

Continuous integration in a social-coding world: Empirical evidence from GITHUB

Updated version with corrections

Bogdan Vasilescu, Stef van Schuylenburg, Jules Wulms, Alexander Serebrenik, Mark G. J. van den Brand
Eindhoven University of Technology,
Den Dolech 2, P.O. Box 513,
5600 MB Eindhoven, The Netherlands
{b.n.vasilescu, a.serebrenik, m.g.j.v.d.brand}@tue.nl, {s.b.v.schuylenburg, j.j.h.m.wulms}@student.tue.nl

Abstract—Continuous integration is a software engineering practice of frequently merging all developer working copies with a shared main branch, e.g., several times a day.

With the advent of GITHUB, a platform well known for its “social coding” features that aid collaboration and sharing, and currently the largest code host in the open source world, collaborative software development has never been more prominent. In GITHUB development one can distinguish between two types of developer contributions to a project: *direct* ones, coming from a typically small group of developers with write access to the main project repository, and *indirect* ones, coming from developers who fork the main repository, update their copies locally, and submit pull requests for review and merger.

In this paper we explore how GITHUB developers use continuous integration as well as whether the contribution type (direct versus indirect) and different project characteristics (e.g., main programming language, or project age) are associated with the success of the automatic builds.

I. INTRODUCTION

Continuous integration (CI) is a software engineering practice of frequently merging all developer working copies with a shared main branch [1], e.g., several times a day, or with every commit. CI is typically supported by build servers that verify each integration automatically, running unit tests and reporting the results back to the developers. The concept, often attributed to Martin Fowler based on a 2000 blog entry [2] but known as “synch-and-stabilize” or “nightly build” already in 1997 [3], is a recommended best practice of agile software development methods like eXtreme Programming [4]. This continuous application of quality control checks aims to speed up the development process and to ultimately improve software quality, by reducing the integration problems occurring between team members that develop software collaboratively [1].

CI as a quality control mechanism in distributed, collaborative contexts is common to both commercial (e.g., Microsoft [5], [6]) and open source software (OSS) development (e.g., FreeBSD and Mozilla [7]). However, CI is especially relevant to OSS projects, due to the difficulties typically associated with imposing structured processes in such projects and on their contributors [7]: requirements documents are often lacking [8] and project contributors are often volunteers [9], typically geographically distributed [10], and rarely motivated by working in a team [11].

With the advent of social media (OSS) software development, recent years have witnessed many changes to how software is developed, and how developers collaborate, communicate, and learn [12]–[14]. One such prominent change is the emergence of the pull-based development model [15], [16], made possible by the distributed version control system Git, and made popular by the “social coding” platform GITHUB, currently the largest code host in the OSS world. In this model one can distinguish between two types of developer contributions to a project: *direct* ones, coming from a typically small group of developers with write access to the main project repository, and *indirect* ones, coming from developers who fork the main repository, update their copies locally, and submit pull requests for review and merger.

GITHUB’s implementation of the pull-based development model enables anyone with an account to submit changes to any repository with only a few clicks. This represents an unprecedented low barrier to entry for potential contributors, but it also impacts testing behavior [12], [17]. For example, GITHUB project owners interviewed by Pham et al. [17] reported scalability challenges when integrating (many) outside contributions, driving them towards automated tests. Automated CI services, such as TRAVIS-CI¹—integrated with GITHUB itself—or JENKINS², facilitate this process by automating a number of steps: whenever a commit is recorded or a pull request is received, the contribution is merged automatically into a testing branch, the existing test suite is run, and the contribution author and project owner are notified of the results. However, despite these potential benefits, CI services are reportedly underused both on GITHUB [17] as well as in OSS in general [18].

In this paper we focus on TRAVIS-CI, arguably the most popular CI service on GITHUB.³ Specifically, we quantitatively explore to what extent GITHUB developers use the TRAVIS-CI service, as well as whether the contribution type (direct versus indirect) or project characteristics (e.g., main

¹<https://travis-ci.com>

²<http://jenkins-ci.org>

³See, e.g., the blog entries <https://blog.codecentric.de/en/2012/05/travis-ci-or-how-continuous-integration-will-become-fun-again/> and <https://blog.futurice.com/tech-pick-of-the-week-travis-ci>, acc. June 2014

programming language, or project age) are associated with the success of the automatic builds.

The remainder of this paper is organised as follows. After discussing the methodology followed in Section II, we present our results in Section III. Then, we review threats to validity in Section IV, and finally we conclude in Section VI.

II. METHODOLOGY

To understand usage of the TRAVIS-CI service in GITHUB projects, we extracted and integrated data from two repositories: (i) GHTORRENT [19], [20], a service collecting and making available metadata for all public projects available on GITHUB; and (ii) the TRAVIS-CI API⁴. This section describes how we collected and analysed these data.

A. Sample Selection

Due to limitations of querying the TRAVIS-CI API, we restricted our attention in this study to a sample of large and active GITHUB projects. Using the GHTORRENT web interface⁵, we selected all GITHUB repositories that: (i) are not forks of other repositories; (ii) have not been deleted; (iii) are at least one year old; (iv) receive both commits and pull requests; (v) have been developed in Java, Python or Ruby; (vi) had at least 10 changes (commits or pull requests) during the last month; and (vii) have at least 10 contributors. We choose projects that receive both commits and pull requests, since we want to understand whether the way modification has been submitted (commit or pull request) can be associated with the build success. Our choice of the programming languages has been motivated by the history of TRAVIS-CI: TRAVIS-CI started as a service to the Ruby community in early 2011, while support for Java and Python has been announced one year later (February 21, 2012 and February 27, 2012, respectively). We expect therefore the use of TRAVIS-CI to be more widespread for Ruby than for Java and Python.

The data were extracted on March 30, 2014.

After filtering our sample contained 223 GITHUB projects, relatively balanced across the three programming languages (Figure 1): 70 (31.4%) were coded in Java, 83 (37.2%) in Python, and 70 (31.4%) in Ruby. The sample includes many large and popular OSS projects, such as rails, ruby, elasticsearch, or gradle.

B. Data Integration

To extract data about the automatic builds, we started by querying the `repos` endpoint of the TRAVIS-CI JSON API (using the repository slug—`username/repo`—as argument), to determine whether TRAVIS-CI is configured for a particular project. Then, if the response was not empty, we iteratively queried the builds associated with this project (25

at a time as per the TRAVIS-CI API) from the `builds` endpoint, collecting the `event_type` fields (that distinguish pull requests from pushes) and the `result` fields (that specify whether the build succeeded—0, or failed—1). Ongoing builds, for which the `result` fields are not set, were ignored.

C. Statistical Analysis

We aggregated the data collected from the TRAVIS-CI API into contingency tables, one for each GITHUB project, with rows corresponding to commits and pull requests, and columns—to passed and failed builds. Then, to test whether the success/failure of the build is independent from the way the modification has been proposed, we applied the χ^2 test of independence. Next, to formalise the strength of this dependence, we calculated the odds ratios and corresponding p -values. Finally, to aggregate the results of the χ^2 tests (one per project), we applied the Stouffer test using the weighted Z-score method [21], [22]. This allows us to lift the results of the individual χ^2 tests to the group level.

III. RESULTS

A. Direct Versus Indirect Contributions

We start by investigating the preference for direct (pushes) and indirect (pull requests) contributions among the projects in our sample. Java projects used the fewest pull requests during the observation month (March 2014), with a maximum of 26. Among the Java projects are also the most projects that do not use pull requests at all out of the 3 languages. Python and Ruby projects both have higher counts of pull requests, with maxima of 97 and 236, respectively. Some Python and Ruby projects even have more pull requests than commits.

The shared repository model (with contributors having write access to the repository) is more popular among Java projects, while Python and Ruby projects have more contributors submitting pull requests. Overall, we see that commits (direct code modifications) are more popular than pull requests (indirect code modifications), with only a small number of projects having more pull requests than commits. Similar findings have been reported by Gousios, Pinzger, and van Deursen in their exploratory study of the pull-based software development model on GITHUB [16].

Direct code modifications (pushed commits) are more popular than indirect code modifications (pull requests).

B. Usage of TRAVIS-CI

Next we investigate usage of TRAVIS-CI among the projects in our sample. First, we observe that an overwhelming majority of the projects are configured to use TRAVIS-CI (206 out of 223 projects, or 92.3%), confirming the anecdotal popularity of the CI service among GITHUB developers. However, slightly less than half of the 206 projects (93, or

⁴<http://docs.travis-ci.com/api/>

⁵Accessible from <http://ghtorrent.org/dblite/>




Figure 1. Distribution of the GITHUB projects considered per programming language: while many Ruby projects are configured for TRAVIS-CI and use it (dark green), most Java projects are configured for TRAVIS-CI but do not use it (middle green). Eleven Java projects are not configured for TRAVIS-CI (light green).

45%) have no associated builds recorded in the TRAVIS-CI database. This shows that while most projects *are ready* to use continuous integration, significantly fewer *actually do*. When we look into the distribution of these projects with no builds with respect to programming language, we observe that Java projects are overrepresented (69.5% of the Java projects configured to use TRAVIS-CI are not actually using it), while Ruby projects are underrepresented (Figure 1).

Although most GITHUB projects in our sample are configured to use the TRAVIS-CI continuous integration service, less than half actually do. In terms of languages, Ruby projects are among the early adopters of TRAVIS-CI, while Java projects are late to use continuous integration.

C. Contribution Type and Build Success

We have observed that the median success rate is 79.5% for commits and 75.9% for pull requests. To better understand whether TRAVIS-CI build failure is independent from the way the modification has been proposed, we focussed on projects with at least 5 failed/successful builds for each contribution type, as required by the χ^2 test of independence (cf. Section II-C). Out of 113 GITHUB projects configured to use TRAVIS-CI and actually using it (206 – 93, cf. the discussion in the previous subsection), 84 projects had sufficient data for the χ^2 test. Among the remaining 29 projects to which the χ^2 test could not be applied, in most cases it was the failed pull requests cell that had insufficient data. In other words, builds fail less frequently when contributions are submitted via pull requests. We believe this is because when a developer does not have commit rights and she suggests a change via a pull request, she will try harder to make sure the change is valid and it will not break the build. However, when instead a developer has commit rights,




Figure 2. Beanplot [23] of the distribution of χ^2 p -values (84 GITHUB projects). The thick vertical line represents the median value.




Figure 3. Beanplot [23] of the distribution of ratios between the odds that commit builds succeed and the odds that pull request builds succeed (45 GITHUB projects with statistically significant χ^2 test results; $p < 0.05$). All odds ratios have $p < 0.05$. The thick vertical line shows the median.

she can try out new things more freely, since she also has the power to reverse the change, if necessary.

Finally, we lift the results of individual χ^2 tests to the group level by applying the Stouffer procedure to obtain a combined significance level, i.e., “overall probability in favor of the outcome of the majority of the studies” [24]. The Stouffer test statistic Z was calculated as 18.34 and the corresponding p -value was too small to be calculated precisely (see Figure 2 for the distribution of p -values of the individual χ^2 tests). This implies that taken together, the data indicates dependence of the build success on the way the modification has been proposed.

To investigate the directionality of this dependence we compute the odds ratios (i.e., the ratios between the odds that commit builds succeed and the odds that pull request builds succeed) for the GITHUB projects that showed statistically significant χ^2 test results at 95% confidence level. Figure 3 displays the distribution of the resulting odds ratios (all statistically significant at 95% confidence level). Inspection of Figure 3 reveals that for the overwhelming majority of projects (39 out of 45, or 87%), builds corresponding to push commits are much more likely to succeed than builds corresponding to pull requests (since the odds ratios are greater than 1). This suggests that pull request evaluation is a complex process; pull requests are prone to integration testing failures, and other mechanisms are needed to ensure quality control, e.g., discussions and code review [16], [25].

Pull requests are more likely to result in integration testing failures than push commits.

D. Impact of Project Differences on the Contribution Type/Build Success Relation

The 84 projects subjected to the χ^2 test in Section III-C have been developed in different languages; have different ages; and involve different numbers of contributors. Table I summarizes differences between those languages, ages, and numbers of contributors in terms of rejecting the null-hypotheses of the χ^2 test, i.e., independence of the build success from the way the modification has been proposed (95% confidence level). The thresholds of 17 and 33 contributors correspond to the 33% and 67% percentiles. Performing Stouffer tests for each group led to very small p -values, indicating that results can be lifted to the group level. The data suggests that null hypotheses can be rejected (✓) for Python and Ruby projects, but not (✗) for Java projects; can be rejected for older projects but not for younger ones; and can be rejected for projects with not too many contributors as opposed to projects with many contributors.

Next we conducted odds ratio tests for projects where the null hypothesis has been rejected (✗) for the group level: **all odds ratio tests turned out to be statistically significant (p -values never exceeded 0.05) and in almost all cases the odds ratios exceeded 1. This means that whenever build success depends on the way the modification has been performed, pull requests are more likely to result in integration testing failures than direct commits.**

For Python and Ruby projects, projects older than two years, and projects with not too many contributors, **pull requests are more likely to result in integration testing failures than direct commits.** No such impact of the way the modification has been performed can be observed for Java projects, projects younger than two years, and projects with many contributors.

IV. THREATS TO VALIDITY

In this section, we discuss the threats to construct validity, internal validity and external validity [26].

Construct validity assesses whether the variables we considered accurately model our hypotheses. One of the threats to construct validity pertains to pull requests that appear as not being merged even if they have been merged [27].

Internal validity means that changes in dependent variables can be attributed to changes in independent variables

	Prog. lang.			Age (years)			Contributors		
	Java	Python	Ruby	<2	2-4	>4	≤17	17-33	>33
# projects	10	34	40	24	42	18	29	27	28
... s.t. $p < 0.05$	3	19	23	9	25	11	18	15	12
H_0	✗	✓	✓	✗	✓	✓	✓	✓	✗
%odds ratio > 1	n/a	89	87	n/a	92	82	89	80	n/a

Table I

COMPARISON OF SUBGROUPS OF 84 GITHUB PROJECTS BASED ON THE PROGRAMMING LANGUAGE, AGE AND THE NUMBER OF CONTRIBUTORS.

instead of to something else. Specifically, we considered if the number of successful and failed builds are related to the type of contribution, the main programming language, the project age, or the number of contributors. We did not consider if other variables can confound this relationship. If such variables exist, they may invalidate our results.

External validity means that the results we found can be generalized to real-world settings. Since we acquired a large set of data from a general-purpose website like GITHUB, we feel that our results can be generalized to large and active open-source projects developed in Java, Python or Ruby and using TRAVIS-CI. Closed-source projects, small projects, projects in very different programming languages or using different CI services, such as JENKINS, may not show the same patterns.

V. FUTURE WORK

We plan to triangulate our quantitative findings through qualitative analysis, such as interviews and questionnaires. Conducting surveys can allow us to obtain insights in the ways continuous integration is used in open-source and proprietary development. A complementary approach will consist in performing a more detailed analysis of the TRAVIS-CI configuration files. Finally, we plan to consider a larger sample of GITHUB projects, including those developed in additional programming languages.

VI. CONCLUSIONS

In this paper we have studied a sample of large and active GITHUB projects developed in Java, Python and Ruby. We started by observing that direct code modifications (commits) are more popular than indirect code modifications (pull requests). Next, we have investigated the use of TRAVIS-CI: although most GITHUB projects in our sample are configured to use the TRAVIS-CI continuous integration service, less than half actually do. In terms of languages, Ruby projects are among the early adopters of TRAVIS-CI, while Java projects are late to use continuous integration. Next, for those projects that actually use TRAVIS-CI, we have studied whether the success or failure of a build is independent on the way code modification has been proposed. **Our overall conclusion is that success or failure of a build does depend on the way the code modification has been proposed: pull requests are more likely to result in integration testing failures than direct commits.** However, we observe differences for projects developed in different programming languages, of different ages, and involving different numbers of contributors.

ACKNOWLEDGEMENTS

Special thanks to Mathias Meyer and the Travis CI team for helping us query their API. Bogdan Vasilescu gratefully acknowledges support from the Dutch Science Foundation (NWO) through the NWO 600.065.120.10N235 project.

REFERENCES

- [1] P. M. Duvall, S. Matyas, and A. Glover, *Continuous integration: improving software quality and reducing risk*. Pearson Education, 2007.
- [2] M. Fowler, “Continuous integration,” <http://martinfowler.com/articles/continuousIntegration.html>, 2006, accessed June 2014.
- [3] M. A. Cusumano and R. W. Selby, “How microsoft builds software,” *Commun. ACM*, vol. 40, no. 6, pp. 53–61, Jun. 1997.
- [4] K. Beck, *Extreme programming explained: embrace change*. Addison-Wesley, 2000.
- [5] M. A. Cusumano, “Extreme programming compared with Microsoft-style iterative development,” *Communications of the ACM*, vol. 50, no. 10, pp. 15–18, 2007.
- [6] A. Miller, “A hundred days of continuous integration,” in *Agile*. IEEE, 2008, pp. 289–293.
- [7] J. Holck and N. Jørgensen, “Continuous integration and quality assurance: A case study of two open source projects,” *Australasian Journal of Information Systems*, vol. 11, no. 1, 2007.
- [8] A. Mockus, R. T. Fielding, and J. D. Herbsleb, “Two case studies of open source software development: Apache and Mozilla,” *TOSEM*, vol. 11, no. 3, pp. 309–346, 2002.
- [9] A. Hars and S. Ou, “Working for free? motivations of participating in open source projects,” in *HICSS*. IEEE, 2001, pp. 1–9.
- [10] B. J. Dempsey, D. Weiss, P. Jones, and J. Greenberg, “Who is an open source software developer?” *Communications of the ACM*, vol. 45, no. 2, pp. 67–72, 2002.
- [11] K. R. Lakhani and R. G. Wolf, “Why hackers do what they do: Understanding motivation and effort in free/open source software projects,” *Perspectives on Free and Open Source Software*, vol. 1, pp. 3–22, 2005.
- [12] L. A. Dabbish, H. C. Stuart, J. Tsay, and J. D. Herbsleb, “Social coding in GitHub: transparency and collaboration in an open software repository,” in *CSCW*. ACM, 2012, pp. 1277–1286.
- [13] A. Begel, J. Bosch, and M.-A. D. Storey, “Social networking meets software development: Perspectives from GitHub, MSDN, Stack Exchange, and TopCoder,” *IEEE Software*, vol. 30, no. 1, pp. 52–66, 2013.
- [14] B. Vasilescu, A. Serebrenik, P. T. Devanbu, and V. Filkov, “How social Q&A sites are changing knowledge sharing in open source software communities,” in *CSCW*. ACM, 2014, pp. 342–354.
- [15] E. T. Barr, C. Bird, P. C. Rigby, A. Hindle, D. M. German, and P. Devanbu, “Cohesive and isolated development with branches,” in *FASE*. Springer, 2012, pp. 316–331.
- [16] G. Gousios, M. Pinzger, and A. van Deursen, “An exploratory study of the pull-based software development model,” in *ICSE*, 2014, pp. 345–355.
- [17] R. Pham, L. Singer, O. Liskin, K. Schneider *et al.*, “Creating a shared understanding of testing culture on a social coding site,” in *ICSE*. IEEE, 2013, pp. 112–121.
- [18] A. Deshpande and D. Riehle, “Continuous integration in open source software development,” in *Open Source Development, Communities and Quality*. Springer, 2008, pp. 273–280.
- [19] G. Gousios, “The GHTorrent dataset and tool suite,” in *MSR*. IEEE, 2013, pp. 233–236.
- [20] G. Gousios, B. Vasilescu, A. Serebrenik, and A. Zaidman, “Lean GHTorrent: GitHub data on demand,” in *MSR*, 2014, pp. 384–387.
- [21] M. Whitlock, “Combining probability from independent tests: the weighted Z-method is superior to Fisher’s approach,” *Journal of Evolutionary biology*, vol. 18, no. 5, pp. 1368–1373, 2005.
- [22] S. A. Stouffer, E. A. Suchman, L. C. DeVinney, S. A. Star, and R. M. Williams Jr., *The American soldier*, ser. Studies in social psychology in World War II. Princeton University Press, 1949, vol. 1, adjustment during army life.
- [23] P. Kampstra, “Beanplot: A boxplot alternative for visual comparison of distributions,” *Journal of Statistical Software, Code Snippets*, vol. 28, no. 1, pp. 1–9, 2008.
- [24] D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*, 4th ed. Chapman & Hall/CRC, 2007.
- [25] J. Tsay, L. Dabbish, and J. Herbsleb, “Let’s talk about it: Evaluating contributions through discussion in GitHub,” in *FSE*. ACM, 2014.
- [26] D. E. Perry, A. A. Porter, and L. G. Votta, “Empirical studies of software engineering: a roadmap,” in *ICSE*. ACM, 2000, pp. 345–355.
- [27] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. Germán, and D. Damian, “The promises and perils of mining GitHub,” in *MSR*, 2014, pp. 92–101.