

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and  
Information Systems

School of Computing and Information Systems

---

7-2023

### The impact of a continuous integration service on the delivery time of merged pull requests

João Helis BERNARDO

Daniel Alencar DA COSTA

Uirá KULESZA

Christoph TREUDE

Singapore Management University, [ctreude@smu.edu.sg](mailto:ctreude@smu.edu.sg)

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Programming Languages and Compilers Commons](#), and the [Software Engineering Commons](#)

---

#### Citation

BERNARDO, João Helis; DA COSTA, Daniel Alencar; KULESZA, Uirá; and TREUDE, Christoph. The impact of a continuous integration service on the delivery time of merged pull requests. (2023). *Empirical Software Engineering*. 28, (4), 1-57.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/8769](https://ink.library.smu.edu.sg/sis_research/8769)

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

## The Impact of a Continuous Integration Service on the Delivery Time of Merged Pull Requests

João Helis Bernardo · Daniel Alencar da Costa ·  
Uirá Kulesza · Christoph Treude

Received: date / Accepted: date

**Abstract** Continuous Integration (CI) is a software development practice that builds and tests software frequently (e.g., at every push). One main motivator to adopt CI is the potential to deliver software functionalities more quickly than not using CI. However, there is little empirical evidence to support that CI helps projects deliver software functionalities more quickly. Through the analysis of 162,653 pull requests (PRs) of 87 GitHub projects, we empirically study whether adopting a CI service (TRAVISCI) can quicken the time to deliver merged PRs. We complement our quantitative study by analyzing 450 survey responses from participants of 73 software projects. Our results reveal that adopting a CI service may not necessarily quicken the delivery of merge PRs. Instead, the pivotal benefit of a CI service is to improve the decision making on PR submissions, without compromising the quality or overloading the project's reviewers and maintainers. The automation provided by CI and the boost in developers' confidence are key advantages of adopting a CI service. Furthermore, open-source projects planning to attract and retain developers should consider the use

---

João Helis Bernardo  
Federal Institute of Rio Grande do Norte (IFRN)  
Federal University of Rio Grande do Norte (UFRN)  
Natal, Brazil  
E-mail: joao.helis@ifrn.edu.br

Daniel Alencar da Costa  
University of Otago  
Dunedin, New Zealand  
E-mail: danielcalencar@otago.ac.nz

Uirá Kulesza  
Federal University of Rio Grande do Norte (UFRN)  
Natal, Brazil  
E-mail: uira@dimap.ufrn.br

Christoph Treude  
University of Melbourne  
Melbourne, Australia  
E-mail: christoph.treude@unimelb.edu.au

of a CI service in their project, since CI is perceived to lower the contribution barrier while making contributors feel more confident and engaged in the project.

**Keywords** Continuous Integration · Pull Request · Delivery Time · Code Review

## 1 Introduction

Development teams are required to deliver software functionalities more quickly than ever to improve the time-to-market and success of their software projects (Debbiche et al., 2014). The quick delivery of functionalities may keep customers engaged with the project while providing valuable feedback. To improve the processes of software integration and packaging, Continuous Integration (CI) has been proposed as part of the Extreme Programming (XP) methodology (Beck, 2000), which claims that CI can provide more confidence for developers and quicken the delivery of software functionalities (Laukkanen et al., 2015).

Continuous Integration is a set of practices that enables development teams to integrate software more frequently (Fowler and Foemmel, 2006). The increased number of integrations is possible by using automated tools such as automated tests. Ideally, CI should automatically compile, test, and package the software whenever code modifications occur. Nowadays, several vendors (e.g., APACHE or GITHUB) have developed tools to provide *CI as a service* for developers to implement their CI pipelines. Examples of these tools are CLOUDBEES, GITHUB ACTIONS, and TRAVISCI. The main philosophy behind CI is that the software must always be in a working state, which is constantly put to test at each integration (Duvall et al., 2007). CI has been widely adopted by the software development community in both open-source and corporate software projects.

Existing research has analyzed the usage of CI in open-source projects hosted on GITHUB (Vasilescu et al., 2014, 2015; Hilton et al., 2016; Bernardo et al., 2018; Nery et al., 2019; Soares et al., 2022; Santos et al., 2022). For instance, Vasilescu et al. (2015) investigated the productivity and quality outcomes of projects that use CI services on GITHUB. They found that projects that use CI merge pull requests (PRs) more quickly if they are submitted by core developers. Also, core developers discover significantly more bugs when they use CI. Although existing research has demonstrated that CI may provide benefits for development teams, Soares et al. (2022) revealed that several studies investigate the benefits of using a *CI service* instead of studying the benefits of CI as a whole practice. For example, instead of checking whether studied projects adopt the full set of practices required by CI (Felidré et al., 2019), most studies have assumed that projects use CI solely because a CI service was employed (such as TRAVISCI). We as a community should be clear about such scenarios when reporting our results. According to Fowler and Foemmel (2006), adopting CI is not using a CI service alone but also adopting and maintaining specific development practices. To adopt CI appropriately, projects must maintain short build durations, fix broken builds as immediately as possible, check-in code frequently, and maintain high code coverage (Duvall et al., 2007; Felidré et al., 2019; Santos et al., 2022).

In this regard, a common claim about adopting CI is that projects are able to release more frequently (Ståhl and Bosch, 2014; Hilton et al., 2016), implying that software updates would be delivered more quickly to their end-users. However, there is no sufficient empirical evidence to show that CI can indeed be associated with a quicker delivery of software functionalities to end users. Studying whether CI can quicken the delivery of software functionalities is important because release delays are frustrating to end users (da Costa et al., 2014, 2016).

In our prior work (Bernardo et al., 2018), we quantitatively analyzed whether the use of a CI service (TRAVIS CI) is correlated with the time to deliver merged *Pull Requests* (PRs) of GitHub projects. Our study investigated 162,653 PRs from 87 GitHub projects, which were implemented in 5 programming languages.<sup>1</sup> We found that the time-to-deliver PRs is shorter *after* adopting TRAVIS CI in only 51.3% of the projects. As we have observed that the use of a CI service is not necessarily associated with a quicker delivery of pull requests, we designed a qualitative study to obtain deeper explanations for our results while deepening our understanding of the potential influence of CI *as a whole practice* on the time-to-market of merged PRs. For example, do developers believe that CI, as a whole, influence the delivery of PRs in their projects? We designed a qualitative study because we could not find answers to such questions in our previous quantitative analyses. To sum up, our qualitative study complements our previous study by providing more explanations and context for the results we observed in the previous study. Therefore, we survey 450 participants from 73 GitHub projects (out of the initial 87 projects of our quantitative study). Our qualitative analysis is composed of:

- Data collection from survey responses of 450 participants of 73 popular open-source projects from GitHub.
- An open-coding analysis of the answers to the open-ended questions of our survey using a thematic analysis technique.
- An analysis of the extent to which the survey participants are in accordance with the quantitative results of our prior work.

## 1.1 Quantitative Study

Our quantitative study addresses the following research questions:

- **RQ1: Are merged pull requests released more quickly using a CI service?** The wide adoption of CI is often motivated by the perceived benefits of this practice. For instance, higher confidence in the software product (Duvall et al., 2007), higher release frequency (Ståhl and Bosch, 2014), and the prospect of delivering software updates more quickly (Laukkanen et al., 2015). However, there is a lack of studies that empirically investigate the association between using a CI service and the time-to-deliver of merged PRs. In RQ1, we study the delivery time of merged PRs *before* and *after* the use of TRAVIS CI.

---

<sup>1</sup> <https://prdeliverydelay.github.io/#studied-projects>

- **RQ2: Does the increased number of PR submissions after adopting a CI service increase the delivery time of pull requests?** In RQ1, we find that only 51.3% of the projects deliver merged PRs more quickly *after* adopting TRAVISCI. This result contradicts the assumption that merged PRs would be delivered more quickly *after* the adoption of a CI service in most of the projects. We then ask the following question: is there another key factor influencing the delivery time of merged PRs *after* TRAVISCI is adopted, such as a significant increase in workload?
- **RQ3: What factors impact the delivery time after adopting a CI service?** In RQ1 and RQ2, we study the impact of adopting a CI service on the delivery time of merged PRs. Nevertheless, it is also important to understand what are the characteristics of the delivery time of merged PRs *before* and *after* the use of a CI service. Such information may help decision makers to track and avoid a high delivery time.

## 1.2 Qualitative Study

In our qualitative study, we address the following research questions:

- **RQ4: What is the perceived influence of CI on the time to deliver merged PRs?**  
In this RQ, we aim to deepen our understanding of how CI may impact the delivery time of PRs. We consult contributors of projects that use CI to obtain qualitative data, which can provide relevant insights to the research community and practitioners.
- **RQ5: What are the perceived causes of delay in the delivery time of merged PRs?**  
In RQ4 we observe that 42.9% of participants are skeptical regarding the impact of CI on the delivery time of merged PRs. Therefore, in this RQ, we further discuss *indirect factors* (i.e., factors that are not necessarily related to CI) that participants believe may also impact the delivery time of PRs.
- **RQ6: What is the perceived influence of CI on the software release process?**  
Given that in RQ2 we observe a substantial increase in the number of delivered PRs per release (*after* the adoption of TRAVISCI), we aim to obtain further insights as to why the increase in the number of delivered PRs occurs. For this purpose, we consult our participants regarding their perceived influence of CI on the release process of their project. For example, is it the case that, because CI encourages the constant packaging of the software, preparing a release is no longer a challenge?
- **RQ7: What is the perceived influence of CI on the code review process?** Intriguingly, in RQ1, we find that PRs are merged faster *before* the adoption of TRAVISCI in 73% (46/63) of the projects. This result motivates us to further investigate factors that influence the merge time when CI is adopted.
- **RQ8: What is the perceived influence of CI on attracting more contributors to open-source projects?** In RQ1 and RQ2, we observe that there exist a higher number of contributors and PR submissions *after* the adoption of TRAVISCI. In RQ8, we consult our participants to better understand whether CI has any influence on attracting more contributors to open-source projects.

**Paper organization.** The remainder of this paper is organized as follows. In Section 2, we discuss the related work. In Section 3, we explain the design of our

quantitative and qualitative studies. In Sections 4 and 5, we present the results of our quantitative and qualitative studies, respectively. We discuss the practical implications of our observations for the research and practice in software engineering in Section 6. In Section 7, we discuss the threats to the validity and limitations of our study. Finally, we draw conclusions in Section 8.

## 2 Related work

Through a systematic literature review of *Agile Release Engineering* practices, Karvonen et al. (2017) highlighted that empirical research in software engineering is crucial to better understand the impact of adopting CI on software development.

### *CI and team productivity*

The study by Hilton et al. (2016) revealed that 70% of the most popular GITHUB projects use CI. The authors identified that CI helps projects to release more often, whereas the CI build status may foster a faster integration of PRs. Vasilescu et al. (2015) studied the potential impact of CI on the quality and productivity of software projects. They found that projects that use CI merge PRs more quickly if these PRs are submitted by core developers. The authors found that core developers identify significantly more bugs when using CI. Regarding the acceptance and latency of PRs in CI (where latency is the time taken to merge a PR), Yu et al. (2016) found that the likelihood of rejecting a PR increases by 89.6% when the PR breaks the build. The results also show that the more succinct the PR is, the greater the probability that the PR is reviewed and merged earlier. Furthermore, Zhao et al. (2017) investigated the transition to TRAVISCI<sup>2</sup> in open-source projects. According to their study, the following changes may occur when TRAVISCI is adopted: (i) a small increase in the number of merged commits; (ii) a statistically significant decrease in the number of merge commit churn; (iii) a moderate increase in the number of closed issues; and (iv) a stationary behavior in the number of closed PRs.

In our study, we use an approach similar to Vasilescu et al. (2015) to identify projects that use TRAVISCI. Our goal with our quantitative study is to understand the association between TRAVISCI and the time taken for PRs to be delivered to end users. Furthermore, our qualitative study investigates how contributors of open-source projects perceive the impact of CI on the review and release processes of their projects. Our work is complementary to prior studies, contributing to a larger understanding of how CI can impact several development activities in software projects (i.e., code review and project release).

### *CI and code review*

Recent studies have investigated the impact of CI on code review. The study by Zampetti et al. (2019) found that PRs that generate successful builds have 1.5 more

---

<sup>2</sup> <https://travis-ci.org/>

chances of being merged. Our qualitative investigation corroborates the results of Zampetti et al. (2019), showing that the CI build status can influence the decisions of code reviewers. Furthermore, Cassee et al. (2020) found that the discussion held before the acceptance of a PR reduced considerably *after* CI was adopted. Conversely, the number of changes developers performed during code review remained roughly the same. The work of Zhang et al. (2022b) investigated the influence of various factors on PR latency. They found that, when using CI, the build status and duration are moderately relevant factors for accepting PRs. In a follow-up study, Zhang et al. (2022a) observed that CI assists the acceptance of PRs by automating the code review process and replacing part of the code inspection work, accelerating the review process. Indeed, our qualitative study reveals that, among the list of CI factors that impact the code review process, the most cited factors were related to an improvement in *automation* and *confidence*. The participants of our study highlighted that CI facilitates the understanding of code decisions, accelerating the code review process.

### ***Adherence to CI best practices***

Vasilescu et al. (2014) studied the use of TRAVISCI in a sample of 223 GITHUB projects. They found that the majority of projects (92.3%) are configured to use TRAVISCI but less than half actually use the CI service. Felidré et al. (2019) analyzed 1,270 open-source projects using TRAVISCI to understand the adherence of projects to the recommended CI practices. The authors observed that 748 (60%) projects perform infrequent check-ins. The study by Nery et al. (2019) studied the relationship between the use of CI and the evolution of software tests. The authors found that the overall test ratio and coverage of projects improved after CI was adopted. In our work, our participants mention that CI impacts the delivery time of merged PRs by improving *project quality*, *automation*, and the *release process*. According to our participants, CI improves the code quality and stability, making developers more confident to ship releases. The confidence in developers can be fostered by comprehensive automated testing, especially when the code coverage is high.

Gallaba and McIntosh (2018) studied 9,312 open-source projects using TRAVISCI to understand how projects are using or misusing the features of TRAVISCI. The authors found that the majority (48.16%) of TRAVISCI configurations is specifying job processing nodes. Furthermore, explicit deployment code is rare (2%), which indicates that developers rarely use TRAVISCI to implement Continuous Delivery. In our qualitative study, our participants often emphasize the relevance of automated tasks in CI to improve the project release process. *Automated tests* and *release automation* (i.e., Continuous Deployment) are frequently mentioned when our participants explain the influence of CI on project releases. However, in addition to many developers understanding CI as a Continuous Deployment enabler, such a feature is misused by many projects that use CI (Gallaba and McIntosh, 2018).

Table 1: Summary of the number of projects and released pull requests grouped by programming language.

Language	Projects	PRs total	PRs <i>before</i> CI	PRs <i>after</i> CI
JavaScript	33	57,104	17,556	39,548
Python	23	55,003	9,107	45,896
Java	11	7,700	3,433	4,267
Ruby	10	22,864	3,197	19,667
PHP	10	19,982	5,817	14,165
<b>Total</b>	<b>87</b>	<b>162,653</b>	<b>39,110</b>	<b>123,543</b>

### 3 Empirical Study Design

We perform two complementary studies: one quantitative and one qualitative. For each study, we describe their studied projects, their data collection process, and their methodology.

#### 3.1 Quantitative Study—Study I

In Study I, we divide our projects in two time periods, *before* and *after* the adoption of TRAVIS CI. Segmenting the data into these two time periods is necessary to study the association between the adoption of TRAVIS CI and the delivery time of PRs.

##### 3.1.1 Studied Projects

Our goal is to identify projects with substantial historical data that adopted TRAVIS CI eventually. We use such projects to better understand the potential influence of adopting a CI service on the delivery time of merged PRs. We use an approach similar to Vasilescu et al. (2015) and Hilton et al. (2016) to select projects that use TRAVIS CI. We use the date of the first build on TRAVIS CI to determine when TRAVIS CI was introduced in a project.

We selected a set of 87 popular GITHUB projects (33 JavaScript, 23 Python, 11 Java, 10 Ruby, and 10 PHP). We collect metrics related to the PRs and releases of each project. The detailed information about all computed metrics for each PR is described in Tables 2 and 3. We believe these metrics can be correlated with the *delivery time* of merged PRs. A total of 162,653 delivered PRs were collected (123,543 PRs were delivered *after* the adoption of TRAVIS CI, whereas 39,110 were delivered *before* the adoption of TRAVIS CI). The unbalanced number of PRs across time periods is a reflection of the duration of the adoption of TRAVIS CI in different projects. The median age of our projects is 5.1 years, where the use of TRAVIS CI accounts for 60.8% (3.1 years) of the age of our projects. Table 1 shows the number of PRs per programming language *before* and *after* the adoption of TRAVIS CI. Our project selection and data collection processes are explained in more detail in an earlier publication of this work (Bernardo et al., 2018).



Table 2: Metrics that are used in our explanatory models (resolver, pull request, and project dimensions).

Dimension	Attributes	Type	Definition (d) — Rationale (r)
Resolver	Contributor Experience	Numeric	<p><b>d:</b> The number of previously released PRs that were submitted by the contributor of a particular PR. We consider the author of the PR to be its contributor.</p> <p><b>r:</b> The greater the experience and participation of a user within a specific open-source project, the greater his/her chance of having his/her PR reviewed and integrated into the codebase of such a project by its core integrators (Shihab et al., 2010).</p>
	Contributor Integration	Numeric	<p><b>d:</b> The average in days of the previously released PRs that were submitted by a particular contributor.</p> <p><b>r:</b> If a particular contributor usually submits PRs that are merged and released quickly, his/her future PR might be merged and released quickly as well (da Costa et al., 2016).</p>
Pull Request	Stack Trace Attached	Boolean	<p><b>d:</b> We verify if the PR report has a stack trace attached in its description.</p> <p><b>r:</b> If the PR provides a bug fix, a stack trace attached may provide useful information regarding the causes of the bug and the importance of the submitted code, which may quicken the merge of the PR and its delivery in a release of the project (Schroter et al., 2010).</p>
	Description Size	Numeric	<p><b>d:</b> The number of characters in the body (description) of a PR.</p> <p><b>r:</b> PRs that are well described might be easier to merge and release than PRs that are more difficult to understand (da Costa et al., 2016).</p>
Project	Queue Rank	Numeric	<p><b>d:</b> The number that represents the moment when a PR is merged compared to other merged PRs in the release cycle. For example, in a queue that contains 100 PRs, the first merged PR has position 1, while the last merged PR has position 100.</p> <p><b>r:</b> A PR with a high <i>queue rank</i> is a recently merged PR. A merged PR might be released faster/slower depending on its queue position (da Costa et al., 2016).</p>
	Merge Workload	Numeric	<p><b>d:</b> The number of PRs that were created and still waiting to be merged by a core integrator at the moment at which a specific PR is submitted.</p> <p><b>r:</b> A PR might be released faster/slower depending of the amount of submitted PRs waiting to be merged. The higher the amount of created PRs waiting to be analyzed and merged, the greater the workload of the contributors to analyze these PRs, which may impact their delivery time.</p>

### 3.1.2 Research Approach

Figure 1 shows the basic life cycle of a delivered PR, where  $t_1$  is the merge phase and  $t_2$  is the delivery phase. We refer to  $t_1 + t_2$  as the lifetime of a PR. In *RQ1*, we analyze the *merge* and *delivery* phases. The merge phase ( $t_1$ ) is the required time for PRs to be

Table 3: Metrics that are used in our explanatory models (process dimension).

Dimension	Attributes	Type	Definition (d) — Rationale (r)
Process	Number of Impacted Files	Numeric	<b>d:</b> The number of files linked to a PR submission. <b>r:</b> The delivery time might be related to the high number of files of a PR, because more effort must be spent to integrate it (Jiang et al., 2013).
	Churn	Numeric	<b>d:</b> The number of added lines plus the number of deleted lines to a PR. <b>r:</b> A higher churn suggests that a great amount of work might be required to verify and integrate the code contribution sent by means of PR (Jiang et al., 2013; Nagappan and Ball, 2005).
	Merge Time	Numeric	<b>d:</b> Number of days between the submission and merge of a PR. <b>r:</b> If a PR is merged quickly, it is more likely to be released faster.
	Number of Activities	Numeric	<b>d:</b> An activity is an entry in the PR’s history. <b>r:</b> A high number of activities might indicate that much work was required to make the PR acceptable, which may impact the integration of such PR into a release (Jiang et al., 2013).
	Number of Comments	Numeric	<b>d:</b> The number of comments of a PR. <b>r:</b> A high number of comments might indicate the importance of a PR or the difficulty to understand it (Giger et al., 2010), which may impact its delivery time (Jiang et al., 2013).
	Interval of Comments	Numeric	<b>d:</b> The sum of the time intervals (days) between comments divided by the total number of comments of a PR. <b>r:</b> A short <i>interval of comments</i> indicates the discussion was held with priority, which suggests that the PR is important, thus, the PR might be delivered faster (da Costa et al., 2016).
	Commits per PR	Numeric	<b>d:</b> Number of commits per PR. <b>r:</b> The higher the number of commits in a PR, the greater the amount of contribution to be analyzed by the project integrators, which might impact the delivery time of the PR.

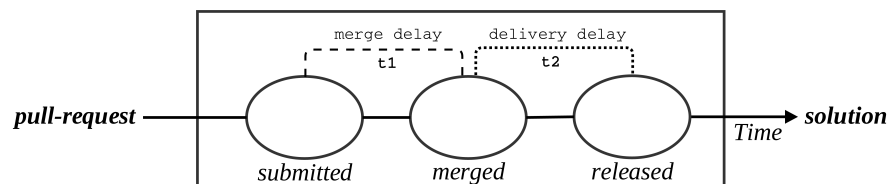


Fig. 1: The basic life-cycle of a delivered pull request.

merged into the codebase, whereas the *delivery phase* ( $t_2$ ) refers to the required time for merged PRs to be released.

We use Mann-Whitney-Wilcoxon (MWW) tests (Wilks, 2011) and Cliff’s delta effect-size measurements (Cliff, 1993) to compare different distributions of values. MWW is a non-parametric test whose null hypothesis is that two distributions come

from the same population ( $\alpha = 0.05$ ). The Cliff's delta is a non-parametric effect-size metric to verify the magnitude of the difference between the values of two distributions. We use the thresholds provided by Romano et al. (2006) to interpret the Cliff's delta, i.e.  $\text{delta} < 0.147$  (*negligible*),  $\text{delta} < 0.33$  (*small*),  $\text{delta} < 0.474$  (*medium*), and  $\text{delta} \geq 0.474$  (*large*). We analyze the entire life-cycle of a PR *before* and *after* the adoption of TRAVISCI. First, we analyze the *delivery time* ( $t_2$ ) and then, we analyze the *merge time* ( $t_1$ ). Lastly, we analyze the *lifetime* of a PR ( $t_1 + t_2$ ).

Similar to RQ1, in RQ2 we use MWW tests (Wilks, 2011) and Cliff's delta measurements (Cliff, 1993) to analyze the data. We use box plots (Williamson et al., 1989) to visually summarize the distributions and perform comparisons. In RQ2, we investigate whether the increase in the lifetime of PRs *after* adopting TRAVISCI is related to a significant increase in PR submission, merge, and delivery rates (also *after* the adoption of TRAVISCI).

We group our dataset into two time periods: *before* and *after* the adoption of TRAVISCI. For each time period, we count the number of PRs that are submitted, merged, and delivered per release. We perform three comparisons in RQ2. First, we compare whether the PR submission, merge, and delivery rates per release significantly increase *after* the adoption of TRAVISCI. Next, we verify whether there is a statistical increase in the release frequency of the projects *after* the adoption of TRAVISCI. We use the Pearson correlation (Best and Roberts, 1975), to test whether two variables are significantly correlated.

In RQ3, we use multiple regression modeling (*Ordinary Least Squares*) to describe the relationship between  $X$  (i.e., the set of explanatory variables, e.g., *churn*, *description length*), and the response variable  $Y$ , i.e., the *delivery time* of merged PRs in terms of days. We control covariates that might influence the results. For each project, we build two explanatory models, one using the PR data *before* the adoption of TRAVISCI, and another using PR data *after* the adoption of TRAVISCI. Tables 2 and 3 show the definition and rationale for each explanatory variable used in our models. Our response variable  $Y$  is the length of time between when a PR was merged and the time at which the same PR was delivered (i.e., delivery time).

We follow the guidelines of Harrell (2015) for fitting linear models. We assess how stable our models are by computing the *optimism-reduced*  $R^2$ . Finally, we use the Wald  $X^2$  maximum likelihood test to evaluate the impact of each explanatory variable in our models. The larger the  $X^2$  value for a variable, the larger the impact of a variable (da Costa et al., 2016). Next, we analyze the direction of the relationship between the most influential variables of our models and the delivery time. The process we use to build our statistical models is explained in more detail in our earlier publication (Bernardo et al., 2018).

### 3.2 Qualitative Study (Study II)

In this section, we explain the data collection and research approach of our qualitative study (Study II).

### 3.2.1 Subject Projects

As the main goal of qualitative analysis is to complement Study I, we select the same 87 GitHub projects used in Study I for Study II. The goal of Study II is to better understand the influence CI can have on the delivery time of merged PRs. We also take the opportunity to better understand the perceived influence of CI on the code review and release processes of our studied projects (i.e., according to the perception of our participants).

### 3.2.2 Data Collection

We first identify contributors who have submitted at least one PR that made into an official release of their project. The release date of the PRs must have fallen between the projects' creation date and November 11th, 2016, i.e., the range used in our search on GITHUB for Study I. By inspecting the PR meta-data of the studied projects, we find a total of 20,698 contributors that fulfill our criteria. To prioritize frequent contributors, for each studied project, we select 15% of contributors that have the highest number of delivered PRs, resulting in 3,105 contributors.

To collect our data, we designed a web-based survey and sent it by email to all 3,105 participants (i.e., the contributors of our subject projects). To encourage participation, we randomly provided six \$50 Amazon gift cards to respondents who explicitly stated their willingness to participate in the draw. To be eligible for the gift cards, the participants needed to answer all questions of the survey. In total, we received 450 responses, resulting in a response rate of 14.5% ( $450/3105$ ). Our invitation letter is available in Appendix B.

Our survey has three major *parts*. The first part concerns the influence of CI on the delivery time of merged PRs, whereas the second and third parts concern the potential influence of CI on the release and review processes of the studied projects, respectively. A complete example of our survey is available in Appendix A, which shows the questionnaire sent to participants of the *haraka/haraka*<sup>3</sup> project. Because our goal was to provide data specific to the projects of our participants, we designed 87 different questionnaires, aiming to obtain richer information about the project and encourage participants to respond more fully to the survey.

Our questionnaire is organized as follows. The first six questions (#3–#8) collect demographic information. Questions #9–#13 tap into the general experience of our participants, whereas questions #14–#26 present data specific to our participants' projects. In terms of questions' goals, questions #9–#13 and #26 capture the potential influence of CI on the delivery time of merged PRs. *Question #14* captures the perceived correctness of our approach to define the TRAVISCI adoption date in our studied projects. *Questions #17–#21* capture the potential influence of CI on the code review process of the projects, whereas *Questions #22–#25* capture the potential influence of CI on the release process of the projects.

---

<sup>3</sup> [github.com/haraka/haraka](https://github.com/haraka/haraka)

### 3.3 Research Approach

We use an inductive thematic analysis, which is designed for identifying, analyzing, and reporting themes found within qualitative data (Braun and Clarke, 2006). In this study, we use the guidelines proposed by Nowell et al. (2017) to perform our thematic analysis.

The first step in our thematic analysis is the coding of our data. This step consists of attaching codes to any piece of relevant qualitative data collected from our questionnaire. The first author conducts three sessions of open coding of the responses to open-ended questions. The second author independently conducts three sessions of open coding for 10% of the responses for each of those questions. Afterward, a new set of codes is generated by the merge of the codes created by each author. We use Cohen's Kappa test to verify the agreement rate between authors when coding the responses to 13 open-ended questions of our questionnaire. We calculate the Kappa value separately for each of the 13 questions. We achieved a median Kappa value of 0.84, indicating substantial agreement (Landis and Koch, 1977). The third author reviews the set of codes to add additional entries and resolve disagreements between the codes from the first and second authors. Next, the first author organizes the codes into *themes* through axial coding. These themes represent higher conceptual constructs (e.g., a theme might group many codes). This categorization was double-checked by the second author. We report the codes and themes generated by our thematic analysis in the result section. When reporting the results of *RQ4—RQ8*, we indicate (in superscript) the number of quotes citing each code and theme. It is important to highlight that the number in superscript does not necessarily indicate the relevance of a code, e.g., a code may be mentioned in more quotes because the code is more easily remembered by our participants. Additionally, when reporting our qualitative results, the frequency with which codes occur across responses can be higher than the total of responses. This is because a response from a participant can be associated with several codes. For example, consider the following quote “*Anything that is considered a critical security fix or major bug fix is generally shipped within 1-2 weeks of submission. This happens frequently*” (C020). We derived two codes from this quote, which are *bug fix* and *security fix*. We use representative quotes from our participants to aid in the understanding of the interpretation of the codes. We omit the participants' names by replacing them with an ID, e.g., *participant 01* receives the “name” C001. In Appendix C, we provide the IDs assigned to our participants and their project.

## 4 Quantitative Study Results

In this section, we present the results of our quantitative study (*RQ1—RQ3*).

### **RQ1: Are merged pull requests released more quickly using a CI service?**

*Only 51.3% of the projects deliver merged PRs more quickly after the adoption of TravisCI.* Out of 87 projects, we observe that 82.7% (72/87) obtained significant

*p-values* (i.e.,  $p < 0.05$ ) when comparing the delivery time of merged PRs *before* and *after* adopting TRAVISCI. Surprisingly, we observe that only 51.3% ( $37/72$ ) of these projects deliver merged PRs more quickly *after* adopting TRAVISCI. Our analyses indicate that 82.7% ( $72/87$ ) of the projects have a statistical difference on the delivery time of merged PRs, but a small median Cliff's delta of 0.304.

***In 73% ( $46/63$ ) of the projects, PRs are merged faster before adopting TravisCI.*** A total of 72.4% ( $63/87$ ) of the projects have a statistical difference on the time to merge PRs with a median Cliff's delta of 0.206 (*small*). With respect to such projects, we observe that 73% ( $46/63$ ) merge PRs more quickly *before* the adoption of TRAVISCI.

***Surprisingly, in 54% of the projects, PRs have a longer lifetime after adopting TravisCI.*** We observe that in 54% ( $47/87$ ) of our projects, PRs have a longer lifetime after the adoption of TRAVISCI. 71.3% ( $62/87$ ) of these projects yield a statistically significant difference ( $p\text{-value} < 0.05$ ) and a *non-negligible* median *delta* between the distributions of PR lifetime ( $\text{delta} \geq 0.147$ ). 37.1% ( $23/62$ ) of such projects yield a large Cliff's delta (median 0.604), while 22.6% ( $14/62$ ) and 40.3% ( $25/62$ ) of the projects obtained medium and small Cliff's deltas, respectively (medians of 0.362 and 0.223). Regarding the projects that yield a  $p\text{-value} < 0.05$ , we observe that 51.6% ( $32/62$ ) have a shorter PR lifetime *before* the adoption of TRAVISCI, while 48.4% ( $30/62$ ) have a shorter PR lifetime *after* the adoption of TRAVISCI.

**Summary:** *Surprisingly, only 51.3% of the projects deliver merged PRs more quickly after the adoption of TRAVISCI. In 54% ( $47/87$ ) of the projects, PRs have a longer lifetime after the adoption of TRAVISCI. Finally, PRs are merged faster before the adoption of TRAVISCI in 71.3% ( $63/87$ ) of the studied projects.*

**Implications:** *If the decision to adopt a CI service is mostly driven by the goal of quickening the delivery time of merged PRs, this decision must be more carefully considered by development teams.*

## **RQ2: Does the increased number of PR submissions after adopting a CI service increase the delivery time of pull requests?**

***71.3% ( $62/87$ ) of the projects receive more PR submissions after the adoption of TRAVISCI.*** Figure 2 shows the distributions of PRs submitted, merged, and delivered per release for the studied projects. We observe that projects tend to submit a median of 42.6 PRs per release *after* the adoption of TRAVISCI, while the median number of PRs submitted per release *before* the adoption of TRAVISCI is 15.3. A Wilcoxon signed rank test reveals that the increase in the number of PR submissions is statistically significant ( $p\text{-value} = 0.0001547$ ), with a Cliff's delta of 0.332 (*medium* effect-size). We also observe a significant increase in the number of merged PRs per release *after* the adoption of TRAVISCI ( $p\text{-value} = 7.897e - 05$ , with a *medium* Cliff's delta of 0.347). The number of merged PRs per release increases from 10.4 (median) (*before* the adoption of TRAVISCI) to 27.9 *after* the adoption of TRAVISCI. Interestingly, we also observe an increase in PR code churn per release *after* the adoption of TRAVISCI. We obtain a  $p\text{-value} = 0.002273$  and a Cliff's delta value of 0.27 (*small*). This significant increase in PR code churn per release may help explain the increased lifetime of PRs *after* the adoption TRAVISCI. Given that more code modifications are performed in

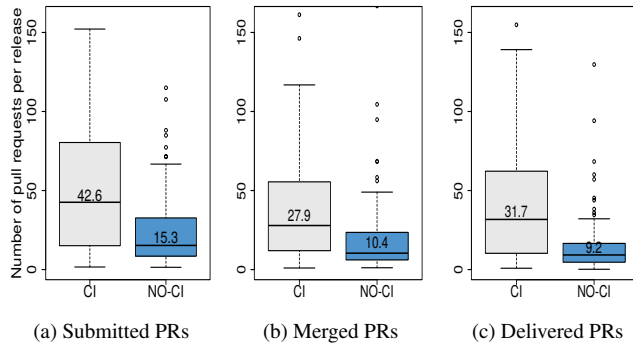


Fig. 2: PR submission, merge, and delivery rates per release.

PRs *after* the adoption of TRAVISCI, they may require more time to be reviewed, merged and delivered.

***After the adoption of TRAVISCI, projects deliver 3.43 times more PRs per release than before the adoption of TRAVISCI.*** When we analyze the PR throughput per release, we find that the number of PRs delivered per release increases significantly *after* the adoption of TRAVISCI. The number of PRs delivered increased from 9.2 to 31.7 *after* the adoption of TRAVISCI (see Figure 2c). Furthermore, the increase in the number of PRs delivered per release is statistically significant ( $p\text{-value} = 1.366e - 05$ , with a *medium* Cliff's delta of 0.3819527).

***We do not observe a significant difference in release frequency after the adoption of TRAVISCI.*** A significant increase in PR submissions may be related to an increase in release frequency *after* the adoption of TRAVISCI. Figure 3 shows the distributions of releases per year *before* and *after* the adoption of TRAVISCI (for each of the studied projects). In the median, projects tend to ship 12.03 releases per year *before* the adoption of TRAVISCI, whereas the median drops to 10.15 *after* the adoption to TRAVISCI. However, we obtain a  $p\text{-value} = 0.146$ , indicating that the differences in release frequency per year *before* and *after* the adoption of TRAVISCI are statistically insignificant. Our results suggest that the high increase in the number of PRs delivered per release is unlikely to be linked with an increase in the number of releases. We investigate whether the increased number of PRs delivered may be due to an increase in the number of contributors *after* the adoption of TRAVISCI.

***We find that 75.9% (66/87) of projects had an increase in the number of contributors per release after the adoption of TRAVISCI.*** Figure 4 shows the distributions of contributors per release both *before* and *after* the adoption of TRAVISCI. The median number of contributors per release increases from 4.4 to 11.2 *after* the adoption of TRAVISCI. We observe that the difference is statistically significant ( $p\text{-value} = 2.525e - 06$  with a *medium* Cliff's delta of 0.413).

Despite the increase in contributors and PRs delivered per release *after* the adoption of TRAVISCI, we did not observe a statistically significant correlation between PRs delivered and number of contributors. Our results show that the number of PRs delivered per release and the number of contributors in PRs per release have small

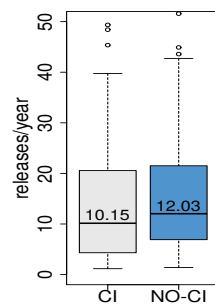


Fig. 3: Releases per year *before* and *after* TRAVISCI.

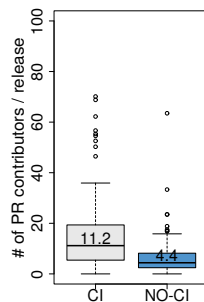


Fig. 4: PR contributors per release.

positive coefficient correlation of 0.1906346. A Pearson correlation test reveals that this correlation is not statistically significant ( $p\text{-value} = 0.07695$ ). Our observations suggest that the increase in PRs delivered *after* the adoption of TRAVISCI is not tightly related to the increase in the number of contributors per release or release frequency. The increase in PRs delivered per release might be due to the quicker feedback of automated tests provided by TRAVISCI. A qualitative study with developers may shed more light upon this matter. We further discuss this issue in Section 7.

**Summary:** *After the adoption of TRAVISCI, projects deliver 3.43 times more PRs per release than before TRAVISCI. The increase in PRs submitted, merged, and delivered after the adoption of TRAVISCI is a possible reason as to why projects may deliver PRs more quickly before the adoption of TRAVISCI.*

**Implications:** *Teams that wish to adopt TRAVISCI should be aware that their projects will not always deliver merged PRs more quickly or release more often. Instead, a pivotal benefit of a CI service is the ability to process more contributions in a given time frame.*

### RQ3: What factors impact the delivery time after adopting a CI service?

*Our models achieve a median  $R^2$  of 0.64 using pull request data before the adoption of TRAVISCI, while achieving 0.67 after the adoption of TRAVISCI.* Moreover, the



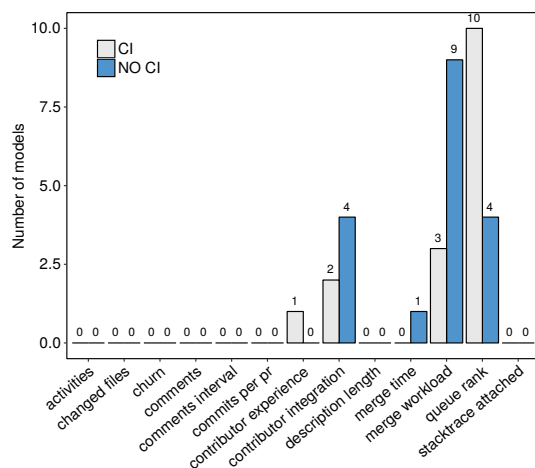


Fig. 5: The number of models per most influential variables.

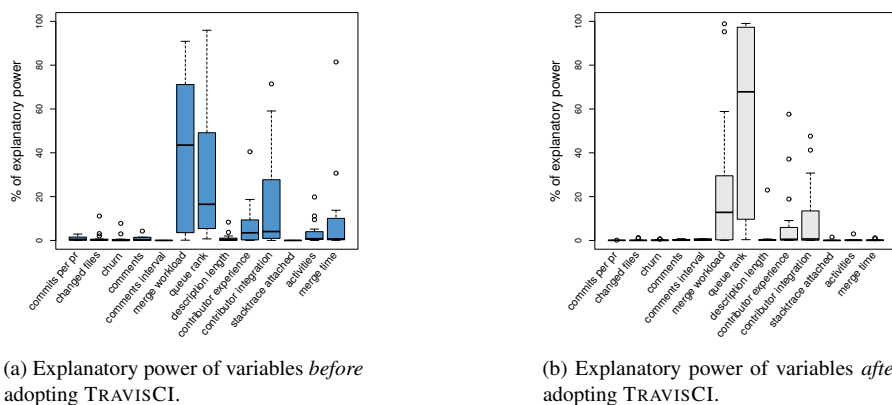


Fig. 6: Distributions of the *explanatory power* of each variable of our models.

median bootstrap-calculated optimism is less than 0.069 for both set of  $R^2$ s obtained by our models.<sup>4</sup> These results suggest that our models are stable enough to perform the statistical inferences that follow.

**The “merge workload” is the most influential variable in the models fit for the time period before the adoption of TRAVISCI.** Merge workload represents the number of PRs competing to be merged (see Table 3) at a point in time. Figure 6 shows the distributions of the explanatory power of each variable of our models. The higher the median explanatory power for a variable, the higher the influence of such a variable on the delivery time of PRs. We observe that *merge workload* has the strongest influence on our models to explain delivery time *before* the adoption of TRAVISCI.

<sup>4</sup> <https://prdeliverydelay.github.io/#rq3-r-squared-and-optimism>

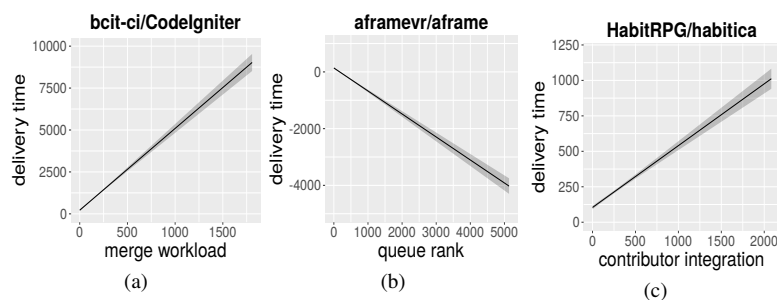


Fig. 7: The relationship between the most influential variables and delivery time.

Our models reveal that the higher the merge workload, the higher the delivery time of a PR. Figure 5 shows each explanatory variable and the number of models for which these variables are the most influential. Indeed, *merge workload* is the most influential variable in (9/18) of models fit for the time period *before* the adoption of TRAVIS CI. Figure 7 shows the relationship between the most influential variables of our models and delivery time. The relationship between *merge workload* and delivery time is shown in Figure 7a. We choose 3 models with the highest  $R^2$ s out of the 34 models to plot the relationships. Indeed, the rest of our models reveal a similar trend.<sup>5</sup>

*The “queue rank” variable is the most influential variable in the models fit for the time period after the adoption of TRAVIS CI.* *Queue rank* represents the moment at which a PR is merged in relation to other merged PRs within the release cycle. Figure 7b shows the relationship between *queue rank* and delivery time. Our models reveal that merged PRs have a lower delivery time when they are merged more recently in the release cycle. In addition, *contributor integration* is the third most influential variable in our models for both time periods, i.e., *before* and *after* the adoption of TRAVIS CI. *Contributor integration* represents the average number of days that previously delivered PRs submitted by a particular contributor took to be merged. Our models also reveal that if a contributor has their prior submitted PRs delivered quickly, their future PR submissions tend to be delivered more quickly (Figure 7c).

**Summary:** *Our models suggest that “merge workload” is the most influential variable to model the delivery time of merged PRs before the adoption of TRAVIS CI. Additionally, our models show that after the adoption of TRAVIS CI, merged PRs have a lower delivery time when they are merged more recently in the release cycle.*

**Implications:** *If software development teams plan to deliver their merged PRs more quickly to their end-users, they should consider having shorter release cycles.*

## 5 Qualitative Study Results

We first discuss the demographics of our participants, focusing on their domain, their main software development activities, and their experience using CI. Afterward, we disclose the findings of each RQ (RQ4—RQ8).

<sup>5</sup> <https://prdeliverydelay.github.io/#rq3-variables-explanatory-power>

The number of responses to each question may vary as none of the questions are mandatory in our survey (to encourage a higher response rate). Hence, not all participants answered all questions. Figure 8a shows the participants' experience in software development. We collect the data in Figure 8a from *Question #3* where the options range from "0 years" to "10 or more years". We observe that 78% ( $348/444$ ) of participants have eight or more years of experience in software development. Finally, Figure 8b shows the experience of participants with CI. 64.2% ( $282/439$ ) of participants have five or more years of experience with CI. Only five (1.3%) participants report less than one year of experience with CI.

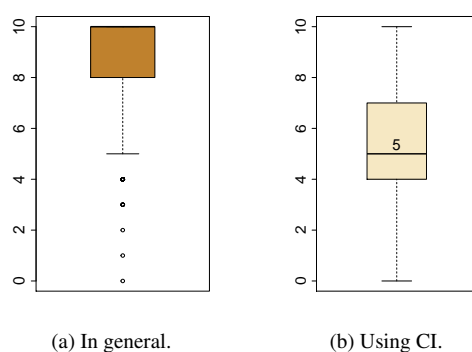


Fig. 8: Participants' experience in software development.

We observe that 51.8% ( $226/436$ ) of participants used CI in 60–100% of their projects (see Figure 9). In terms of participants' main activities, we observe that the development of new features is the most common activity among them, followed by test and review. A total of 92.9% of participants state that developing new features is one of their main activities in their projects. Another 229 participants (50.9%) state that code review is another main activity they perform in their projects. Figure 10 shows the participants' main activities according to their own classification (*Question #7*). Given that a participant can take on several roles, the sum of the percentages in Figure 10 can be greater than 100%. Considering the demographics, we were able to collect a diverse set of participants in terms of experience with CI, domain area, and development activities (e.g., bug fixing, developing new features, or code review).

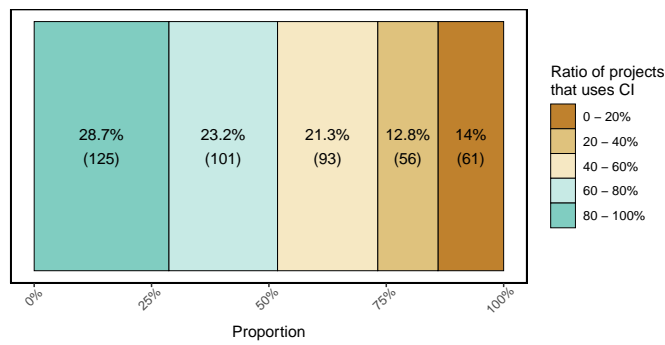


Fig. 9: Proportion of participants' projects that use CI (*Question #5*).

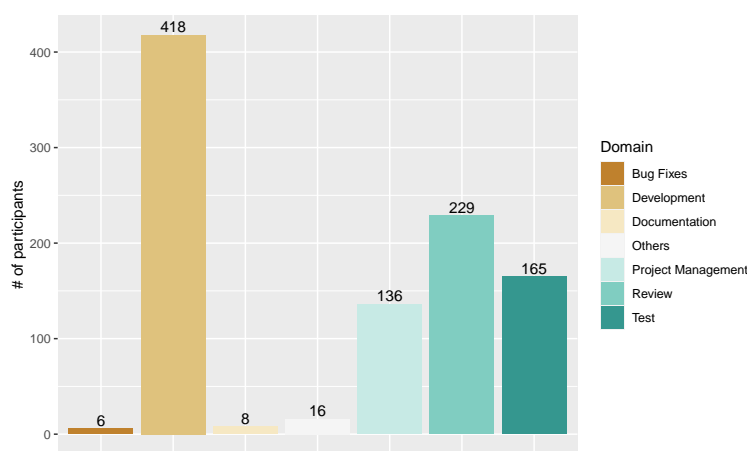


Fig. 10: Developers' main activities.

With respect to *domain expertise*, we observe that *web development* is the most common domain of our participants. 77.8% ( $350/450$ ) of participants state that web development is one of their main domains (*Question #8*). The second and third most common domains are business software development (35.3%,  $159/450$ ) and mobile applications (32.2%,  $145/450$ ). Business software is a system used to measure and improve enterprise productivity and to perform other business functions. Document Management Systems, Employee Scheduling Software, and Enterprise Resource Planning (ERP) are examples of business software. Additionally, a significant number of participants are from areas such as scientific development (i.e., those who develop software systems to analyze, visualize, or simulate processes or data) and big data (i.e., those who use scientific software to process and analyze data). This diversity of domains demonstrates that we obtain insights from several roles and development areas.

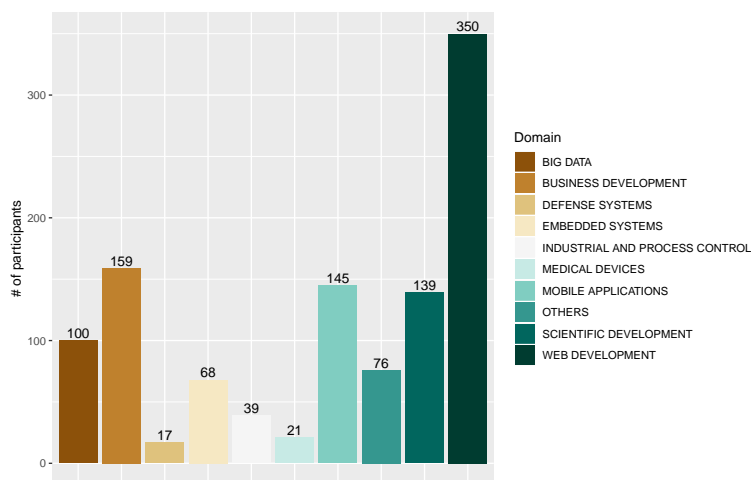


Fig. 11: Participants' domain area.

In Table 4 we present the summary of Themes that were generated by our thematic analysis for each research question (RQ). For instance, in *RQ4* we investigate the perceived influence of CI on the delivery time of merged PRs, which is associated with the following themes: *automation*, *project quality*, and *release process*. Additionally, a theme may emerge in the results of more than one RQ. For instance, the *automation* theme emerges in the results of both *RQ4* and *RQ6*. This is because participants of our study believe that *automation* impacts the delivery time of merged PRs (*RQ4*) while also impacting the release process of projects (*RQ6*). Table 4 presents an overview of the findings of the qualitative study. The detailed analysis for each theme is described in the result section of *RQ4—RQ8*.

#### RQ4: What is the perceived influence of CI on the time to deliver merged PRs?

**77% (338/441) of participants agree with the statement that CI shortens the delivery time of merged PRs.** In question #12 of our survey, we ask participants to express the extent to which they agree with the following statement: “the adoption of CI shortens the time to deliver merged PRs to end users.” Most of our participants (77%, 338/441) agree with the statement, while 16% (72/441) are neutral, and 7% (31/441) disagree or strongly disagree with the statement (see Figure 12). However, it is not clear whether these latter participants perceive CI as not having any influence on the delivery time of merged PRs or whether they perceive CI as having a negative influence on the delivery time of merged PRs.

After analyzing our participants' responses, the influence of CI on the delivery time of merged PRs was captured through the following themes: *release process*, *project quality*, and *automation*. The description and examples of mentions for each

Table 4: High-level Overview of the Themes per RQ of the qualitative study.

Research Question	Theme
RQ4: Impact of CI on the PR delivery time	Automation
	Project quality
	Release process
RQ5: Themes that impact the PR delivery time in general	PR characteristics
	Project maintenance
	Release process
	Team characteristics
	Contributors
RQ6: Influence of CI on the project release process	Testing
	Automation
	Project Stability
	Release characteristics
RQ7: Influence of CI on the project review process	CI does not impact code review
	CI impacts code review
RQ8: Impact of CI on attracting more contributors	Attractive project characteristics
	Lower contribution barrier

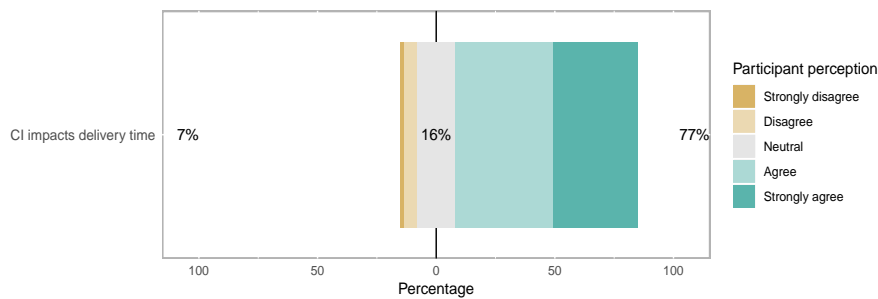


Fig. 12: Developer’s perception about the influence of continuous integration on the delivery time of merged pull requests (*Question #12*).

theme are presented in the following. Additionally, Table 5 shows the frequency of mentions for each code and theme related to how CI may influence the delivery time of merged PRs.

**Release process.**<sup>(49)</sup> Several responses to our questionnaire indicate that the adoption of CI influences the delivery time of merged PRs because CI promotes *faster release cycles*.<sup>(29)</sup> For instance, C346 declares that “*Good use of CI could help in faster release cycle, because you can be more confident in shipping something that works.*” Some practitioners of CI (Goodman and Elbaz, 2008) have claimed that some benefits of using CI are the improved release frequency and predictability. However, our quantitative

Table 5: Frequency of mentions in participants’ responses for each code and theme related to how CI may influence the delivery time of merged PRs.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
Automation	Automated testing	39	99
	Earlier feedback	24	
	Reduced testing time	10	
	Reduced burden on reviewers	9	
	Automated building	7	
	Less manual work	6	
	Improved automation	4	
Project quality	Code quality	14	130
	Code stability	23	
	Higher test coverage	6	
	Better confidence	83	
	Reduced regression risk	4	
Release process	Faster release cycle	29	49
	Automated deployment	17	
	Smaller release	3	

study (RQ2) does not support such a claim. We do not observe a significant difference in release frequency *after* the adoption of a CI service (e.g., TRAVISCI) for the studied projects. Furthermore, we found that *after* the adoption of TRAVISCI, projects delivered 3.43 times more PRs per release than *before* the adoption of TRAVISCI. We observe that although the release frequency was not significantly affected by the adoption of a CI service, projects process substantially more PRs per release than *before* the adoption of TRAVISCI. Furthermore, *automated deployment*<sup>(17)</sup> can be another step in the adoption of CI which helps projects to rapidly deliver software changes to end users (Humble and Farley, 2010). Automated deployment refers to the process of making developers’ code available to end users automatically (Rahman et al., 2015). C347 explains that “*CI is the only way to automate deployment, thus speeding up customer delivery*”. Finally, developers also mentioned *smaller releases*<sup>(3)</sup> as an influencing factor of CI on the delivery time of merged PRs. For instance, C076 states that it “*makes sense for the releases to be smaller and more frequent as a project reaches a level of stability.*”

**Project Quality.**<sup>(130)</sup> This is the theme most mentioned by our participants. According to participants, CI influences the delivery time of merged PRs by increasing *code quality*,<sup>(14)</sup> providing *code stability*,<sup>(23)</sup> and reducing *regression risks*.<sup>(4)</sup> For example, C280 states that CI promotes a “*much better quality of contributions and therefore much shorter release cycles.*” Also, CI influences code stability, as declared by C270:

“it makes the required time to deliver shorter because the maintainer can be relatively sure the change does not break other use cases.” According to Vasilescu et al. (2015), core developers using CI can discover more bugs than developers in projects not using CI. *Better quality confidence*<sup>(83)</sup> is the most mentioned code when it comes to the adoption of CI. According to participants, the delivery time of merged PRs is positively influenced by CI because “it’s much easier to trust a PR that was built and tested at a CI environment than having to do everything manually on my own machine” (C361). This trust in CI is also related to a *higher test coverage*.<sup>(6)</sup> For instance, C151 states that “CI requires a good test suite, which gives confidence in the correctness of the project.” Indeed, poor test coverage may make successful builds misleading (Felidré et al., 2019) (i.e., the builds may still contain unidentified bugs).

**Automation.**<sup>(99)</sup> According to our participants, CI also influences the delivery time of merged PRs by *improving automation*.<sup>(4)</sup> Automation is a key aspect of CI. Projects implementing proper CI must at least automate their build and testing processes. Automation leads to *less manual work*,<sup>(6)</sup> as explained by C203 when stating that “Yes, the amount of manual work involved in a release is much less when you use CI.” The *automated testing*<sup>(39)</sup> code is mentioned several times as influencing the delivery time of merged PRs, which is in accordance with the study by Rahman et al. (2015). For example, C361 states that “I worked on the Bokeh project both before and after the full integration of TRAVIS CI. Previously, major releases took months of planning due to a lack of manpower needed for running tests. With the adoption of CI, new versions can be released semi-monthly thanks to CI greatly reducing the number of man-hours needed for testing.”

The automated test execution provides *earlier feedback*,<sup>(24)</sup> which is also recurrently mentioned by our participants. For instance, C439 states that “early errors are identified by CI, so it facilitates the delivery process.” Also according to our participants, CI contributes to a *reduced testing time*.<sup>(10)</sup> C062 states that “when you have CI (+ a strong test suite) you can in some cases shorten a lot the manual test of that bug fix, or in some cases even skip it entirely (when the fix is simple enough).” Finally, it was also mentioned that CI contributes to a *reduced burden on reviewers*.<sup>(9)</sup> For instance, C026 declares that “relying on reviewers to build to test will increase time to ship massively and will be a drain on the already scarce resource of reviewers.” This observation is interesting as it may explain our results in RQ2 related to the higher number of PRs delivered per release (see Section 4), i.e., as there is less burden to reviewers, they have more capacity to review PRs that, otherwise, would have waited longer in the delivering queue. Previous studies observed that there is less discussion in PRs *after* the adoption of CI (Cassee et al., 2020), which could also save reviewers’ time. The use of CI automates several tasks in software development (i.e., build and test), thus saving time from maintainers, so they can focus on the content of the proposed software changes and launch software releases. Indeed, several participants stated that *automated building*<sup>(7)</sup> quickens the delivery time of merged PRs. For instance, C400 declares that “it [automated building] reduced the required time [to deliver] since CI built code, run tests etc.”



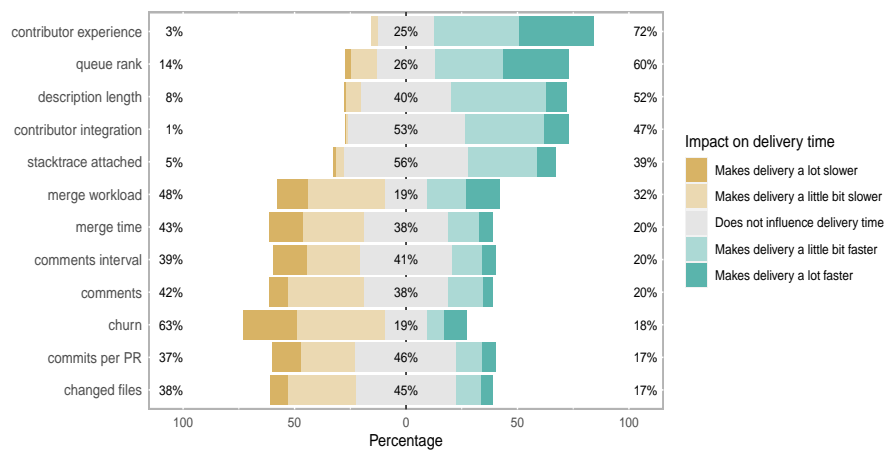


Fig. 13: Developers' perception about factors' impact on the delivery time of merged PRs (*Question #12*).

According to our participants, *PR churn* and the number of PRs waiting to be merged (*merge workload*) are the most important variables of our models in **Study I**. In question #12, we request our participants to rate the degree to which 12 variables used in our regression models (see *RQ3*, Tables 2 and 3) may influence the delivery time of merged PRs. We present the following variables: (i) a number that represents the moment at which a PR is merged compared to other merged PRs within the release cycle (queue rank); (ii) contributor integration; (iii) stack-trace attached; (iv) description size; (v) contributor experience; (vi) merge workload; (vii) changed files; (viii) churn; (ix) merge time; (x) number of commits per PR; (xi) number of comments; and (xii) interval of comments.

Our participants were invited to rate their perception regarding the impact of each above-mentioned variable on a 5-point Likert scale. The options were the following: (i) makes delivery a lot slower; (ii) makes delivery a little bit slower; (iii) does not influence the delivery time; (iv) makes delivery a little bit faster; and (v) makes delivery a lot faster. Figure 13 shows the perception of participants regarding the influence of each variable on the delivery time of merged PRs. Figure 14 shows the frequency of each rating per variable. The lower the percentage of "does not influence delivery time," the higher the perceived influence of a variable on the delivery time of merged PRs.

The variables that are most rated as having influence on the delivery time are *churn* and *merge workload*. This finding is partially in agreement with our regression models (see *RQ3*). The *merge workload* is also one of the most influential variables in our models. According to our regression models, the higher the merge workload the higher the delivery time of merged PRs. This is in agreement with the perception of 81% of participants of our survey. The influence of *merge workload* is explained by C093, when they mention that a longer delivery time can be due to "not enough developers and too many pull requests & issues to manage". Contrasting our models, participants rated *PR churn* as one of the most influential codes on delivery time. When asked

about examples of merged PRs that took long to be delivered, C109 mentioned PRs that have a “big change for a core function”. However, our regression models (RQ3) do not rate *code churn* as an influential variable to explain delivery time. Another agreement between our models and participants is that 73.8% of participants rate *queue rank* as influential to explain the delivery time of merged PRs. This corroborates the results from RQ3, which reveals that *queue rank* is influential when explaining delivery time.

Although the responses from participants confirm the influence of certain variables used in our regression models (RQ3), in Figure 14, we observe that the option “Does not influence delivery time” is frequently chosen for many other variables. 6 out of 12 variables (*stacktrace attached*, *description length*, *contributor integration*, *commits per PR*, *comments interval* and *changed files*) were ranked above 40% (2 being over 50%), which means that a substantial number of participants is skeptical regarding the influence of such variables on delivery time. This finding supports our regression model’s results, except for the contributor integration metric. The contributor integration is the third most influential variable in our models for time periods *before* and *after* the adoption of TRAVIS CI. Our models suggest that if a contributor has their prior PR delivered quickly, their future PRs are more likely to be delivered quickly.

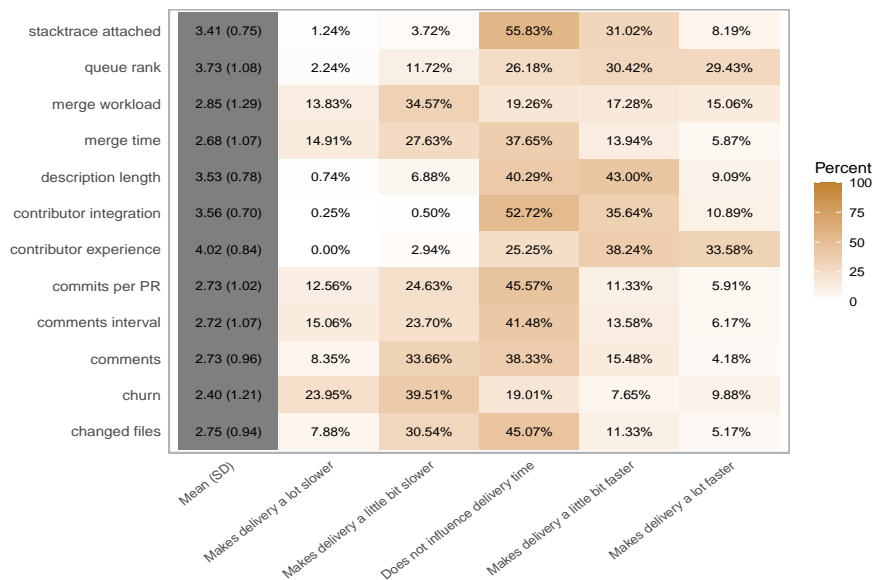


Fig. 14: Rating of codes related to delivery time of merged PRs (Question #13).

**When presented with the median delivery time of merged PRs before and after the adoption of TRAVIS CI, 42.9% (140/326) of participants state that TRAVIS CI**

**has little influence on the delivery time, attributing the change in delivery time to other unrelated factors.** According to responses to Question #26 of our survey, 42.9% of participants attribute the delivery time of PRs to other factors not directly related to the adoption of TRAVIS CI (i.e., project maintenance, release strategy, and PR characteristics). As an example, we observe that the median delivery time of PRs in the *rails/rails* project increased from 120 to 184 days *after* the adoption of TRAVIS CI. However, according to 13 participants of the *rails/rails* project, “*this is definitely not related to CI. rails/rails has its own release schedules*” (C063). Furthermore, when we presented the data for participants of the *ansible/ansible* project, where the median delivery time increased from 36 to 121 days (*after* the adoption of TRAVIS CI), C355 explained that “*all depends on the project owners decision when to release.*” Additionally, in RQ5, we further discuss the factors that participants argue to not being directly related to CI, but can influence the delivery time of merged PRs.

**Summary:** *The general perception is that CI influences the delivery time by improving automation, the release process, and the project quality. Furthermore, code churn and merge workload are the variables with the highest perceived influence on the delivery time of merged PRs. However, when showing specific project data to the respective participants, 42.9% of participants are skeptical regarding the influence of CI on the delivery time of the merged PRs.*

**Implications:** *Our participants consider that the key benefit of CI is to improve the mechanisms by which project contributions are processed (e.g., facilitating decisions related to PR submissions), without compromising the quality or overloading the reviewers and maintainers of the projects.*

#### RQ5: What are the perceived causes of delay in the delivery time of merged PRs?

The following themes are generated by our thematic analysis of the potential causes of delivery delay in merged PRs: ***PR characteristics; project maintenance; release process; team characteristics; contributors; testing and automation.*** Table 6 shows the frequency of mentions in our participants’ responses for each of the codes and themes that influence the delivery time of merged PRs. We describe all themes and the most mentioned codes in the following.

***PR Characteristics.***<sup>(330)</sup> The *characteristics of merged PRs* theme is the theme most mentioned by our participants. According to participants’ responses, *PR prioritization*<sup>(57)</sup> is one of the main factors that can shorten or lengthen the time to deliver merged PRs. The priority of a PR is recurrently associated with *bug fixes*<sup>(91)</sup> and *security fixes*.<sup>(12)</sup> According to C020, “*Anything that is considered a critical security fix or major bug fix is generally shipped within 1-2 weeks of submission. This happens frequently.*” In contrast, several participants<sup>(22)</sup> state that PRs with a longer delivery time are frequently associated with non-urgent features. As stated by C111, “*The most frequent reason [for a longer delivery time] is that the PR is not business-critical.*” In a similar vein, the study by Gousios et al. (2015) found that integrators commonly prioritize contributions by examining their criticality or urgency, e.g., bug fixes or new important features are commonly assigned a higher priority.

Table 6: Frequency of mentions in the developer responses for each code and theme that impact the delivery time of merged PRs.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
PR characteristics	Bug fix	91	330
	PR position in the release cycle	32	
	Change complexity	86	
	PR prioritization	57	
	PR size	13	
	Security fix	12	
	Code quality	11	
	Guideline adherence	7	
	Backward compatibility	6	
	Feature dependence	4	
	Good PR description	4	
	Feature improvement	4	
Breaking change	3		
Project maintainance	Maintainers availability	41	112
	Maintainers activeness	34	
	Volunteer based	12	
	Interest gauging of maintainers	11	
	Maintainer responsiveness	7	
	Maintainers workload	7	
Release process	Release cycle	86	137
	Automated deployment	16	
	Batching	11	
	Business rules	10	
	Manual release process	5	
	Misuse of CI	5	
	Release early, release often culture	4	
Team characteristics	Team size	21	40
	Small project	13	
	Open source	4	
	Paid staff	2	
Contributors	Contributor trustworthiness	4	9
	Contributor experience	3	
	Contributor and maintainers relationship	2	
Testing	Test coverage	12	35
	Testing time	9	
	Lacking tests	7	
	Broken tests	3	
	Manual testing	2	
	Build duration	2	

Additionally, the *code quality*<sup>(11)</sup> of the PR, the *change complexity*<sup>(10)</sup> and the *guideline adherence*<sup>(7)</sup> are commonly mentioned codes related to the delivery time of merged PRs. For example, C168 exemplified that “*adherence to pull request guidelines. Small fix. Clearly defined solution*” are factors that quicken the delivery time of merged PRs. Along the same lines, Gousios et al. (2015) argue that contributions conforming to project style and architecture, source code quality, and test coverage are top priorities for integrators. Finally, the *change complexity*<sup>(86)</sup> is also mentioned to help PRs to be quickly evaluated and delivered. For instance, C431 states that “*the PR I issued to Crafty was integrated very quickly, mainly because it was a trivial, absolutely non-breaking change.*” The study by Yu et al. (2016) also identified that the complexity of PRs is a factor that influence the PR latency (i.e., the time taken for a PR to be merged). Weißgerber et al. (2008) observed that smaller PRs are more likely to be accepted. Indeed, according to our survey participants, the less complex or the more trivial a PR is, the greater the more likely that the PR will be quickly delivered as less effort is needed.

**Project maintenance.**<sup>(112)</sup> Project maintenance is associated with the project maintainers’ activities. When considering the influence of the maintainers on the delivery time of merged PRs, most participants mentioned *maintainers’ availability*<sup>(41)</sup>. The study by Yu et al. (2016) found that integrators’ availability has a significant effect on PR latency. Additionally, our study suggests that *maintainers’ availability* influences the delivery of PRs. For instance, C115 considers that a longer delivery time is associated with “*long times between releases, mostly due to maintainer availability*”. Another important and frequently mentioned cause of delay is that open-source projects are *volunteer based*,<sup>(12)</sup> i.e., contributors are often volunteers (Hars and Shaosong, 2002). For instance, C336 stated that “*OSS projects are staffed by volunteers who come and go and then the priorities of the project shift and some feature become less important*”. Furthermore, *maintainers’ workload*,<sup>(7)</sup> *maintainer activeness*<sup>(34)</sup> and *maintainers’ engagement*<sup>(11)</sup> are believed to influence the delivery time of merged PRs. For instance, C419 states that “*if the maintainers of the project are interested, it [the PR] will get processed quickly.*” Previous work also observed that *workload* is a factor that plays a key role in the delivery time of addressed issues of three large open-source projects (da Costa et al., 2018), which corroborates our results. The *maintainer responsiveness*<sup>(7)</sup> also influences the delivery time of merged PRs. C090 states that “*it also helps [to deliver PRs more quickly] if maintainers can be easily contacted (IRC/Slack/Twitter)*”. Furthermore, codes related to project maintenance should be carefully considered in project management, since they might influence not only the delivery time of merged PRs, but also project success. The study by Coelho and Valente (2017) elucidates that lack of maintainers’ time and interest are factors that might lead open-source projects to fail.

**Team Characteristics.**<sup>(40)</sup> Team characteristics are also believed to influence the delivery time of merged PRs. Several participants explained that a long delivery time may occur due to *team size*.<sup>(21)</sup> As stated by C393 “*the team doing reviews were (and is still) understaffed.*” The study by Vasilescu et al. (2015) also identified that larger teams can process more PRs (i.e. merge or reject PRs). This is common in *open-source*<sup>(4)</sup> projects, as explained by C238 when stating that “*Open source projects*

*tend to delay the publication. Private projects suffer from this problem with much less impact. I wonder if it is due to the lack of a dedicated team in the open-source project, or maybe the focus isn't necessary in the part of the software I contributed."* Small project<sup>(13)</sup> is also believed to quicken the delivery of merged PRs, as explained by C321, "*on small projects some PRs might be released as a hotfix release very quickly. So I think the speed of delivery is usually in direct proportion to the size of the project."*

Overall, the codes related to project characteristics should be carefully observed for projects attempting to decrease the time to deliver their merged PRs. Our participants believe that small projects tend to deliver their PRs more quickly, as they can manage the incoming contributions more easily. With project growth (e.g., an increased number of PRs and project complexity), a small core *team size* can become a bottleneck when delivering merged PRs. In open-source projects, the bottleneck may be exaggerated as projects are volunteer-based (Hars and Shaosong, 2002) with most developers working in their free time. To overcome these barriers, projects may consider adopting strategies to deal with the large increase in contributions, as well as to deal with maintainers' inattentiveness, like *transfer the project to new maintainers* or *accept new core developers* (Coelho and Valente, 2017). Additionally, adding *paid staff*<sup>(2)</sup> to the project could be an alternative to deal with the project workload and quicken the delivery time of the merged PRs. For instance, C225 explained the importance of paid staff on the software project by stating the following: "*I have submitted PRs to very large open-source projects, like sklearn or AWS CLI. These projects typically get released frequently on established schedules by maintainers who are, in part, employed to release the projects."*

**Contributors.**<sup>(9)</sup> This theme reflects the potential influence that contributors (i.e., those who submit PRs) have on the delivery time of PRs. The contributors' social status is important when it comes to the delivery of their PRs. When analyzing the time length between the submission and merge of a PR, Yu et al. (2016) found that open-source projects prefer to quickly accept PRs originating from trusted contributors. After PRs are merged, the contributors' social status also influences the time to deliver PRs. For example, *contributor trustworthiness*<sup>(4)</sup> is often evaluated before a PR is delivered. C421 states that "*after you work with people for a while, you recognize and trust those that have proved to be good at what they do."* According to our participants, the greater the *contributor experience*,<sup>(3)</sup> the more likely the contributor will have their PR delivered more quickly. Furthermore, Soares et al. (2018) observed that the social relationship between contributors and reviewers influences the evaluation of a PR. In fact, we also identify that the *contributor and maintainer relationship*<sup>(2)</sup> positively influences the delivery time of merged PRs (i.e., contributors that are socially closer to a core team member have a higher chance to have their PRs delivered more quickly). C403, for example, states that "*some [PRs] were shipped quickly because one of maintainers is my friend then he merged immediately."*

**Testing.**<sup>(35)</sup> The delivery time of merged PRs can also be associated with testing. When asked about factors that might cause a PR to be delayed, our participants listed *testing time*,<sup>(9)</sup> *lacking tests*,<sup>(7)</sup> *broken tests*,<sup>(3)</sup> and *manual testing*.<sup>(2)</sup> For instance, in relation to testing time, C030 states that "*release time can be quite long but not due to a specific PR, but to an overall review and testing process of a whole software release."*

In this respect, reducing manual steps is a must for projects wishing to release code more frequently (Neely and Stolt, 2013). Test planning is very important in the Quality Assurance (QA) process. For instance, with automated test suites, the QA team no longer needs to manually execute the tests for the majority of the system, which would be more error-prone and slow (Neely and Stolt, 2013). However, beyond the automated test execution, projects must also be concerned with *test coverage*<sup>(12)</sup>. Automated test execution for projects with low test coverage potentially leads to certain bugs not being identified during build time (Felidré et al., 2019). Several participants mention *test coverage*<sup>(12)</sup> as influencing the delivery time of PRs. C347 declares that “*code that had a large test coverage and small PR are generally deployed safely and quickly.*” Additionally, C287 mentions issues related to building duration when testing their PRs: “*NixOS nixpkgs stable channel: CI for deep dependencies take a huge amount of computation time.*” Indeed, keeping the build and test process short (ideally by not taking more than 10 minutes) is one of the prerequisites for CI adoption (Humble and Farley, 2010). A long build duration may lead to a set of problems, for example, developers may check-in their code less often, as they have to sit around for a long time while waiting for the build (and tests) to run.

**Release process.**<sup>(132)</sup> The delivery time of merged PRs may also be associated with the release process of projects. The *release cycle*<sup>(86)</sup> code is the most cited code of this theme. For instance, C029 states that “*code was merged quickly, but had to wait for the test, review, and release cycle to complete. So had to wait for months to see the code I needed released publicly.*” Indeed, a shorter release cycle has been mentioned to shorten the time-to-market and quicken the users’ feedback loop (da Silva et al., 2016). da Costa et al. (2018) also found that traditional release cycles (which are longer cycles) could actually deliver new functionalities more quickly by using minor releases. Therefore, it seems that the higher the frequency of *user-intended* releases, the quicker the delivery of merged PRs.

When explaining potential reasons for the quick delivery of PRs, our participants recurrently mentioned *automated deployment*.<sup>(16)</sup> For instance, C100 states “*we also implemented continuous deployment so that when a change is merged, it is automatically deployed.*” Automated deployment is mandatory for the adoption of continuous deployment (CD). The goal of CD is to automatically deploy every change to the production environment (Shahin et al., 2017). In the context of the pull-based development model, CD is said to have a substantial influence on the delivery time of the proposed changes, since each merged PR is automatically deployed.

Additionally, *batching*<sup>(11)</sup> and *business rules*<sup>(10)</sup> are also important codes when it comes to the delivery time of PRs. The *batching* code is associated with the process of queuing up PRs to launch bigger releases. For instance, C092 states that a PR that experiences a long delivery time is linked to the fact that “*Usually they [project maintainers] wait for a good amount of fixes or an important fix like the ones involving security [to launch a release].*” However, project managers should be careful about the risks of bigger releases. Usually, big releases are a result of a longer release cycle, which may delay the delivery of PRs for a longer time. In contrast, smaller batch sizes would help the production environment to have fewer defects as smaller code changes may lead to faster feedback from the CI system (Neely and Stolt, 2013). Regarding the business impact on delivery time, C353 declared that “*If the change introduces*

a new feature that has an impact on end users, the marketing and customer success team need to communicate to them in advance and it takes weeks.” Indeed, the work of da Costa et al. (2018) also found that the delivery time may also be associated with the collaboration with other teams. They also mention that the marketing team is recurrently cited when delays to release occur due to other teams’ collaboration. For instance, a PR may be delayed due to the need of aligning the software release with external events for marketing reasons. Additionally, the *misuse of CI*<sup>(5)</sup> is mentioned as a factor that negatively impacts the delivery time of PRs, i.e., C118 states the following: “I’ve worked in a project that hadn’t CI nor automated tests, so each release took at least one week to be deployed. Once, we had some issues with our deploy system that made us delay the deploy for one month.” Lastly, the release culture also impacts delivery time. For example, C090 stated that “projects with release early, release often<sup>(4)</sup> culture are usually the fastest to deliver.” Overall, our results suggest that projects should reduce manual processes to quicken the release process and increase the release frequency. The adoption of continuous software engineering practices (Shahin et al., 2017), i.e., CI, Continuous Delivery (CDE), and Continuous Deployment (CD), should be considered in this regard.

**Summary:** *The delivery time of merged PRs is impacted by several factors. 87.3% (579/663) of the mentions associate the delivery time of PRs with their characteristics, the project release process, and project maintenance. According to our survey responses, simple PRs and PRs that fix bugs are delivered more quickly. The PR delivery time is also often linked to the availability of maintainers and the size of the release cycle.*

**Implications:** *Teams that wish to deliver their merged PRs more quickly to their users should also be concerned with other aspects beyond CI, such as encouraging their contributors to submit simple PRs and maintaining short release cycles.*

## RQ6: What is the perceived influence of CI on the software release process?

We report the results of RQ6 in two subsections. First, we report on the release processes of our participants’ projects in general (i.e., not considering influences of CI yet). This is important to obtain an overview of the variety of release processes in our data and will provide us with more context when interpreting the influence of CI on release processes. Afterward, we show the perceived influence of CI on the release processes of our participants’ projects.

### Release processes in general.

According to our participants, the release process of their project are *goal oriented*,<sup>(97)</sup> *maintainer oriented*,<sup>(18)</sup> follow a specific *release strategy*,<sup>(94)</sup> such as *continuous delivery*,<sup>(28)</sup> and are driven by *business*<sup>(18)</sup> and *user demand*.<sup>(7)</sup> In addition, some participants perceive an *ad hoc*<sup>(7)</sup> approach to the release process of their project. Table 7 shows the citation frequency of each theme and code related to the release process of projects.

*Goal oriented*<sup>(97)</sup> is the theme that emerged the most when it comes to the release process of our participants’ projects. An example of a goal that projects strive to



Table 7: Frequency of themes and codes as captured from our participants' responses.

Theme	Code	Frequency	
		Frequency per Code	Frequency per Theme
Goal oriented	Code stability	28	97
	New feature	22	
	Tested code	13	
	Feature completeness	10	
	Enough content	10	
	Project roadmap	7	
	Project milestone	7	
Release strategy	Fixed periods	49	94
	Continuous delivery	28	
	Release schedule	10	
	Release early, release often practice	4	
	Time-based release	3	
Maintainer oriented			18
Business-driven			18
User demand			7
Ad hoc			7

achieve is *code stability*.<sup>(28)</sup> As C370 explains, their project creates a release “*when the API is stable and extensively tested.*” Other participants also explain that their projects produce a release when a desired *new feature* has been developed,<sup>(22)</sup> the *code has been tested*,<sup>(13)</sup> or *enough content*<sup>(10)</sup> has been developed to launch a new version of the software. For example, C178 explains that “*it [the release] is usually done when enough new features and patches have been made.*” Therefore, we observe that some participants perceive that their projects adopt a more traditional release strategy (as opposed to rapid releases (da Costa et al., 2016)) to deliver new versions of their project. This strategy may also be called *feature-based releases*, where a project launches a new release when a set of bug fixes and new features are ready (Michlmayr et al., 2015). However, *feature-based releases* may not be ideal in a volunteer-based open-source project. For example, there is a risk that projects have an unpredictable release schedule, where releases may take a long time to be launched or never happen due to certain features never being completed (Michlmayr et al., 2015).

In a different vein, when explaining the release process of certain projects, our participants perceive that they adopt a more modern *release strategy*.<sup>(94)</sup> For example, in certain projects, there exist *fixed periods*<sup>(49)</sup> and a predictable *release schedule*<sup>(10)</sup> for launching new releases. Other projects use *continuous delivery*<sup>(28)</sup> or follow the *release early, release often*<sup>(4)</sup> practice. C032 explains that in the *saltstack/salt* project there is a *fixed period*<sup>(49)</sup> for its releases, “*feature release every 6 months, bug fix*

releases when necessary”. Practices such as *release early*, *release often* are well established in open source development, which leads to benefits related to quality and consistency as errors can be detected sooner (Fitzgerald and Stol, 2017).

Some participants have recurrently mentioned that the release process of their project is *maintainer oriented*,<sup>(18)</sup> *business driven*,<sup>(18)</sup> or dependent on the *user demand*.<sup>(7)</sup> We also observe seven citations that state that the process is not clearly defined or is *ad hoc*.<sup>(7)</sup> For instance, C444 explains that the *bokeh/bokeh* project “*is volunteer based, so [the release is launched] when we can and think it’s reasonable.*” C117 also states that the *fog/fog* project produces its release “*when the maintainer decided it was time to do so.*” Additionally, C003 states that the act of launching a release in *grails/grails-core* “*is a decision between business and the development team.*” The release process may also depend on the outreach of the release. For instance, C302 explains that an important factor to produce a release in *boto/boto* is “*when [they] need a wider audience.*”

*Our participants perceive different release strategies for their project. Some projects follow a feature-based release strategy, whereas others adopt a time-based release approach (e.g., release based on business needs or user demands).*

**The influence of CI on release processes.**

**52.3% (168/321) of the developers perceive an increase in release frequency after the adoption of CI.** 15.9% (51/321) of the developers do not perceive any influence from CI on release processes. Also, 3.4% (11/321) agree that CI leads to a decrease in the number of releases while 28.3% (91/321) refrain from stating an opinion (i.e., the participants do not know of or are unsure about an influence of CI on release processes, see Figure 15).

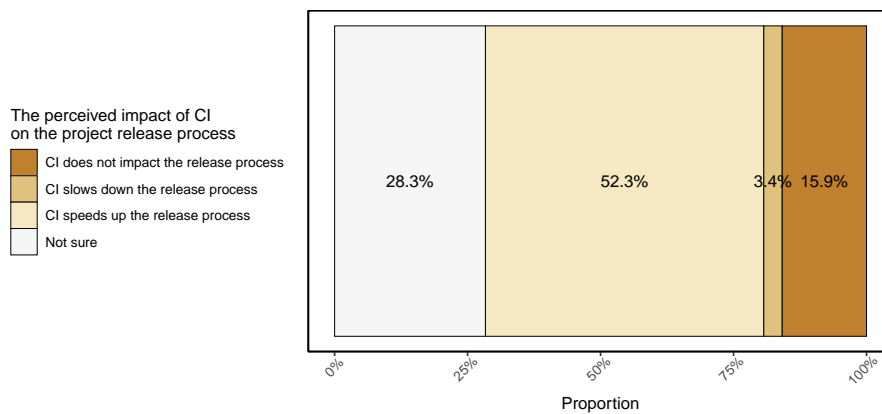


Fig. 15: Percentage of the perceived influence of CI on release processes (*Question #18*).

Table 8: Frequency of citations for each theme and code related to CI that might impact the project release process.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
Automation	Automated testing	24	59
	Release automation	17	
	Earlier feedback	11	
	Easier to produce a re-release	3	
	Automated building	2	
	Earlier integration	2	
Project Stability	Confidence	35	47
	Code stability	6	
	Releasable master	3	
	Less regressions	3	
Release characteristics	Minor releases	5	13
	Smaller releases	3	
	Bug fix releases	3	
	Security releases	2	

CI increases the release frequency by improving automation,<sup>(59)</sup> project stability,<sup>(47)</sup> and release characteristics.<sup>(13)</sup> Table 8 shows the frequency of citations identified for each theme related to the use of CI that may impact the project release process. Each theme is described in the following.

**Automation.**<sup>(59)</sup> In the previous subsection, we identified *automation* as one of the themes that may quicken the time to deliver merged PRs. In this subsection, we identify that 59 of our participants draw a relationship between *automation* improvements brought by CI and the increase in release frequency. For example, *automated tests*<sup>(24)</sup> and *release automation*<sup>(17)</sup> are frequently cited when participants explain the increase in release frequency. As explained by C070, “*The testing becomes much simpler and automated, thus it takes less time to validate a release.*” C79 complements the previous answer when they state that “*CI helps with automated tests, so we can merge and release faster.*” Furthermore, when discussing the benefits of CI in relation to release automation, C252 states: “*we can do release often with CI. It’s automated.*” CI is seen as a continuous deployment enabler, as explained by C327: “*with CI, you can increase the frequency because CI can also deploy automatically.*”

**Project stability.**<sup>(47)</sup> The release process of projects is impacted by project stability. The most cited code related to project stability is *confidence*.<sup>(35)</sup> For instance, C084 states that “*with the confidence gained from CI jobs, maintainers can reduce their time with testing tasks and focus on releases/new features.*” C056 also states that “*CI allows to keep a releasable master-branch at all times and allows external parties to*

*rely on the quality of master.*” The feedback from our participants reveals that *code stability*<sup>(35)</sup> allows releases to be prepared more easily. C331 expresses that, because CI reduces the number of blocking issues, creating a release becomes easier: “*it [CI] in general contributes to fewer bugs so there are less blocking issues so it is easier to follow a release schedule.*”

**Release characteristics.**<sup>(13)</sup> The adoption of CI is also mentioned to have changed certain characteristics of releases, which might lead to a higher frequency of releases over time. Several participants perceive that the adoption of CI started a trend of *smaller*<sup>(3)</sup> and *minor releases*.<sup>(5)</sup> Participant C203 from *chef/chef* explains the following: “*Yes it [CI] made releases and deployments much more frequent. It removed a lot of manual testing effort and validation. It also changed the culture of the team in such a way that people delivered changes in a smaller and more incremental way.*” Additionally, participants also perceived an increase in *bugfix*<sup>(3)</sup> and *security fix*<sup>(2)</sup> releases.

15.9% (<sup>5</sup>/<sub>321</sub>) of participants do not perceive any influence from CI on the release process of their projects. According to such participants, the release frequency is *maintainer oriented*<sup>(2)</sup>, depends strictly on the project *release policy*<sup>(1)</sup> or depends on the *project maturity*.<sup>(1)</sup> According to participants, instead of influencing the release frequency, CI influences the *merge time*<sup>(1)</sup>, *quickens the testing process*<sup>(1)</sup>, and provides *better quality*<sup>(5)</sup>. For instance, C317 states that “*CI may help add more features. It should not increase the frequency of releases essentially. It is a matter of policy I think.*” Additionally, C110 expresses: “*It looks like it depends on wish of repository owners and their plan*”. On another note, C169 perceives that the increase in release frequency is a side-effect of the maturation of a project, which can occur due to the adoption of CI: “*Only insofar as adoption of CI indicates the professionalization of a project, which can be correlated with more frequent releases.*”

Finally, only 3.4% (<sup>1</sup>/<sub>321</sub>) of participants perceive that CI decreases the release frequency in their projects. According to 11 participants, the decrease in release frequency can be related to the influence of CI on *code stability*<sup>(2)</sup> and *quality*.<sup>(2)</sup> For instance, C094 explains that in *ipython/ipython* “*The releases are bit less frequent. The use of CI makes the code more tested and lowers a need for bugfix-minor releases.*” Additionally, C257 declares that “*CI made the number of releases smaller and less frequent, because it helped to catch errors in the code during review. While this made releasing slower, it made the quality of those releases much better.*”

**Summary:** 52.3% (<sup>168</sup>/<sub>321</sub>) of participants perceive an increase in the release frequency after the adoption of CI. This increase is related to improvements brought by CI, such as better automation, project stability, and changes in release characteristics (e.g., smaller releases). Furthermore, only 3.4% (<sup>1</sup>/<sub>321</sub>) of participants agree that CI decreases the release frequency of their projects.

**Implications:** Teams planning to improve their release process should consider the adoption of CI, which will not always release more often. However, the automation provided by CI fosters more confidence in releasing the software.

Table 9: Frequency of citations for each theme and code related to review processes.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
Project review strategy	Peer review	109	184
	GitHub standard review	37	
	Expert review	11	
	Ad-hoc code review process	9	
	Review checklist	5	
	Development branch review	5	
	Project goals	3	
	Mailing list discussion	3	
	Pair programming	2	
Quality assurance metrics	CI check	44	130
	Tests verification	38	
	Code style	20	
	Proper documentation	7	
	Linter check	6	
	Security	3	
	Data coverage check	3	
	Efficiency check	3	
	Avoiding conflicts	2	
	Code line inspection	2	
	Error detection	2	
NA			92

### RQ7: What is the perceived influence of CI on the code review process?

We first request our participants to detail the review process of their project before we gauge their perceptions about how CI may influence the code review process. Our participants highlight that the projects have specific *review strategies*<sup>(184)</sup> that are followed. Examples of these strategies are *peer review*<sup>(109)</sup> and *expert review*.<sup>(11)</sup> Additionally, review processes have to be mindful of *quality assurance metrics*.<sup>(130)</sup> Table 9 shows the frequency of codes generated from the responses of our participants.

*Peer review*<sup>(109)</sup> and *CI check*<sup>(44)</sup> are the most frequently mentioned codes. *Peer review* is the process of manually checking violations of code standards and logical errors in a patch submitted by developers (Rahman and Roy, 2017). In the existing literature, peer review has been demonstrated to be effective for improving design quality and the overall quality of software projects (Rahman and Roy, 2017). In a similar vein, many quotes<sup>(109)</sup> from our participants highlight the use of peer review in their projects. For example, C420 states that in the *dropwizard* project, “All changes are reviewed by one or two peers depending on self-assessed complexity of the change”.

Other developers state that their projects follow the *GitHub standard review*<sup>(37)</sup> flow. For instance, C308 states: “*We use GitHub flow, open PR, require review(s), review with suggestions or concerns, discuss and revise if needed and review again, then merge.*”

We also identify more specific review strategies in some projects, e.g., *review checklist*,<sup>(5)</sup> *development branch review*,<sup>(5)</sup> *project goals*,<sup>(3)</sup> *mailing list discussion*<sup>(3)</sup> and *pair programming*.<sup>(2)</sup> For instance, when talking about the *review checklist*, C203 states that “*Code review occurs when a PR is opened. We have a checklist for both the reviewer and committer to go through for each change. Tests run on each PR and the tests need to pass before a change can be merged.*” Additionally, the submitted PR should be aligned with the project goals, as explained by C129 when declaring that “*The reviewer should read and understand all of the changes, and the changes should be in-line with the project’s conventions and goals.*” We also receive 9 responses in which participants do not identify a specific review process in their projects. These participants state that their projects follow an *ad-hoc code review process*.

The *quality assurance metrics*<sup>(130)</sup> theme has also been frequently cited by our participants when explaining the review process of their projects. Indeed, improving the quality of patches to software projects is one of the main motivators of modern code review (Bacchelli and Bird, 2013; Bavota and Russo, 2015). When observing quality assurance metrics, reviewers often use the results of *CI checks*<sup>(44)</sup> to ensure the quality of submitted PRs. Other verifications are frequently mentioned by participants when checking code quality, e.g., *tests verification*,<sup>(38)</sup> *code style*,<sup>(20)</sup> and *proper documentation*.<sup>(7)</sup> Project maintainers, in general, rely on automated tools to support the process of code review (Vasilescu et al., 2015). CI is often used in popular open-source projects to check whether the PR breaks the build. Moreover, CI verifies whether the tests pass and automatically checks whether the PR matches the project style guide (Casseo et al., 2020). According to C160, “*If CI is green and PR looks sane, merge.*” Additionally, C408 explains that, in their project, the review process also has to “*check CI jobs (lint, test, build).*”

*Our participants perceive that their projects have specific code review strategies (e.g., peer review and expert review). Moreover, the review process of our participants’ projects rely on quality assurance verification (e.g., CI check, code style, proper documentation).*

### The perceived impact of CI on the project review process

**Most of our participants’ quotes (58%, 335/578) agree that CI influences the review process of software projects.** Among the 578 quotes related to the influence of CI on the code review process, we observe that the majority (58%, 335/578) state that CI has some influence on the code review process. 15% (87/578) of quotes state that CI has no influence on code review, while 27% (156/578) of quotes did not express a clear position.

*Automation*<sup>(139)</sup> is the most cited code when it comes to how CI influences code review. As expressed by C100, “*It made it [the code review process] more efficient because the amount of manual testing that needs to happen reduced a lot. It also democratized the process, so the whole team is able to get started doing reviews.*”

Table 10: Frequency of citations for each theme and code related to the impact of CI on review processes.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
<b>CI does not influence code review</b>			87
	Automation	139	
	PR filtering	51	
	Higher confidence	43	
	Better focus	25	
	Faster review	15	
<b>CI impacts code review</b>			335
	Earlier feedback	38	
	Less review workload	7	
	Regression identification	4	
	Easier to reference failure	3	
	Improved testability	4	
	Smaller PR granularity	6	
NA			156

Indeed, automated processes (as brought by CI) are often combined with manual code reviews made by the quality assurance team (Rahman and Roy, 2017). The infrastructure of CI is frequently used with automated builds and quality checks, involving static analysis tools and automated testing (Zampetti et al., 2019). In that respect, our participants argue that CI influences code reviews because reviewers enjoy a *better focus*<sup>(25)</sup> during the review, e.g., a better focus on the code logic, security, and design. C352 elaborates on this focus when stating the following: “*We try to use linters/style checkers to remove the style nit part of the code review process. It means we spent more time thinking about the architecture and logic vs. the formatting*”.

In addition, reviewers can have a better focus on specific checks because of the *PR filtering*<sup>(51)</sup> process promoted by the use of CI. The process of PR filtering reduces the reviewers’ burden by filtering out PRs that break the build. For instance, C059 highlights that “*we do not even start reviews before CI is green.*” The study by Zampetti et al. (2019) reveals that PRs with green builds have slightly more chances to get merged than broken builds, although other process-related factors have a stronger correlation with the merge process. Table 10 shows the complete list of codes related to how CI influences code review processes.

*Higher confidence*<sup>(43)</sup> is recurrently cited when it comes to how CI influences code review. C370 argues that, with CI, there is a “*reduced time of the code review because you are more confident the PR works and focus more on code quality during the review than checking the logic.*” Another important point offered by CI within the reviewer tasks is the *earlier feedback*<sup>(26)</sup> of the proposed PRs. C316 argues that CI “*increased the speed [of code review], as certain issues are pointed out immediately*”.

**Although most of the quotes related to RQ7 agree that CI influences code review, still, 15% (87/578) of quotes state that CI has no influence on code review processes.** For instance, C215 perceives that CI influences the release process of software projects, but does not influence code review, “*I think CI does not have much to do with code review speed. Reviewers only receive one bit of information from CI, and they are expected to look at the code carefully, not the CI results. But the project manager tends to be more confident in releasing new versions with CI enabled*”.

**Summary:** Most participants agree that CI influences code review processes. 58% (335/578) of the quotes related to CI and code review agree that CI has an impact on code review. According to participants, automation is a key aspect of CI that speeds up code review.

**Implications:** CI may quicken the process of sorting which PRs are worth reviewing, e.g., a PR with green build status, which may improve the decision-making process of software projects.

### **RQ8: What is the perceived influence of CI on attracting more contributors to open-source projects?**

**59% (227/383) of quotes in RQ8 argue that projects using CI are more attractive to receive external contributors.** The most recurrent themes in RQ8 reveal that projects using CI have more *attractive characteristics*<sup>(73)</sup> (e.g., easier PR acceptance) and a *lower contribution barrier*.<sup>(154)</sup> Conversely, 26% (98/383) of the quotes argue that CI does not influence the number of project contributors, attributing the increase of contributors to other factors, such as *project growth*,<sup>(35)</sup> *maturity*<sup>(14)</sup> and *popularity*.<sup>(22)</sup> The remaining 15% (58/383) of quotes refer to answers, such as “No Answer (NA)”. Table 11 shows the complete list of themes and codes related to the influence of CI on attracting contributors to open-source projects.

**Attractive project characteristics.**<sup>(73)</sup> Several quotes state that developers feel more attracted to the characteristics of projects that use CI. They argue that, when using CI, projects tend to have a *reduced review effort/time*,<sup>(22)</sup> which may lead to *PRs being accepted* more easily.<sup>(7)</sup> As explained by C033, “*It’s easier to handle incoming changes from people, so it’s possible to have more of them.*” Additionally, the *project quality*<sup>(7)</sup> and *stability*<sup>(7)</sup> of the projects are frequently mentioned by our participants as a consequence of using CI, which may motivate developers to contribute more. In this regard, C178 declares that “*People want to work in high-quality projects. CI is a mark of quality.*” According to our participants, potential contributors also look for stability in projects and CI provides this sense of stability. As explained by C051, “*Projects that use CI tend to look more stable and serious. It might be a reason for some contributors to be attracted to more serious projects.*” Finally, potential contributors prefer projects that follow industry best practices. In this regard, C355 states: “*Maybe because the project looks more professional, following best industrial practices.*<sup>(2)</sup> *I would not contribute to a project without CI, or I would setup CI first*”.

**Lower contribution barrier.**<sup>(154)</sup> The majority of quotes stating that CI attracts more contributors explain that this is due to CI projects having a lower contribution barrier.



Table 11: Frequency of citations per theme and code related to the influence of CI on attracting contributors to open-source projects.

Theme	Code	Frequency	
		Frequency per code	Frequency per theme
Attractive project characteristics with CI	Reduced review effort/time	22	73
	Clear development process	12	
	Project quality	7	
	Project stability	7	
	Easier PR acceptance	7	
	Actively maintained project	5	
	Faster delivery	4	
	Welcoming for contributions	3	
	Regular releases	2	
	Projects with CI seems more mature	2	
	Best industrial practices followed	2	
	Lower contribution barrier with CI	CI confidence	
Build status awareness		57	
Build and test automation		28	
Engagement to contribute		6	
Lowered entry barrier		3	
Not related to CI	Project growth	35	98
	Project popularity	22	
	Project maturity	14	
	Project activeness	3	
	Non-causal correlation	24	
NA			58

This lower barrier promotes an *increased confidence*<sup>(60)</sup> in contributors when a project is using CI. The quotes also argue that CI promotes a *better build status awareness*<sup>(57)</sup> for contributors. Regarding confidence, C146 declares that “*Developers like CI, it adds confidence to your work and it is pleasurable to work in this highly structured and coordinated way.*” These observations corroborate the study by Coelho and Valente (2017), which also found CI as one of the most important maintenance practices in top open-source projects. Regarding CI promoting a better build status awareness, C084 states: “*In my case, I like to have my changeset reviewed and tested soon as possible, and CI jobs are really fast for that.*” Moreover, several participants argue that it is *easier to contribute* when a project uses CI. In this matter, C100 declares that “*As a contributor not part of the core team, CI makes it easier to understand the code, because you can look at the tests to understand the design. It makes it easier to contribute without a lot of prior knowledge of the project.*”

**Although most quotes agree that CI attracts more contributors, 26% (98/383) of quotes state that there is no causal relationship between CI and the increase in contributors.** Many participants argue that there is a *non-causal relationship*<sup>(24)</sup> between the increase in contributors and the adoption of CI. For instance, C352 declares that “*I think that people adopt CI because contributions become difficult to manage due to increasing quantity (probably driven by popularity). I would imagine these two variables are not causally related but simply correlated.*”. Indeed, many of our participants attribute the increase in contributors to *project growth*<sup>(35)</sup>, *maturity*,<sup>(14)</sup> and *popularity*.<sup>(22)</sup> For instance, C423 states that the increase in contributors is “*probably not related [to CI], but accidentally relates to a hype curve and or maturity level of the project.*” Indeed, the study by Hilton et al. (2016) shows that popular projects are more likely to use CI. Hilton et al. (2016) also found that the first CI build in their investigated projects occurred around 1 year (median) after the project creation. They argue that this is the case because the adoption of CI may not always provide a large amount of value during the very initial phases of the development of a project.

**Summary:** *Although most quotes from participants (59%, 227/383) argue that projects using CI are more attractive to potential contributors, 26% (98/383) of quotes argue that there is no causal relationship between CI and the increase in contributors, i.e., other factors play a role instead, such as project growth and increase in project popularity over time.*

**Implications:** *Open-source projects intending to attract and retain external contributors should consider the use of CI in their pipeline since CI is perceived to lower the contribution barrier while making contributors feel more confident and engaged in the project.*

## 6 Discussion

In this section, we outline the implications of our results to both research and practice in software engineering.

**Using a Continuous integration service is not a silver bullet.** Through our quantitative analyses, we observe that a CI service does not always reduce the time for delivering merged PRs to end users. In fact, analyzing 87 projects, we observe that only 51% of the projects deliver merged PRs more quickly *after* the adoption of TRAVISCI (Section 4 - RQ1). Additionally, our qualitative study reveals that there is no consensus regarding the impact of CI on the delivery time of merged PRs. 42.9% of our participants declared that CI does not have impact on the delivery time of merged PRs (Section 4 - RQ4), instead, factors such as *project release process*, *project maintenance* and *PR characteristics* (i.e., *trivial PRs*) are believed to influence the delivery time of merged PRs (Section 4 - RQ5). If the decision to adopt CI is mostly driven by the goal of quickening the delivery of merged PRs (Laukkanen et al., 2015), such a decision must be more carefully considered by development teams. Finally, previous research suggests that the adoption of CI increases the release frequency of a software project Hilton et al. (2016). However, we did not observe such an increase in our quantitative analyses (Section 4 - RQ2). Our study only considers user-intended releases, so we do not consider pre, beta, alpha, and rc (release candidate) releases in

our analyses. It might be the case that when considering only established releases, the release frequency does not statically increases *after* the adoption of TRAVISCI.

***If CI is a CD enabler, why is CD seemingly rare?*** CI, Continuous Delivery (CDE), and Continuous Deployment (CD) are complementary practices that can be used in the agile releasing engineering environment (Karvonen et al., 2017). CDE is a practice that automates the software delivery process, and it is often considered to extend CI. Therefore, a project that uses CDE, the delivery can occur at any time, with little manual effort required. Furthermore, CD is a step further from the adoption of CDE, which is a practice where projects release each successful build to end users automatically. In this context, several participants of our study consider that *automated deployment* is a subsequent step in CI adoption that can help projects rapidly deliver software changes to end users (e.g., CD can deliver merged PRs automatically). However, through the analysis of 9,312 open-source projects that use TRAVISCI, the study by Gallaba and McIntosh (2018) found that explicit deployment code is rare (2%), which suggests that developers rarely use TRAVISCI to implement Continuous Deployment. An interesting future study is to better understand the gap between CI and CD as well as how to bridge this gap. Furthermore, we observe that *before* the adoption of TRAVISCI, the merge workload is the most influential variable to model the delivery time of PRs, while *after* the adoption of TRAVISCI, the most influential variable is the moment at which a PR is merged in the release cycle (i.e., queue rank metric). One possible reason for the change in most influential variables in the time periods *before* and *after* TRAVISCI, is that *after* the adoption of TRAVISCI, the merge workload could have been better managed, leading the queue rank to be more influential on the delivery time of merged PRs. This indicates that the delivery time of merged PRs is more dependent on when the PR was merged in the release cycle than whether the project adopts a CI service. Therefore, projects that wish to quicken the delivery of merged PRs need to foster the culture of frequent release instead of solely relying on the adoption of a CI service (i.e., Travis CI) in their pipeline.

***Automation and confidence are key aspects for the throughput generated by CI.*** We observe that the adoption of a CI service is associated with many benefits, such as a higher number of contributors, PR submissions, and a higher PR churn per release (Section 4 - RQ2). However, the release frequency is roughly the same as *before* using TRAVISCI. Therefore, teams that wish to adopt a CI service should be aware that their projects will not always deliver merged PRs more quickly or release them more often, but that the pivotal benefit of a CI service is the ability to process substantially more contributions in a given time frame, which is closely tied to the automation and confidence that release managers (Section 5 - RQ6), reviewers (Section 5 - RQ7), and external contributors (Section 5 - RQ8) feel towards their codebase and project environment.

***CI may improve the decision-making process of software projects.*** Our results (Section 5 - RQ7) reveal that most contributors' quotes (58%, <sup>87/578</sup>) agree that CI impacts the time required to review PRs. According to our participants, CI may quicken the process of sorting which PRs are worth reviewing, e.g., a PRs with green builds. However, project maintainers should be concerned with the test coverage in their project, as it is essential for reviewers to be more confidence in the CI feedback to submitted PRs. We observe that contributors whose previously submitted PRs were

merged and delivered quickly, are also likely to have their future PRs delivered quickly (4 - *RQ3*). Hence, we recommend that the first PR submissions of a new contributor should be carefully crafted to maintain a successful track record in their projects (so they can build trust, causing their future PRs to be delivered more quickly). An interesting future work would be to investigate how CI can influence the decision-making process involved in different development tasks, i.e., from requirements engineering to project delivery (Sharma et al., 2021).

## 7 Threats to the validity and Limitations

In this section, we discuss the threats to the validity and limitations of our studies.

### 7.1 Threats to Validity – quantitative study

**Construct Validity.** The threats to construct validity are concerned with the extent to which the operational measures in the study really represent what researchers intended to measure. We define the delivery time of a PR as the time between when a PR is merged and the moment at which a PR is delivered through a release. However, the way we link PRs to releases may not match the actual number of delivered PRs per release. For instance, if a version control system of a project has the following release tags *v1.0*, *v2.0*, *no-ver*, and *v3.0*, we remove the *no-ver* tag. If there are PRs associated with the *no-ver* release, such PRs will be associated with release *v3.0*. However, only 5.36% (403/7519) of our studied releases are affected by this bias.

**Internal Validity.** Internal threats are concerned with the ability to draw conclusions from the relationship between the dependent variable (delivery time of merged PRs) and independent variables (e.g., release commits and queue rank). Regarding our models, we acknowledge that our independent variables are not exhaustive. Although our models achieve sound  $R^2$  values, other variables may be used to improve performance (e.g., a *boolean* indicating whether a PR is associated with an issue report and another *boolean* that verifies whether a PR was submitted by a core developer or an external contributor). Nevertheless, our set of independent variables should be approached as a starting set that can be easily computed rather than a final solution.

At the time the quantitative data was extracted from GitHub, our projects had a median of 5.1 years of life (2 without using TRAVISCI and 3.1 using TRAVISCI). As we extracted data from the entire lifetime of the projects, we collected more PRs in the *after*-TRAVISCI time period. However, the Mann-Whitney-Wilcoxon (MWW) test and Cliff's delta effect size, which were the statistical methods used to perform the comparisons, are suitable for groups of different sizes (Mann and Whitney, 1947). Furthermore, we are aware that factors not related to the adoption of TRAVISCI may have impacted the delivery time of merged PRs as well (i.e., project maturity and team size). We argue that this concern was alleviated by the conceptual replication of our study conducted by Guo and Leitner (2019). First, they replicated our study using the same subject projects and methodology. Then, they addressed the same research question of our study, by using the Regression Discontinuity Design (RDD),

which allows verifying whether there is a trend of PR delivery times over time and whether this trend changes significantly when TRAVISCI is introduced. Finally, they introduced a control group of comparable projects that never used TRAVISCI. They found that the results of our quantitative study largely hold in their replication study.

**External Validity.** External threats are concerned with the extent to which we can generalize our results (Perry et al., 2000). In this work, we analyzed 162,653 PRs of 87 popular open-source projects from GitHub. All projects adopt TRAVISCI as their CI service. We acknowledge that we cannot generalize our results to any other projects with similar or different settings (e.g., private software projects). Nevertheless, in order to achieve more generalizable results, replications of our study in different settings are required. For replication purposes, we publicize our datasets and results to the interested researcher.<sup>6</sup>

## 7.2 Limitation of our qualitative study

The main limitation of our qualitative study is the thematic analysis of the survey responses. The first author coded all responses to the 13 open-ended questions of the survey, and the second author coded a sub-set of 10% of the responses to each question, which was selected randomly. Although the systematic analytical process was used to analyze the data, the credibility of the findings is enhanced if more than one researcher completely analyzes the data. Nevertheless, we used the Cohen's Kappa test to verify the agreement level of the two authors when coding the answers. Given that we achieved an almost perfect level of agreement between them (median Kappa's value of 0.84), we believe that the coding process has been consistent.

While we achieved theoretical saturation concerning the responses to the research questions of our qualitative study, we are aware that the self-selection of our participants may have biased our results. From the subset of 3,105 individuals, we received 450 responses to the survey (14.5% response rate). Therefore, the contributors that did not respond to our survey invitation may have different views on the questions, which could lead us to different results. Additionally, the participants that responded to our survey may have been affected by social desirability bias, which is the tendency to answer questions in a way that will be seen as advantageous by others (i.e., participants responding according to what they think the "correct" answer should be, making themselves and their software development look better than it actually is).

In our qualitative study, we investigated the perception of CI practitioners regarding the influence of CI on the delivery time of PRs, and its potential influence on the review and release processes of software projects. The participants of our qualitative study are contributors from 73 out of the 87 GitHub projects studied in our quantitative analysis. However, we cannot measure the degree to which the studied projects use CI. Hence, selecting participants of projects that use a CI service, but not necessarily CI as a whole practice, may bias our analysis. This is because we may have received responses from participants that did not fully witness the benefits of CI when it is fully implemented. Nevertheless, 73.2% of the participants reported that they used CI in

---

<sup>6</sup> <https://prdeliverydelay.github.io/#datasets>

60–100% of their projects, which suggests that most of our participants have a varied experience with CI. Another issue is that we did not distinguish our participants by core and external contributors. Had we considered the responses from different types of contributors separately, different conclusions could have been drawn. On the other hand, we were able to collect a diverse set of participants, i.e., while 77.8% of the participants have web development as one of their main activities, 52% of participants consider code review as one of their main activities, and 30.9% are involved in project management (see Figure 10), which provides us with insightful feedback.

## 8 Conclusion

Our work consists of two studies that quantitatively and qualitatively investigate the influence of a CI service (e.g., TRAVISCI), and CI as a practice, on the time-to-delivery of merged PRs, respectively. In our quantitative study, we analyze 162,653 PRs of 87 GitHub projects to understand the factors that influence (and improve) the delivery time of merged PRs. In our qualitative study, we analyze 450 survey responses from participants of 73 projects (out of the initial 87 projects). We investigate the perceived influence of CI on the delivery time of merged PRs. We also study the perceived influence of CI on the code review and release processes.

As a key takeaway, our studies demonstrate that the adoption of TRAVISCI will not necessarily deliver or merge PRs more quickly. Instead, the pivotal benefit of a CI service is to improve the mechanisms by which contributions to projects are processed (e.g., facilitating decisions on PR submissions), without compromising the quality of the project or overloading developers. The automation provided by CI and the boost in developers' confidence are key aspects of using CI. For instance, CI may help the process of sorting which PRs are worth reviewing (e.g., PRs with green builds). Furthermore, open-source projects wishing to attract and retain external contributors should consider the use of CI in their pipeline, since CI is perceived to lower the contribution barrier while making contributors feel more confident and engaged in the project.

## Acknowledgments

This work is partially supported by INES ([www.ines.org.br](http://www.ines.org.br)), CNPq grants 465614/2014-0 and 425211/2018-5, CAPES grant 88887.136410/2017-00, FACEPE grants APQ-0399-1.03/17, and PRONEX APQ/0388-1.03/14.

## Data Availability

For replication purposes, we publicize our datasets and results to the interested researcher: <https://prdeliverydelay.github.io/#datasets>

## **Declarations**

**Conflict of interests.** The authors declare that they have no conflict of interest.

## Appendix A Project Survey Example

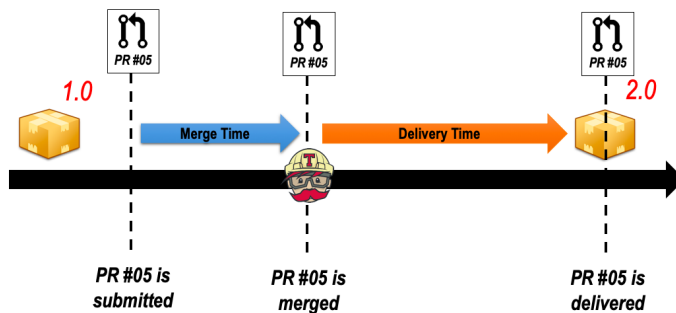
### Understanding Delivery Time of Pull-Requests

Once a pull request is merged (i.e., reviewed, tested and ready to be delivered), this pull request may still wait a long time to be delivered to end users of a software project. We call this time delivery time (see Figure 1). In this study, we intend to understand the factors that may accelerate or delay the delivery of merged pull requests. Our research team is from Brazil, New Zealand, and Australia.

This questionnaire is part of broad research effort that studied 87 large open-source projects to investigate the factors that most impact the time that is necessary to deliver merged pull requests to end users of a software project through an official release.

The survey will ask you a few general and specific questions (i.e., we will show you data that is derived from haraka/Haraka project).

Figure 1 - The basic life-cycle of a pull request.



1. By answering this survey, you will be entered into a draw of six \$50 Amazon gift cards. If you'd like to participate in the draw, please leave your email address below.

---



2. Have we correctly identified you as someone who contributes or contributed to a project that uses Continuous Integration? \*

Marcar apenas uma oval.

Yes

No *Pular para a pergunta 31*

#### Demographic Questions

3. For how long have you been developing software?

Marcar apenas uma oval.

0 1 2 3 4 5 6 7 8 9 10  
Years             or more

4. Out of your projects which use Continuous Integration (CI), which project have you most contributed over the last years?

\_\_\_\_\_

5. What is the ratio of projects that you have worked on that used CI?

Marcar apenas uma oval.

0 - 20%

20 - 40%

40 - 60%

60 - 80%

80 - 100%

6. For how long have you been developing projects that use CI?

Marcar apenas uma oval.

0 1 2 3 4 5 6 7 8 9 10  
Years            or more

7. What are your main activities on the projects you contribute to?

Marque todas que se aplicam.

- Test
- Review
- Development
- Management

Outro:  \_\_\_\_\_

8. In which domains of software development have you worked?

Marque todas que se aplicam.

- Mobile applications
- Web development
- Scientific development
- Business development
- Medical devices
- Industrial and process control
- Embedded systems
- Big Data
- Defense systems

Outro:  \_\_\_\_\_

#### Perceptions about delivery time of pull-requests

9. Can you tell us about a pull request that you submitted or reviewed and that eventually got merged, but took a long time to be effectively shipped in an official release? What were the reasons for this delay?

\_\_\_\_\_

10. Also, can you tell us about a pull request that you submitted or reviewed (and that eventually got merged), which was quickly shipped in an official release? What were the reasons for such a quick shipment?

---

11. According to your experience, do you think that the adoption of Continuous Integration (CI) may have some influence on the required time to deliver merged pull requests to end users? If so, how does CI influence this required time?

---

12. To what extent do you agree with the following statement: The adoption of Continuous Integration shortens the time to deliver merged pull requests to end users?

*Marcar apenas uma oval.*

- Strongly agree
- Agree
- Neutral
- Disagree
- Strongly disagree

13. To what extent do you agree that the factors listed in the table below influence the time necessary to release a merged pull request that is waiting to be delivered?

Marcar apenas uma oval por linha.

	Makes delivery a lot faster	Makes delivery a little bit faster	Does not influence delivery time	Makes delivery a little bit slower	Makes delivery a lot slower
The contributor experience	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The time that previously submitted PRs of a contributor were delivered	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
A PR has a stacktrace attached	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The description size of a PR	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The time at which a pull request is merged during a release cycle	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The number of PRs waiting to be delivered at the moment that a new PR was merged	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The number of comments recorded on a PR	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The time interval between the PR comments	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The number of files attached to a PR	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Number of lines	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

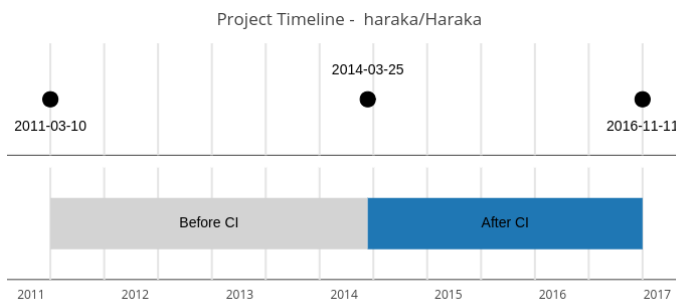
of code in a PR

Number of commits in a PR

The time that a pull request waited to be merged

Specific questions for the project haraka/Haraka

In the next questions of this survey we will show specific data of the project haraka/Haraka that you contribute to. We group such data into two buckets: before and after the adoption of CI. We considered that the project started using CI on the date of its first build on Travis-CI (2014-03-25). Also, we used the GitHub API to collect data of the haraka/Haraka project from its creation on GitHub on March, 10th, 2011 to November 11th, 2016.



14. Does the way we divided the project timeline into "Before CI" and "After CI" make sense to you?

\_\_\_\_\_

15. For how long have you worked in the haraka/Haraka project?

*Marcar apenas uma oval.*

	0	1	2	3	4	5	6	7	8	9	10	
Years	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	or more

16. Did you participate in the adoption of CI in haraka/Haraka?

*Marcar apenas uma oval.*

Yes  
 No

17. What is the releasing process/policy of your team? For example, how do you decide when to ship a release of your project?

---

18. In your opinion, did the adoption of continuous integration have any influence on the frequency of releases? Could you describe this influence to us?

---

19. When analyzing the data of the project haraka/Haraka, we observed that this project tended to ship 18 releases per year before CI, while it decreased to 6 per year after CI. How would you explain this change?

---

20. We quantitatively studied 87 popular projects on GitHub and we observed that, after adopting continuous integration, these projects delivered 3 times as many pull requests compared to before. How would you explain this change?

---

21. In our study, we also observed that, after adopting CI, the projects tended to receive contributions from more developers than before adopting CI. How would you explain this difference?

---

22. Could you briefly describe to us the code review process adopted by your project?

---

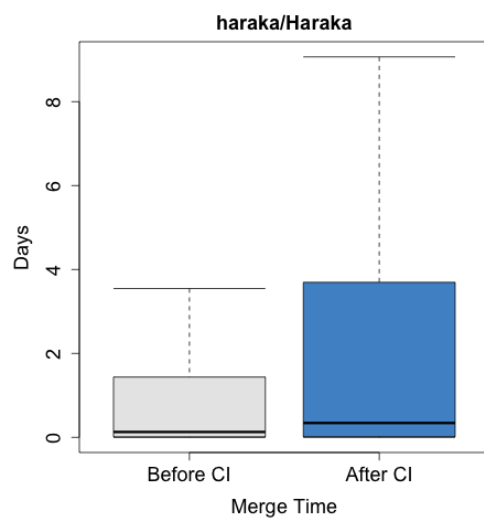
23. Did the adoption of Continuous Integration influence the code review process of your project? If so, how?

---

24. In your opinion, did the adoption of Continuous Integration increase or decrease the speed of code reviews? Why?

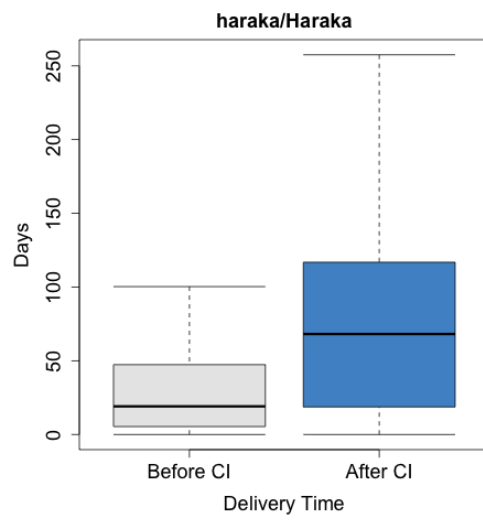
---

25. When analyzing the data of the project you contributed to (haraka/Haraka), we observed that before the adoption of travis-CI the submitted pull requests take a median of 3 hours to be merged while this time increased to 8 hours after CI. Do you think the adoption of this CI service had any influence on the time to review and merge a pull request? If so, how?





26. According to the data we collected for the project haraka/Haraka, once a pull request is reviewed and merged, this pull request still waited a median of 19 days to be shipped to end users before CI, while after CI this number increases to 68 days. Do these numbers make sense to you? Why?



### Ending our questionnaire

27. Would you like to be informed about our findings?

Marcar apenas uma oval.

Yes

No

28. Would you be willing to be contacted for a quick online follow-up interview (at a time convenient for you)?

*Marcar apenas uma oval.*

Yes

No

29. If you checked yes on one of the two questions above and have not yet left your e-mail, please leave it below.

---

30. Do you have further comments for us?

---

Sorry, we are working on our questionnaire respondents selection algorithm

As we are seeking to understand the impact of adopting continuous integration on the delivery time of pull-requests, we are looking for respondents that have used CI in some part of their career.

We are trying to improve our questionnaire respondents selection algorithm. You do not have to answer the rest of the questions. Thank you.

3. If you'd like to be informed about our findings and have not yet left your e-mail, please leave it below.

---

## Appendix B Invitation Letter

MAIL SUBJECT: **Why do PRs take so long to be delivered? Research survey**

Dear **\$contributor.name**,

We are a group of researchers from universities based in Brazil, Australia, and New Zealand. We are studying the impact of Continuous Integration on the time to release merged pull requests to end users of open source projects.

We have collected public data from the project **\$project.fullName** in the period from **\$project.creationDate** to 2016-11-11. According to our data, you have contributed **\$contributor.deliveredPRsCount** pull-requests to **\$project.fullName** which were effectively merged and delivered to end users.

As you were a contributor of the project **\$project.fullName**, we would appreciate if you shared your experience with us by answering a few questions in the following survey:

Google Form: [Understanding Delivery Time of Pull Requests](#)

The survey has 24 questions (all of them are optional) and will take less than 15 minutes to complete. To compensate you for your time, all participants that answer all questions will be entered into a draw of six \$50 Amazon gift cards.

Best Regards,

João Helis Bernardo.

PhD student at the Federal University of Rio Grande do Norte, Brazil.

## Appendix C Number of participants per project

The number of participants per project are distributed in Tables 12 and 13.

Table 12: Number of participants per project and their IDs (PART I)

#	project	participants IDs	total of participants
1	grails/grails-core	C001 – C005	5
2	saltstack/salt	C006 – C035	30
3	mozilla-b2g/gaia	C036 – C042	7
4	rails/rails	C043 – C066	24
5	owncloud/core	C067 – C079	13
6	cakephp/cakephp	C080 – C092	13
7	ipython/ipython	C093 – C097	5
8	ansible/ansible	C098 – C113	16
9	fog/fog	C114 – C125	12
10	appcelerator/titanium_mobile	C126	1
11	TryGhost/Ghost	C127 – C133	7
12	mozilla/pdf.js	C134 – C139	6
13	elastic/kibana	C140 – C143	4
14	AnalyticalGraphicsInc/cesium	C144 – C147	4
15	twbs/bootstrap	C148 – C150	3
16	sympy/sympy	C151 – C157	7
17	matplotlib/matplotlib	C158 – C169	12
18	scipy/scipy	C170 – C185	16
19	divio/django-cms	C186 – C191	6
20	woocommerce/woocommerce	C192 – C201	10
21	chef/chef	C202 – C206	5
22	puppetlabs/puppet	C207 – C211	5
23	Theano/Theano	C212 – C217	6
24	frappe/erpnext	C218 – C221	4
25	scikit-learn/scikit-learn	C222 – C228	7
26	callemall/material-ui	C229 – C231	3
27	zurb/foundation-sites	C232 – C240	9
28	laravel/laravel	C241 – C243	3
29	Leaflet/Leaflet	C244 – C251	8
30	BabylonJS/Babylon.js	C252 – C254	3

Table 13: Number of participants per project and their IDs (PART II)

#	project	participants IDs	total of participants
31	HabitRPG/habitica	C255 – C258	4
32	hapijs/hapi	C259 – C263	5
33	getsentry/sentry	C264 – C266	3
34	elastic/logstash	C267 – C268	2
35	kivy/kivy	C269 – C278	10
36	apereo/cas	C279 – C283	5
37	jashkenas/underscore	C284	1
38	ether/etherpad-lite	C285 – C289	5
39	mantl/mantl	C290	1
40	Pylons/pyramid	C291 – C298	8
41	boto/boto	C299 – C309	11
42	request/request	C310	1
43	jhipster/generator-jhipster	C311 – C318	8
44	refinery/refinerycms	C319 – C321	3
45	Netflix/Hystrix	C322 – C323	2
46	square/picasso	C324	1
47	humhub/humhub	C325 – C326	2
48	bundler/bundler	C327 – C329	3
49	isagalaev/highlight.js	C330 – C338	9
50	haraka/Haraka	C339 – C342	4
51	ReactiveX/RxJava	C343 – C345	3
52	andypetrella/spark-notebook	C346 – C348	3
53	TelescopeJS/Telescope	C349 – C351	3
54	robolectric/robolectric	C352 – C357	6
55	fchollet/keras	C358 – C362	5
56	photonstorm/phaser	C363 – C370	8
57	siacs/Conversations	C371 – C373	3
58	jsbin/jsbin	C374 – C381	8
59	buildbot/buildbot	C382 – C385	4
60	cython/cython	C386 – C390	5
61	spinnaker/spinnaker	C391	1
62	openhab/openhab	C392 – C399	8
63	jashkenas/backbone	C400 – C408	9
64	aframevr/aframe	C409 – C413	5
65	androidannotations/androidannotations	C414 – C415	2
66	dropwizard/dropwizard	C416 – C423	8
67	scikit-image/scikit-image	C424 – C425	2
68	invoiceninja/invoiceninja	C426 – C430	5
69	craftyjs/Crafty	C431 – C433	3
70	serverless/serverless	C434	1
71	bokeh/bokeh	C435 – C444	10
72	vanilla/vanilla	C445 – C448	4
73	Yelp/mrjob	C449 – C450	2

## References

- Bacchelli A, Bird C (2013) Expectations, outcomes, and challenges of modern code review. In: 2013 35th International Conference on Software Engineering (ICSE), IEEE, pp 712–721
- Bavota G, Russo B (2015) Four eyes are better than two: On the impact of code reviews on software quality. In: 2015 IEEE International Conference on Software Maintenance and Evolution (ICSME), IEEE, pp 81–90
- Beck K (2000) *Extreme Programming Explained: Embrace Change*. Addison-Wesley Professional
- Bernardo JH, da Costa DA, Kulesza U (2018) Studying the impact of adopting continuous integration on the delivery time of pull requests. In: 2018 IEEE/ACM 15th International Conference on Mining Software Repositories (MSR), IEEE, pp 131–141
- Best D, Roberts D (1975) Algorithm as 89: the upper tail probabilities of spearman's rho. *Journal of the Royal Statistical Society Series C (Applied Statistics)* 24(3):377–379
- Braun V, Clarke V (2006) Using thematic analysis in psychology. *Qualitative research in psychology* 3(2):77–101
- Cassee N, Vasilescu B, Serebrenik A (2020) The silent helper: the impact of continuous integration on code reviews. In: 2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER), IEEE, pp 423–434
- Cliff N (1993) Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin* 114(3):494
- Coelho J, Valente MT (2017) Why modern open source projects fail. In: *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, pp 186–196
- da Costa DA, Abebe SL, McIntosh S, Kulesza U, Hassan AE (2014) An empirical study of delays in the integration of addressed issues. In: *Software Maintenance and Evolution (ICSME), 2014 IEEE International Conference on*, IEEE, pp 281–290
- da Costa DA, McIntosh S, Kulesza U, Hassan AE (2016) The impact of switching to a rapid release cycle on the integration delay of addressed issues: An empirical study of the mozilla firefox project. In: *Proceedings of the 13th International Conference on Mining Software Repositories*, ACM, New York, NY, USA, MSR '16, pp 374–385
- da Costa DA, McIntosh S, Treude C, Kulesza U, Hassan AE (2018) The impact of rapid release cycles on the integration delay of fixed issues. *Empirical Software Engineering* 23(2):835–904
- Debbiche A, Dienér M, Svensson RB (2014) Challenges when adopting continuous integration: A case study. In: *International Conference on Product-Focused Software Process Improvement*, Springer, pp 17–32
- Duvall P, Matyas SM, Glover A (2007) *Continuous Integration: Improving Software Quality and Reducing Risk (The Addison-Wesley Signature Series)*. Addison-Wesley Professional
- Felidré W, Furtado L, da Costa DA, Cartaxo B, Pinto G (2019) Continuous integration theater. In: 2019 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM), IEEE, pp 1–10
- Fitzgerald B, Stol KJ (2017) Continuous software engineering: A roadmap and agenda. *Journal of Systems and Software* 123:176–189
- Fowler M, Foemmel M (2006) *Continuous integration*. Thought-Works) [http://www.thoughtworks.com/Continuous Integration.pdf](http://www.thoughtworks.com/Continuous%20Integration.pdf) p 122
- Gallaba K, McIntosh S (2018) Use and misuse of continuous integration features: An empirical study of projects that (mis) use travis ci. *IEEE Transactions on Software Engineering* 46(1):33–50
- Giger E, Pinzger M, Gall H (2010) Predicting the fix time of bugs. In: *Proceedings of the 2nd International Workshop on Recommendation Systems for Software Engineering*, ACM, pp 52–56
- Goodman D, Elbaz M (2008) "it's not the pants, it's the people in the pants" learnings from the gap agile transformation what worked, how we did it, and what still puzzles us. In: *Agile 2008 Conference*, IEEE, pp 112–115
- Gousios G, Zaidman A, Storey MA, Van Deursen A (2015) Work practices and challenges in pull-based development: the integrator's perspective. In: *Proceedings of the 37th International Conference on Software Engineering-Volume 1*, IEEE Press, pp 358–368
- Guo Y, Leitner P (2019) Studying the impact of ci on pull request delivery time in open source projects: a conceptual replication. *PeerJ Computer Science* 5:e245
- Harrell F (2015) *Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis*. Springer

- Hars A, Shaosong O (2002) Working for free? motivations for participating in open-source projects. *International journal of electronic commerce* 6(3):25–39
- Hilton M, Tunnell T, Huang K, Marinov D, Dig D (2016) Usage, costs, and benefits of continuous integration in open-source projects. In: *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering - ASE 2016*
- Humble J, Farley D (2010) *Continuous delivery: reliable software releases through build, test, and deployment automation*. Pearson Education
- Jiang Y, Adams B, German DM (2013) Will my patch make it? and how fast? case study on the linux kernel. In: *Mining Software Repositories (MSR), 2013 10th IEEE Working Conference on*, IEEE, pp 101–110
- Karvonen T, Behutiye W, Oivo M, Kuvaja P (2017) Systematic literature review on the impacts of agile release engineering practices. *Information and Software Technology* 86:87 – 100
- Landis JR, Koch GG (1977) The measurement of observer agreement for categorical data. *biometrics* pp 159–174
- Laukkanen E, Paasivaara M, Arvonen T (2015) Stakeholder perceptions of the adoption of continuous integration – a case study. In: *Proceedings of the 2015 Agile Conference*, IEEE Computer Society, AGILE '15, pp 11–20
- Mann HB, Whitney DR (1947) On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18(1):50 – 60, DOI 10.1214/aoms/1177730491, URL <https://doi.org/10.1214/aoms/1177730491>
- Michlmayr M, Fitzgerald B, Stol KJ (2015) Why and how should open source projects adopt time-based releases? *IEEE Software* 32(2):55–63
- Nagappan N, Ball T (2005) Use of relative code churn measures to predict system defect density. In: *Software Engineering, 2005. ICSE 2005. Proceedings. 27th International Conference on*, IEEE, pp 284–292
- Neely S, Stolt S (2013) Continuous delivery? easy! just change everything (well, maybe it is not that easy). In: *2013 Agile Conference*, IEEE, pp 121–128
- Nery GS, da Costa DA, Kulesza U (2019) An empirical study of the relationship between continuous integration and test code evolution. In: *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, IEEE, pp 426–436
- Nowell LS, Norris JM, White DE, Moules NJ (2017) Thematic analysis: Striving to meet the trustworthiness criteria. *International journal of qualitative methods* 16(1):1609406917733847
- Perry DE, Porter AA, Votta LG (2000) Empirical studies of software engineering: A roadmap. In: *Proceedings of the Conference on The Future of Software Engineering*, ACM, ICSE '00, pp 345–355
- Rahman AAU, Helms E, Williams L, Parnin C (2015) Synthesizing continuous deployment practices used in software development. In: *2015 Agile Conference*, IEEE, pp 1–10
- Rahman MM, Roy CK (2017) Impact of continuous integration on code reviews. In: *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, IEEE, pp 499–502
- Romano J, Kromrey J, Coraggio J, Skowronek J (2006) Appropriate statistics for ordinal level data: Should we really be using t-test and Cohen'sd for evaluating group differences on the NSSE and other surveys? In: *annual meeting of the Florida Association of Institutional Research*, pp 1–3
- Santos J, Alencar da Costa D, Kulesza U (2022) Investigating the impact of continuous integration practices on the productivity and quality of open-source projects. In: *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, pp 137–147
- Schroter A, Schröter A, Bettenburg N, Premraj R (2010) Do stack traces help developers fix bugs? In: *Mining Software Repositories (MSR), 2010 7th IEEE Working Conference on*, IEEE, pp 118–121
- Shahin M, Babar MA, Zhu L (2017) Continuous integration, delivery and deployment: a systematic review on approaches, tools, challenges and practices. *IEEE Access* 5:3909–3943
- Sharma P, Savarimuthu T, Stanger N (2021) Influence of roles in decision-making during oss developmenta study of python. In: *Evaluation and Assessment in Software Engineering*, pp 50–59
- Shihab E, Ihara A, Kamei Y, Ibrahim WM, Ohira M, Adams B, Hassan AE, Matsumoto Ki (2010) Predicting re-opened bugs: A case study on the eclipse project. In: *Reverse Engineering (WCRE), 2010 17th Working Conference on*, IEEE, pp 249–258
- da Silva ACBG, de Figueiredo Carneiro G, de Paula ACM, Monteiro MP, e Abreu FB (2016) Agility and quality attributes in open source software projects release practices. In: *2016 10th International Conference on the Quality of Information and Communications Technology (QUATIC)*, IEEE, pp 107–112
- Soares DM, de Lima Júnior ML, Plastino A, Murta L (2018) What factors influence the reviewer assignment to pull requests? *Information and Software Technology* 98:32–43

- Soares E, Sizilio G, Santos J, da Costa DA, Kulesza U (2022) The effects of continuous integration on software development: a systematic literature review. *Empirical Software Engineering* 27(3):1–61
- Stähl D, Bosch J (2014) Modeling continuous integration practice differences in industry software development. *J Syst Softw* 87:48–59
- Vasilescu B, Van Schuylenburg S, Wulms J, Serebrenik A, van den Brand MG (2014) Continuous integration in a social-coding world: Empirical evidence from github. In: *Software Maintenance and Evolution (ICSME)*, 2014 IEEE International Conference on, IEEE, pp 401–405
- Vasilescu B, Yu Y, Wang H, Devanbu P, Filkov V (2015) Quality and productivity outcomes relating to continuous integration in GitHub. In: *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2015*
- Weißgerber P, Neu D, Diehl S (2008) Small patches get in! In: *Proceedings of the 2008 international working conference on Mining software repositories*, pp 67–76
- Wilks DS (2011) *Statistical methods in the atmospheric sciences*, vol 100. Academic press
- Williamson DF, Parker RA, Kendrick JS (1989) The box plot: A simple visual method to interpret data. *Annals of Internal Medicine* 110:916–921
- Yu Y, Yin G, Wang T, Yang C, Wang H (2016) Determinants of pull-based development in the context of continuous integration. *Sci China Inf Sci* 59(8)
- Zampetti F, Bavota G, Canfora G, Di Penta M (2019) A study on the interplay between pull request review and continuous integration builds. In: *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, IEEE, pp 38–48
- Zhang X, Yu Y, Georgios G, Rastogi A (2022a) Pull request decisions explained: An empirical overview. *IEEE Transactions on Software Engineering*
- Zhang X, Yu Y, Wang T, Rastogi A, Wang H (2022b) Pull request latency explained: An empirical overview. *Empirical Software Engineering* 27(6):1–38
- Zhao Y, Serebrenik A, Zhou Y, Filkov V, Vasilescu B (2017) The impact of continuous integration on other software development practices: a large-scale empirical study. In: *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, IEEE Press, pp 60–71