u> Development (WoPDaSD) at The 4th International Conference on Open Source Systems, Milan, Italy, 2008.
- [53] Stanley Wasserman and Katherine Faust. <u>Social Network Analysis: Methods and Applications</u>. Number 8 in Structural analysis in the social sciences. Cambridge University Press, 1 edition, 1994.
- [54] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Bjöorn Regnell, and Anders Wesslén. Experimentation in software engineering: an introduction. Kluwer Academic Publishers, Norwell, MA, USA, 2000.

- [55] Lu Xiao, Yuanfang Cai, Rick Kazman, Ran Mo, and Qiong Feng. Identifying and quantifying architectural debt. In <u>Proceedings of the 38th</u> International Conference on Software Engineering, 2016.
- [56] Jin Xu, Scott Christley, and Gregory Madey. Application of social network analysis to the study of open source software. In <u>The economics of</u> open source software development, pages 247–269. Elsevier, 2006.
- [57] Yunwen Ye and Kouichi Kishida. Toward an understanding of the motivation open source software developers. In <u>Proceedings of the 25th</u> international conference on software engineering, pages 419–429. IEEE Computer Society, 2003.