

Towards Predicting the Impact of Software Changes on Building Activities

Michele Tufano

The College of William and Mary
Williamsburg, Virginia, USA
mtufano@cs.wm.edu

Hitesh Sajjani

Microsoft
Redmond, Washington, USA
hitsaj@microsoft.com

Kim Herzig

Microsoft
Redmond, Washington, USA
kimh@microsoft.com

Abstract—The pervasive adoption of Continuous Integration practices – both in industry and open source projects – has led software building to become a daily activity for thousands of developers around the world. Companies such as Microsoft have invested in in-house infrastructures with the goal of optimizing the build process. CLOUDBUILD, a distributed and caching build service developed internally by Microsoft, runs the build process in parallel in the cloud and relies on caching to accelerate builds. This allows for agile development and rapid delivery of software even several times a day. However, moving towards faster builds requires not only improvements on the infrastructure side, but also attention to developers’ changes in the software. Surely, architectural decisions and software changes, such as addition of dependencies, can lead to significant build time increase. Yet, estimating the impact of such changes on build time can be challenging when dealing with complex, distributed, and cached build systems.

In this paper, we envision a predictive model able to preemptively alert developers on the extent to which their software changes may impact future building activities. In particular, we describe an approach that analyzes the developer’s change and predicts (i) whether it impacts (any of) the Longest Critical Path; (ii) may lead to build time increase and its delta; and (iii) the percentage of future builds that might be affected by such change.

Index Terms—software building, predictive models

I. INTRODUCTION

Compiling and verifying software is a process involving several types of *tasks* with the end goal of translating source code into an efficient executable program. The process usually starts by fetching source code from a repository managed by a version control system such as Git, Subversion, Bazaar, *etc.*. Next, optional static program analysis tasks can be performed in order to assess code quality requirements defined in the organization’s build process. These include, but not limited to, security vulnerability checks, adherence of the code to stylistic and formatting rules, comments, code clones, and code quality metrics checks such as cohesion and complexity. Subsequently, source code is compiled and turned in executable (or intermediate) objects that are combined to generate potentially different versions of the executable program. Automated unit test cases execute in parallel to compilation tasks to ensure code quality. Finally, additional tasks may be executed such as storage of the program drops, cleaning of temporary files, logs, and notifications.

Continuous Integration (CI) is a development practice that involves frequent integration of code changes into a shared repository. Developers are encouraged to integrate often and daily, while each integration is verified by an automated build and tests [1]. CI aims to avoid the problems caused by a separate integration phase in the software process: unpredictability and large integration effort [1], [2]. The adoption of CI has been steadily growing, both in industry [3], [4] and open source projects [5], thanks to its ability to facilitate agile software development and allowing faster delivery cadence of software products. On the other hand, building software daily, several times a day, potentially for many different developers and teams, requires fast and reliable builds.

Companies such as Microsoft, Google, and Facebook have invested in infrastructures with the goal of accelerating the build and verification processes to enable their teams to build, integrate, and iterate faster—a procedure to ensure teams and products stay competitive. All of these modern build system rely on two main principles: distribution and pluralization as well as caching. Distribution build tasks minimizes resource limitation while caching reduces the amount of resource to be spent in the first place. All of these modern build systems helped teams to accelerate their development process. In fact, nowadays, build times no longer depend on infrastructure and resources availability, but on architectural constraints. The more dependencies a software project has between its individual components, the lower the ability to make use of parallelism (we can only compile two independent components in parallel) and caching (build targets depending on recompiled tasks cannot come from cache). We provide more details on these concepts later in this paper.

In other words, optimizing build speed nowadays becomes more and more a question of designing software systems and dependency structures to allow modern build systems to make full use of distribution and caching. Every code change adding a dependency between two previously independent software modules can impact build speed and slow down a development teams release cycles.

In this paper, we shifted the focus of build performances on the *developer-side*, by envisioning an approach able to alert developers on the extent to which their software changes may impact future building activities. The goal is to empower developers by raising awareness of the impact of their software

changes. Developers can then decide whether to perform corrective operations or confirm the current change. Such approach could be integrated in the Pull Request (PR) process, where code changes are reviewed not only using classic guidelines, but also on the impact of these changes to build time [6]. Changes that are likely to have a significant impact on build speed may need special approval and may trigger more carefully code reviews whether the newly introduced dependencies are actual necessary or whether solutions with less build speed impact could be found. We think of this process as a kind of “stay-fast” process to maintain build agility by preventing code changes negatively impacting build speed. It should raise awareness that simple changes can have significant consequences to development processes. Lebeuf et al. [7] provided a visualization framework for corresponding “get-fast” efforts.

In particular, we designed this approach to work in conjunction with CLOUDBUILD, which poses several challenges given its distributed and cached nature. In details, our approach is intended to analyze the developer’s change immediately before the build, and predict: (i) whether it impacts any of the most frequent LCPs of the branch; (ii) whether the changes may lead to build time increase and an estimation of such time delta; and (iii) an estimation of the percentage of future builds that might be affected by such change and experience build time increase. The contributions of the paper can be summarized as follows:

- we advocate for assisting developers in understanding the impact of their changes on build activities, so that corrective operations can be performed early in the development process;
- describe an approach which aims at predicting and estimating the extent to which developers’ changes may impact future build activities, in terms of time and percentage of affected builds;
- illustrate how we plan to perform the evaluation of such predictive model.

In this paper we focus our study on Microsoft’s CLOUDBUILD [8] system—a Microsoft internal cached, distributed build system. Please note that the basic concepts of CLOUDBUILD are very similar to those of Buck [9] (Facebook) and Bazel [10] (Google). Thus, we strongly believe that the overall concepts presented in this paper are not Microsoft specific but can be applied for other build systems. However, the technical details of this work remain Microsoft specific.

II. PROBLEM SCALE

In this section, we provide a discussion on how painful and impacting code changes adding new dependencies can be. Please note, we are sharing Microsoft specific experiences.

In recent years, build speed regressions have become a major issue for many product teams. The CLOUDBUILD team introduced a specific task-force responsible for these very expensive and time consuming investigations helping 1st party customer to overcome these issues. Some of the co-authors of this paper are part of this investigation team.

Changes to the dependency structure of the system under build represent one of the most common pattern of build time regression. In particular, adding new dependencies between existing modules or adding new modules that depend on already long dependency chains can cause build speed degradation of up to 50%. In nearly all cases engineers were not aware of the impact their code changes would have on build speed. Removing these dependencies after the fact was nearly always painful and expensive, with cascading dependency effects.

III. APPROACH

In this section we describe the proposed approach, which aims at analyzing a developer’s change to a code branch and estimating:

- whether or not it may negatively impact future build activities;
- the Δ in build time increase;
- the percentage of future builds that might experience this time increase.

This estimation shall be performed before the actual build, when only the source code, the changes, and the Dependency Graph (DG) are available.

At first glance, it seems that such an approach could be implemented with simple checks on the LCP extracted from the DG or, even more simply, by running the build on the changed code and measuring the build time difference with the previous builds. Unfortunately, these simple implementations would not work in a distributed and cached build environment such as CLOUDBUILD.

A. Challenges of Distributed and Cached Build Systems

Figure 1 shows an example of a DG – a directed graph representing the dependencies of a build – where nodes represent build targets, and edges represent dependencies between targets. With *target* we refer to any atomic piece of execution, such as a compilation task, a unit test, a drop *etc.*. When a full build is performed (*i.e.*, all the targets are executed), the unique LCP can be predicted statically by analyzing the DG and the execution time of each target. For example, in Figure 1 if we assume that each target t_i has the execution time, the LCP would be the following $LCP_{full} : \{t_0, t_2, t_4, t_6\}$.

Conversely, in a cached build system such as CLOUDBUILD, a full build is rarely performed, since caching allows to reuse the output of a previous build and execute only a subset of the targets. For example, if the target t_3 is changed, the build system only needs to execute t_3 and all the dependent targets (*e.g.*, t_5), while can reuse the cached version of the remaining targets (*e.g.*, the nodes in gray in Figure 1). In this case, the actual LCP would be $LCP : \{t_3, t_5\}$.

In such a build environment, there is not a single LCP, but rather a variety of possible LCPs depending on the changed

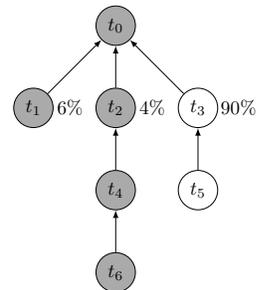


Fig. 1. Dependency Graph

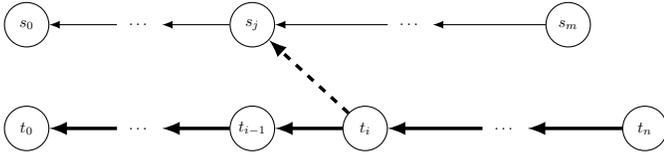


Fig. 2. Outward Dependency

Algorithm 1 Estimation for outward dependency

```

1: function OUTWARDESTIMATION
2:   if  $Time(s_0, \dots, s_j) \leq Time(t_0, \dots, t_{i-1})$  then
3:      $LCP \leftarrow t_0, \dots, t_n$ 
4:     return 0
5:   else
6:      $LCP \leftarrow s_0, \dots, s_j, t_i, \dots, t_n$ 
7:     return  $\Delta Time\{(s_0, \dots, s_j), (t_0, \dots, t_{i-1})\}$ 

```

targets and the caching status. Let $b(t_i)$ be the probability that the target t_i is built, while $c(t_i) = 1 - b(t_i)$ the probability that the same target comes from cache, then we could also compute the probability of each LCP. For example, in Figure 1, we show the probability $b(t_i)$ for the nodes t_1, t_2, t_3 . In this example, t_3 is built 90% of the times, which means that the $LCP: \{t_3, t_5\}$ is more frequent than the $LCP: \{t_2, t_4, t_6\}$. In this situation, a new dependency towards the nodes t_3 or t_5 could potentially have a greater negative effect on build activities than a new dependency on the longer LCP_{full} , which is rarely executed.

In summary, caching allows for optimized builds, which requires to execute only a subset of the targets. Thus, depending on which targets need to be re-built, a different LCP can be experienced for a build. The distributed environment allows multiple developers to perform builds in the cloud, which means that a developer's change could affect the build activity of many other developers.

Therefore, the impact of a software change should be measured not only on the build time increase introduced, but also on how often this time increase will be experienced, based on the probability of the affected LCP. In our preliminary analysis, we observed that the top-5 most frequent LCPs cover around 20-40% of the total builds performed in the last 3 months for a given project. Clearly, monitoring software changes involving those LCPs would be crucial. In our proposed approach, we consider the top- k most frequent LCPs, where k is a project-dependent and user-defined value.

B. Build Impact Estimation

Let $LCPs$ be the top- k LCPs in a project, DG_{prev} and DG_{curr} be the DG before and after the change, the approach aims at estimating whether the change impacts one of the $LCPs$, the potential build time increase Δ , and the percentage of builds affected in the future.

The approach starts by computing a graph diff between DG_{prev} and DG_{curr} , in order to detect any new edge and node added in the change. If the diff detects a newly added edge (*i.e.*, dependency), the approach checks whether one of the two endpoints of the edge (*i.e.*, dependent or dependency node)

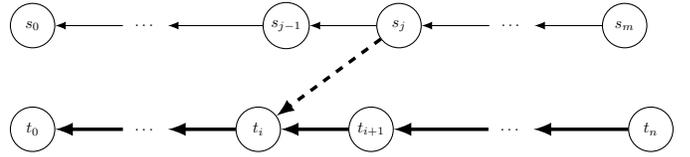


Fig. 3. Inward Dependency

Algorithm 2 Estimation for inward dependency

```

1: function INWARDESTIMATION
2:   if  $Time(s_0, \dots, s_j) \geq Time(t_0, \dots, t_i)$  then
3:      $LCP \leftarrow t_0, \dots, t_n$ 
4:     return 0
5:   else
6:     if  $Time(s_j, \dots, s_m) \leq Time(t_{i+1}, \dots, t_n)$  then
7:        $LCP \leftarrow t_0, \dots, t_n$ 
8:       return 0
9:     else
10:       $LCP \leftarrow t_0, \dots, t_i, s_j, \dots, s_m$ 
11:      return  $\Delta Time\{(s_j, \dots, s_m), (t_{i+1}, \dots, t_n)\}$ 

```

is a target node $t_i \in LCPs$. If so, we classify the dependency in two major categories: *outward* or *inward* based on whether the target node t_i is the dependent or dependency node.

1) *Outward Dependency*: Figure 2 shows the $LCP: t_0, \dots, t_i, \dots, t_n$ in bold and represented as a Dependency Chain (DC), where the execution order is from left to right. The newly added outward dependency $t_i \rightarrow s_j$ is represented as a dashed edge. The node s_j is also represented in its own DC, where s_m is the last node to be executed in the sub-tree rooted at s_j in the DG. Algorithm 1 shows the steps used by the approach to estimate the potential impact of the new dependency. In algorithm, we use a proxy function $Time(\dots)$ that estimates the execution time of a sequence of build targets. If $Time(s_0, \dots, s_j) \leq Time(t_0, \dots, t_{i-1})$ (lines 2-4), the new dependency is estimated not to introduce any delay on the LCP, instead if the opposite is true (lines 5-7), the target t_i will need to wait additional time before being executed (since s_j has not been completed yet), therefore the new dependency is estimated to generate a new $LCP \leftarrow s_0, \dots, s_j, t_i, \dots, t_n$. The delta in build time (line 7) will be experienced for the percentage of future builds that involve the build of the target s_j (or any of the previous nodes in its DC).

Note that a special case for outward dependency is when t_i is the head of the LCP. In this case, the dependency always introduce a build time delay.

2) *Inward Dependency*: Figure 3 shows the introduction of an inward dependency $s_j \rightarrow t_i$. Algorithm 2 provides the pseudo-code of the steps followed by the approach in order to estimate the potential impact of the new dependency. In particular, in order for this new dependency to increase the build time, s_j needs to experience a delay (*else* branch at line 5) and $Time(s_j, \dots, s_m) > Time(t_{i+1}, \dots, t_n)$ (lines 9-11). This will introduce a build time increase Δ (line 11) and generate a new $LCP \leftarrow t_0, \dots, t_i, s_j, \dots, s_m$. The build time increase will be experienced for the future builds that involve the execution (not from cache) of target t_i (or any of

the previous nodes in its DC).

Note that a special case for outward dependency is when t_i is the tail of the LCP. In this case, the dependency always introduce a build time delay.

C. Approximation using Historical Data

The proposed approach performs its estimation by approximating execution time and probabilities using build historical data. In particular, CLOUDBUILD logs execution time of each target and other metadata for every build, which can be statistically analyzed for future predictions. Execution time statistics can be used to approximate the function $Time(\dots)$, Top- k most frequent LCPs can be identified by observing the build logs in the recent history of the project, similarly, probabilities $b(t_i)$ can be approximated by caching statistics.

IV. EXPERIMENTAL DESIGN PLAN

While in this paper we only formally describe the idea of the proposed approach, and we have not yet evaluated the model, this section illustrates the experimental design we intend to follow in the future. In particular, we plan to evaluate the accuracy of the approach by executing it across the change history of several software projects built using CLOUDBUILD, and validating its estimations using the historical build data. In details, given an historical evaluation period of a project (e.g., the last 3 months of build activities), we execute the approach on each and every build submission (i.e., when the build is requested, before the build is executed) and evaluate its accuracy by comparing the estimation, in the current build and the future build impact, with the real build data.

Let b_i be a build in which the approach estimates that a software change generated a new LCP_{new} (from the original LCP_{old}) which introduced a build time increase Δ and it is estimated to affect $p\%$ of future builds.

A. Current Build

In the current build session, we evaluate whether the estimated LCP_{new} is actually the LCP obtained during the build execution, as reported by the historical logs and metadata.

B. Past-Future Builds

If the current build evaluation confirms the estimation, a build time period before (i.e., *past*) and after (i.e., *future*) the build b_i is selected approximately of the same length. Builds with LCP_{old} are selected from the *past* period, while builds with LCP_{new} are selected from the *future* period. The build execution time of the two sets is statistically analyzed in order to identify whether there is a statistically significant difference in build time, and compared it with the estimated Δ . Next, the amount of builds experiencing the LCP_{new} in the *future* period is compared against the predicted percentage $p\%$.

C. Historical Parameters

As discussed in III-C, the approach's estimation is based on the approximations of historical data. The amount of historical data to consider is a sensible choice. On one hand, considering only few recent data points could lead to inaccuracies due to

outliers, on the other hand, considering too much historical data could introduce imprecision due to data obsolescence. We plan to experiment and tune the historical parameters such as: (i) the number of most frequent LCPs k and their build coverage, (ii) the historical period length when computing the targets' execution time, and (iii) the caching probabilities.

V. THREATS TO VALIDITY

Threats to internal validity relate to result bias from confounding factors. The proposed approach analyzes the impact of each new dependency independently. In future work, we plan to consider the potential impact of multiple dependencies added in the same code change. Additionally, in this paper we assume that the configuration for the build environment (i.e., number of machines, cores) is stable or similar across different builds.

Threats to external validity concerns the generalizability of the research. In our case, while we envisioned this system to work with CLOUDBUILD, the approach is generalizable to other distributed cached build systems.

VI. CONCLUSIONS

In this paper we envision a predictive model able to alert developers on the extent to which their software changes may impact future build activities.

As future work, we plan to evaluate the proposed approach and test its utility and usability for developers. Additionally, we plan to also incorporate positive feedback in the prediction, such as when a software change could lead to faster builds.

VII. ACKNOWLEDGMENT

We thank Kıvanç Muşlu and Christopher DuVarney from Tools for Software Engineers group for their valuable help.

REFERENCES

- [1] M. Fowler and M. Foemmel, "Continuous integration," *Thought-Works* [http://www.thoughtworks.com/Continuous Integration.pdf](http://www.thoughtworks.com/Continuous%20Integration.pdf).
- [2] E. Laukkanen, M. Paasivaara, and T. Arvonen, "Stakeholder perceptions of the adoption of continuous integration – a case study," in *2015 Agile Conference*, 2015.
- [3] D. G. Feitelson, E. Frachtenberg, and K. L. Beck, "Development and deployment at facebook," *IEEE Internet Computing*.
- [4] "On the journey to continuous deployment: Technical and social challenges along the way," *Information and Software Technology*, vol. 57.
- [5] J. Holck and N. Jørgensen, "Continuous integration and quality assurance: A case study of two open source projects," *Australasian Journal of Information Systems*, vol. 11, no. 1, 2003.
- [6] R. Wen, D. Gilbert, M. G. Roche, and S. McIntosh, "BLIMP Tracer: Integrating Build Impact Analysis with Code Review," in *Proc. of the International Conference on Software Maintenance and Evolution (ICSME)*, 2018.
- [7] C. Lebeuf, E. Voyloshnikova, K. Herzig, and M.-A. Storey, "Understanding, debugging, and optimizing distributed software builds: A design study," in *Proceedings of the 34th International Conference on Software Maintenance and Evolution*, ser. ICSME '18, 2018.
- [8] H. Esfahani, J. Fietz, Q. Ke, A. Kolomiets, E. Lan, E. Mavrinac, W. Schulte, N. Sanches, and S. Kandula, "Cloudbuild: Microsoft's distributed and caching build service," in *Proceedings of the 38th International Conference on Software Engineering Companion*, ser. ICSE '16, 2016, pp. 11–20.
- [9] Facebook buck. [Online]. Available: <https://buckbuild.com/>
- [10] Google bazel. [Online]. Available: <http://www.bazel.io/>