To test, or not to test: A proactive approach for deciding complete performance test initiation

1st Omar Javed
Department of Informatics, Università della svizzera italiana
Lugano, Switzerland
omar.javed@usi.ch

2nd Prashant Singh
Department of Information Technology, Uppsala University
Uppsala, Sweden
prashant.singh@scilifelab.uu.se

3rd Giles Reger
School of Computer Science, University of Manchester
Manchester, United Kingdom
giles.reger@manchester.ac.uk

4th Salman Toor
Department of Information Technology, Uppsala University
Uppsala, Sweden
salman.toor@it.uu.se

Abstract—Software performance testing requires a set of inputs that exercise different sections of the code to identify performance issues. However, running tests on a large set of inputs can be a very time-consuming process. It is even more problematic when test inputs are constantly growing, which is the case with a large-scale scientific organization such as CERN where the process of performing scientific experiment generates plethora of data that is analyzed by physicists leading to new scientific discoveries. Therefore, in this article, we present a test input minimization approach based on a clustering technique to handle the issue of testing on growing data. Furthermore, we use clustering information to propose an approach that recommends the tester to decide when to run the complete test suite for performance testing. To demonstrate the efficacy of our approach, we applied it to two different code updates of a web service which is used at CERN and we found that the recommendation for performance test initiation made by our approach for an update with bottleneck is valid.

Index Terms—Recommendation system, Performance analysis, Unsupervised learning, Software testing.

I. INTRODUCTION

performance issues during the early stage of software development is crucial to avoid software failures [1]. It has been shown that performance bugs can cause critical software failures that can result in the abandonment of big projects [2]. Furthermore, performance bugs are often very costly to diagnose because their consequences are not easily observable, which leads to many hours of diagnosis and fixing [1].

One approach to performance bug detection is unit testing. However, anecdotal evidence shows that unit testing for performance analysis is not nearly as common as unit testing of code correctness (i.e., functional bug detection) [3]. One of the reasons is that performance issues arise only with particular inputs. Hence, they can easily escape into production.

To overcome this issue, studies have shown that profiling unit tests with different inputs can lead to identifying performance bugs [1], [4]. Furthermore, these studies use inputs that are gathered from analyzing bug patches. However, they do not consider large-scale input sets used by real users.

A big challenge for performance testing is to identify interesting input space i.e., inputs that can expose performance bugs [4]. This makes performance testing over a large set of inputs a very time-consuming process. The reason is that unit tests are executed repeatedly on different inputs to evaluate the performance of the software on code changes (i.e., performance regression testing).

Furthermore, the high testing time can be more problematic for systems in organizations that deals with generation of growing data. For example, CERN (the European Organization for Nuclear Research) processes prodigious amounts of data every day (in the order of Petascale) [5]. Moreover, the Large Hadron Collider (LHC) which is the "most powerful particle accelerator" facilitate different detectors such as the Compact Muon Solenoid (CMS), ATLAS etc in carrying out experiments, produces a large volume of data i.e., 90 PetaByte each year [5]. The issue for such system is the dependency on the stored data that grows over time. Therefore, it is necessary for performance testing to account for that growth.

To avoid the overhead and turnaround time involved in checking for performance issues on code change. Existing studies [6], [7] have proposed different strategies for optimizing performance regression testing based on either test selection or code metrics. However, these studies do not consider the effect of data growth.

Therefore, in this article, we use a web service, employed by the CMS experiment at CERN called Conditions Uploader service [8] as a use case to demonstrate the problem of performance regression testing with increasing number of test inputs. We use the Conditions uploader service because it adds a level of complexity to the traditional performance testing paradigm i.e., the data growth problem. Here, the growth problem is because of the continued running of LHC that further generates more data. This growing data is used as test inputs for the testing of CMS uploader service. Consequently, making performance regression testing on increasing inputs a challenging task.
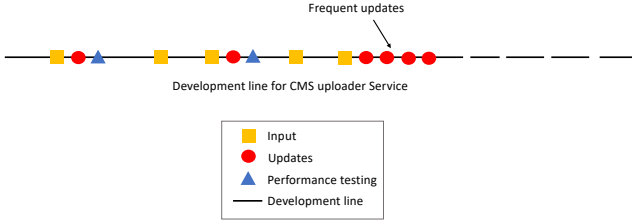
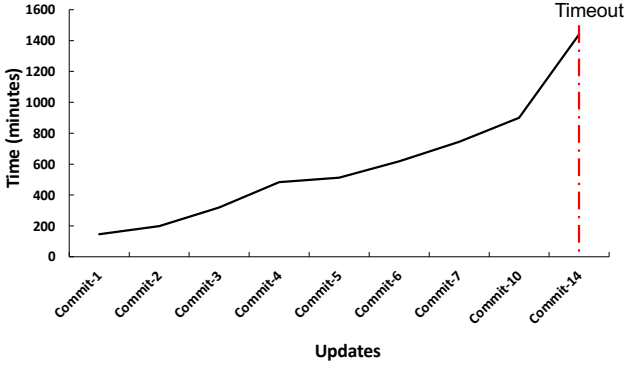Fig. 1. Figure showing the development timeline of CMS uploader service.

Fig. 2. Figure demonstrating performance testing with increasing inputs. It shows the increase in execution time of testing with the increase in inputs at each commit. Commit-1 is the oldest update whereas Commit-14 is the latest.

Figure 1 and 2 uses the CMS uploader service to demonstrate two key problems of performance regression testing with growing data. The problem is demonstrated in the context of Continuous Integration (CI) which is now a widely used practice, as it performs automated builds and regression tests (i.e., functional and non-functional) of software. We discuss these two problems as follows:

1) Frequent Update Problem: The first problem is the execution of a test suite on frequent updates (as demonstrated by red circles in the Figure 1). The figure provides an illustration of the practice of running a test suite on frequent updates to check for performance bugs. This practice is often impractical and challenging for developers, since the waiting time for performance tests is high. Therefore, developers run tests in batches i.e., testing is done on a weekly or monthly basis or when there is an important update [9]. This infrequent execution of the tests makes it difficult for developers to detect which update caused the issue [9].

2) Data growth Problem: The second problem is the execution of the test suite after the addition of new data (or test inputs) leading to a higher execution time; this can be seen in Figure 2 which represents the actual part of update history of CMS uploader web service. The figure highlights the time taken when increasing data i.e., number of inputs at each update. Furthermore, we observe from the figure that regression testing with growing test inputs can quickly lead to

a timeout in CI; this is because existing CI services limit the use of resources for a longer duration of time.

Based on this premise, we propose an approach to overcome the aforementioned problems of performance regression testing.

B. Contribution of the article

We present three key contributions of this study, which are as follows:

1. We discuss and address the issue of (ever) growing data in the context of web service testing. This study is timely since many organizations are facing (or will face) this imminent issue.

2. We demonstrate test input minimization approach. With our approach, testers can overcome the problem of high test execution time with growing data.

3. We propose a recommendation approach to help developers in deciding when it is feasible to initiate performance testing. This recommendation approach can address the problem of costly test execution at each update.

C. Structure of the article

The article is divided into separate sections, which are as follows:

• Section II presents the background and related work in this area of study.

• Section III explains our approach for the recommendation of performance regression testing.

• Section IV and V presents evaluation and our findings.

• Section VI provides our conclusion for the study.

• In appendix, we present our analytic system for CI and also significance test.

D. Resources

We provide implementation of our algorithms for recommending performance testing, which are available in a public repository website [10]. Furthermore, we also provide source code of our CI analytic system as well as the instructions for setting up the front-end and back-end of our analytic system. This can also be found in a public repository [11].

II. BACKGROUND AND RELATED WORK

In this section, we provide a background of the CMS uploader service. Furthermore, we explore the existing state-of-the-art in the area of performance regression tests and input selection. Finally we show existing CI tools and services as well as highlight CI tools used in performance regression testing.

A. CERN CMS Uploader service

Data growth is becoming a big problem for systems in organizations that deals with ever growing data. One such example is CERN which has world's largest and most powerful particle accelerator known as Large Hadron Collider (LHC). Furthermore, the Compact Muon Solenoid (CMS) is a general-purpose detector at one of the interaction points on the LHC [12]. Physics analysis requires a process of reconstruction that makes use of collision events as well as alignment and calibration or "conditions" data. During LHC runs, conditions data begins with its computation and ends with its upload to a central Conditions database.

The service responsible for uploading conditions data is the CMS Uploader service. It performs correctness checks before conditions data are uploaded. The CMS uploader service has 65 test cases and approximately 40K test inputs representing conditions data that continues to grow due to LHC runs. Therefore, adequate test size, growing input space, and its usage by real users at CERN make the CMS uploader service a reasonable use case for our study.

Furthermore, the importance of automated performance analysis and testing over different inputs for the CMS uploader service has already been demonstrated in previous works [13], [14]. In this study, we will address the issues and challenges of performance regression testing of the CMS uploader service with growing data.

B. Regression performance test selection

Traditional regression test selection techniques focus mostly on correctness tests. There are few studies that focus on performance regression test selection [6], [7], [15]. These studies use previous commit information on making test selection. In contrast, our approach helps developers in deciding when to run performance tests. Our approach is based on approximating the slowdown of the system by sampling minimal information of code behavior with respect to test inputs. Therefore, our approach, does not require historical code change information to make predictions, rather uses information about current code modifications to make a decision.

Furthermore, there are studies based on heuristics to identify potential performance changes e.g., assigning costs to source code changes [16], prioritizing performance tests based on the impact a code change has on the system's performance [17] and selection of test inputs to determine a performance model (e.g. the function of the response time) [18], [19]. We draw inspiration from these studies to build a recommendation approach that identifies regression performance test opportunities.

C. Input selection for testing

Existing study on input selection [4] tries to understand the effectiveness of performance bug detection. A specific set of inputs is generated based on the information present in the bug report. To select inputs based on program behavior, Qi Luo [20] proposed an approach that analyzes the performance behavior of the program and tries to connect the behavior with test input data. This work is similar to ours in terms of selecting inputs based on program behavior. Furthermore, one of the state-of-the-art approach DeepGini [21] prioritizes test inputs by measuring the confidence of Deep Neutral Networks (DNN) for classifying each test input. Test inputs are prioritized higher based on similar probabilities for all classes. Moreover, a study [22] aims to estimate the accuracy of the DNN model by selecting a small set of test inputs. Compared to these studies, we use unsupervised learning to reduce the number of test inputs that exercise similar program behavior based on unit test execution.

D. Continuous Integration services

Continuous integration (CI) is the practice of integrating different code changes by multiple developers. It consists of automatically checking correctness of the code before integration. There are a number of popular CI options available such as CircleCI [23], TravisCI [24], Codeship [25], Bitbucket pipelines [26], SemaphoreCI [27], and many more. These CI tools allow developers to automatically execute unit tests (in a remote server) before pushing changes to the code on a code repository service like GitHub. Hence, a CI process provides a separation of concern (i.e., one service handles code version control, while a different service handles fault detection).

However, many existing CI services only provide functionality for correctness tests. There are separate tools and plugins that support performance testing in CI e.g., Jenkins uses the Taurus tool for load testing [28]. PeASS [29] is a tool which identifies performance degradation between two code versions. Similarly, PerfCI [14] is a toolchain for automated performance testing which allows developers to specify the number of test inputs to be used for performance testing. Our approach could be helpful in such scenarios where the cost of performance tests is quite high due to a large number of inputs.

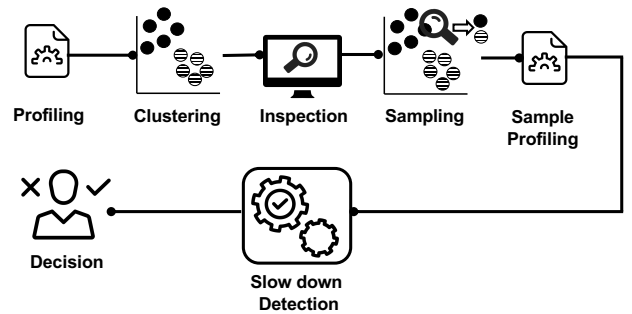## III. Approach for recommending performance tests

Fig. 3. Our approach for recommending performance test initiation

This section describes our approach for recommending performance test initiation. Our approach consists of several stages as shown in figure 3. Therefore, we explain different stages that work together to assist developers in making informed decisions about test initiation.

Our approach works on two different types of updates; 1) addition of test inputs and 2) code modifications. Profiling, clustering and inspection are three stages that are applied to cluster test inputs. On the other hand, sampling, sample profiling and slow down detection are applied on code updates.

A. Profiling

We apply static and dynamic analysis to observe program behavior and collect different program information such as execution time, memory usage, size of the input etc. In machine learning terminology, we refer to each program information (such as execution time, memory usage etc) as attributes.

A test suite consists of several test cases; each test case executes a part of the application code. Therefore, the test suite is executed repeatedly on different test inputs to exercise different section of the application code.

Therefore, profiling (or code analysis) is performed during test execution to apply unsupervised machine learning algorithms to group test inputs based on similar attributes. The selection of attributes (or program information) is based on our experience with performance analysis of CMS uploader service and from previous test case minimization studies [30]. We further explain each one of the attributes as follows:

- *Execution time:* The total time it takes for the test suite to complete execution on each test input.

- *Memory usage:* The memory used during the execution of the test suite on each test input.

- *Total number of iterations:* This represents the accumulated number of loop iterations executed in different sections of the code exercised by the test cases in a test suite on each test input.

- *Total statements executed:* Statements which are executed by the test suite on each test input.

- *Function calls:* Function calls made during the execution of the test suite on each test input.

- *Conditional executed:* Sequence of conditions taken in the execution of the test suite on each test input.

- *Input Size:* Input represent various aspects of the CMS detector's configuration. Therefore, the size of each input varies based on the information it contains about the detector configuration.

1) Minimizing measurement perturbations: One of the issues of profiling program's resources such as execution time is measurement perturbation. Instrumentation inserted in the code can perturbate measurements resulting in incorrect measurement of execution time [31]. To overcome this issue, we use PerfCI [14] to carry out our profiling task. PerfCI allows developers to write pluggable user-defined analysis for CI. It leverage stages [32] in a CI process to allow different profiling analysis to execute separately. Therefore, one can measure resources such as execution time in a different stage and other attributes can be collected in a separate stage. Furthermore, stages can be executed in parallel to quickly complete profiling and ensure that profiling itself does not lead to perturbations of profiling results.

B. Clustering

Identifying similar code behavior by clustering has been carried out in previous software testing studies [33]–[35]. However, our study addresses the issue of performance regression testing in growing test inputs. Furthermore, clustering of test inputs provides the benefit of gaining useful insights about the test inputs with respect to code behavior. This information would be useful for developers when selecting test inputs.

Reduction of test inputs is carried out by clustering data points collected from profiling. These data points represents program behavior on each test input. After clustering of test inputs, there is a stage called inspection which allows developer to check whether the clustering applied is appropriate.

C. Inspection

It is not always easy to identify clear clusters. The order in which data points are arranged may affect the clusters. Furthermore, another issue with identification of clear clusters is with data which has many missing values [36].

The arrangement of initial clusters is also important. A careful and comprehensive analysis of data is required. If the initial clusters are not carefully and properly chosen, then after some iterations, clusters may become empty [36]. Therefore, cluster inspection is necessary to ensure correct clustering is done. In this regard, different clustering algorithms and parameters are used for correct clustering. In appendix, we present our dashboard that allows inspection of clustering data.

We now explain sampling, sample profiling and slow down detection with the help of fig 4 along with algorithm 1 and 2. This makes the core decision making part of our approach.

D. Sampling

After clustering of inputs, our approach is prepared for identifying performance regression testing opportunities based on code updates. We use the number of executed statements and execution time as the two key attributes for decision making. Intuitively, these two attributes should exhibit correlation as executed statements has an impact on the execution time [37]. We further demonstrate the correlation of these attributes in the evaluation section.

Figure 4A shows selection of data points by random sampling. Sampling is done to collection minimum program information based on new code updates. Therefore, we kept the sampling size at a minimum e.g., we use 3 data points from each cluster. These data points contain information about specific test inputs which are used by sample profiling.

E. Sample profiling

In this stage, profiling is done with sampled test inputs on code update. Program information such as "number of executed statements" and "execution time" is collected in this
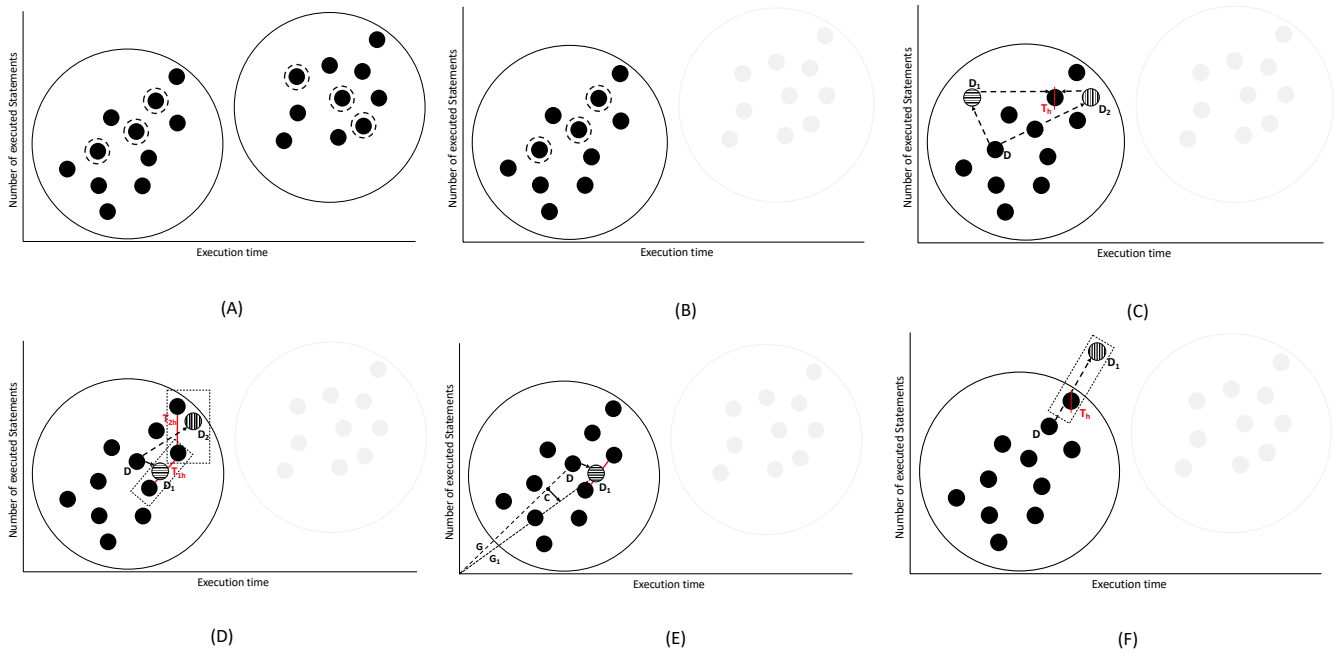
Fig. 4. Figure demonstrating different steps taken by our approach for recommending performance regression testing. Fig.3 A) shows sampling of datapoints in clusters, B) selects cluster (in normal black color) one at a time, C) & D) identification of time threshold, E) gradient measurement and F) handling of outlier datapoint.

Algorithm 1: Decision to Perform Testing

Decide(D_u:updated data point, T_h: time threshold,
C: current cluster, A_l: acceptable limit)

1. Initialize decision D to False
2. **IF** time information T_u in $D_u > T_h$
3.   Set D to True
4. **ENDIF**
5. **IF** D is not True
6.   Get time information T_p in previous data point from current cluster C
7.   Get no. of executed statement information S_p in previous data point from current cluster C
8.   Get time information T_u in D_u
9.   Get no. of executed statement information S_u in D_u
10.   Calculate gradient G on points $(0, T_p)$, $(0,S_p)$
11.   Calculate gradient G_1 on points $(0, T_u)$, $(0,S_u)$
12.   Calculate gradient change G_c between G and G_1
13.   **IF** $G_c > A_l$
14.   Set D to True
15.   **ENDIF**
16. **ENDIF**
17. Return D

---

stage. Further optimization can be achieve in sample profiling. Instead of running all tests on sampled test inputs, one can execute test cases which exercise application code that has been updated. This can be achieved by identifying difference between commits. A **git diff** command [38] can be used for

this purpose.

F. Slow down detection and decision making

After sampling and sample profiling stage, our algorithm maps new updated data points to their respective clusters. This mapping is done by selecting one cluster at a time as shown in figure 4B, where selected cluster is highlighted in black.

To detect a slow-down, our algorithm initially tries to find a time threshold value for the updated data point. Therefore, it attempts to detect the location of the updated data point in its cluster. We demonstrate this in figure 4C, if the "number of executed statements" attribute of the updated data point matches with the "number of executed statements" attribute in any previous data point. Our algorithm tries to find a data point that has the highest execution time" attribute among all data points sharing the same "number of executed statement" attribute. This will be adjudged as the time threshold (indicated as t_h) in figure 4C.

As shown in figure 4C, there are two cases where the updated data point can lie; one is that the "execution time" attribute is decreased which is indicated as D_1. In the second case, there is an increase in the "execution time" beyond the time threshold T_h as indicated by D_2 that is approximated as a slowdown. Therefore, a decision to initiate performance testing is made by the algorithm. This operation is shown on line 3-5 of algorithm 1.

On the other hand, if the "number of executed statement" attribute of the updated data point does not match with the "number of executed statement" attribute of any previous data point then time threshold is identified based on the location

Algorithm 2: Cluster data point evaluation

DataPointEvaluation(C_l: list of clusters, D_d: list of updated data points)

1. **FOR EACH** c in C_l
2.   **FOR EACH** d in D_d
3.   **IF** num. of executed statements s_e in d equal to s_e in c
4.   Get time information t for s_e in c.
5.   return Decide (d, t, c)
6.   **ELSE**
7.   Get data points above and below the datapoint d in cluster c
8.   **IF** datapoint above (or below) d does not exist
9.   Get max time information t_m for datapoint below (or above) in cluster c
10.   return Decide (d, t_m, c)
11.   **ELSE**
12.   Find max time information t_a of datapoint in c above d
13.   Find max time information t_b of datapoint in c below d
14.   Calculate mid-point t_{md} of t_a and t_b
15.   return Decide (d, t_{md}, c)
16.   **ENDIF**
16.   **ENDIF**
17.   **END FOR**
18. **END FOR**

of two data points which are above and below the updated data point. Moreover, data points (above and below) that have maximum "execution time" attribute are used to determine the time threshold; this is done by calculating the mid point of the "execution time" attribute of the data points above and below shown as T_{1_h} for D_1 or T_{2_h} for D_2 in figure 4D. After setting the time threshold, same conditions are applied to make a decision i.e., does the execution time attribute of updated data point cross the time threshold value as indicated for D_2 in figure 4D. This operation is accomplished in line 7,12-15 of algorithm 2.

Furthermore, figure 4D demonstrates a case where "execution time" attribute of the updated data point is close to the time threshold. Even though there is an increase in the execution time for the data point but algorithm will not decide in favor of initiation performance testing from this observation because the updated data point does not cross the time threshold. Hence missing out on a performance test opportunity.

To handle this issue, the gradient of the lines from the origin to either data points (i.e., original and updated) are measured. This is shown in figure 4E as G and G_1. To check whether there is a decrease of gradient from G to G_1, a change is measured which is indicated as C in figure 4E. This allows one to estimate the slow-down of the time between the original and updated data point. In section V, we demonstrate the advantage

for using gradient in making decisions for performance testing.

Our algorithm for slow-down detection also takes outliers into account i.e., the updated data point falls outside of the boundary of cluster. Depending on where the updated data point lies, the data point in the cluster closet to the updated data point is selected. For example, the data point having maximum execution time attribute below the outlier data point is selected as shown in figure 4F and the operation is shown on line 8-10 in algorithm 2.

Therefore, cluster data point evaluation algorithm (i.e., algorithm 2) identifies the position of updated data point with reference to original data point. The decision making is then done by Decision to Perform Tests algorithm. To make a decision, the algorithm will always first check whether time threshold condition is met line 2-3. On a false condition, it will measure gradients of the data points and identify whether the change is within the acceptable limit or not. Acceptable limit is the limit on slow-down i.e., how much slow-down is acceptable for the software. This parameter can be tuned by the developer.

IV. EVALUATION

The analysis of this study has been conducted on a SNIC science cloud [39] with Ubuntu Linux 16.04.4 LTS operating system. Furthermore, the configuration of the virtual machine is 4x 2.2 GHz vCPUs, 8 GB of RAM.

We use Gitlab CI pipeline at CERN [40] for executing test suite. The web service that we have used for our case study is CMS uploader service which is developed at CERN [8]. Furthermore, time measurements are taken an average of 3 repeated runs.

In our evaluation, we will begin by demonstrating the relationship between the variables (i.e., attributes) used for clustering. We used correlation test which is most widely used statistical measure to assess relationships among variables.

Secondly, we demonstrate minimization of test inputs by applying different clustering algorithms. Finally, we use two different code updates for demonstrating the feasibility of our recommendation approach.

A. Correlation between variables representing program behavior

In figure 5, we demonstrate a correlation between different variables that are selected by profiling. These variables represents the behavior of the code during the execution of test suite. The relationship in figure 5 is based on executing test suite with 4000 different test inputs. This is because addition of more test inputs was not changing the behavior of our analysis. This is demonstrated in figure 11.

We can see from the figure that 5 out of 7 variables (or attributes) demonstrates a positive correlation i.e., variables move in the same direction. However, there is also a negative correlation between two variables; this is between input size and executed conditions. This is quite intuitive since increasing input size should not have an effect on how conditionals are executed since conditionals are based on the evaluation of

Fig. 5. Correlation matrix between different variables; where labels such as **Size** represents Input size, **FunCalls** represents total function calls, **TExeStmt** represents total executed statements, **TNoItr** represents number of iterations, **ExeCond** represents executed conditionals, **Mem** represents memory used during execution of test suite, and **Time** represents the total execution time

expressions. Furthermore, we also see that other variables have negligible (or very low) correlation with executed conditionals.

Furthermore, input size demonstrates a strong correlation with function calls, executed statements, iterations as well as execution time. This means that increasing input size will increase the number of fuction calls, executed statements, number of loops and executed time of the test suite. On the other hand, there is a weak correlation between input size and memory usage.

This indicates that CMS uploader service is not a memory intensive application. We try to see whether the correlation are statistically significant by conducting a P-value test. This is shown in appendix-B.
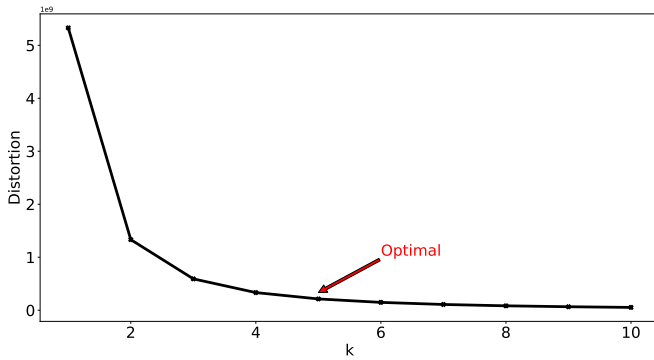

Fig. 6. Elbow method for determining the number of clusters in dataset

B. Identification of clusters

We use clustering to minimize inputs based on similar code behavior. We apply most popular distance-based clustering such as K-means [41]. Moreover, identification of the optimal number of clusters is done by a heuristic technique called the elbow method [42].

This method consists of plotting the variations (or distortions) representing how well the dataset is clustered by the
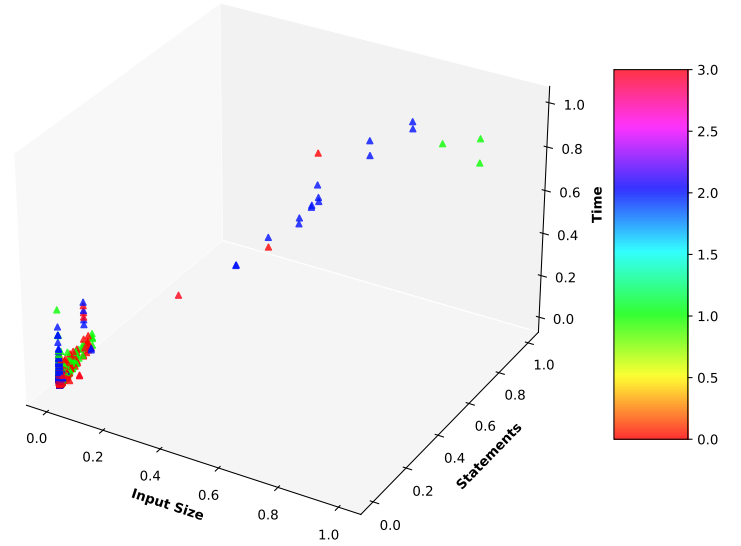

Fig. 7. Three dimensional view of clustering done by distance-based clustering algorithm

model. Therefore, the number of clusters is selected when the variations stop or the curve becomes smooth. Figure 6 shows the elbow method applied on our dataset. It can be seen that distortion is minimized at 5 which is shown with an red arrow in figure 6.

After identifying optimal number of clusters by elbow method the result of clustering by KMeans is shown in figure 7. The figure shows a 3D scattered plot based on 3 dimensions such as input Size, executed statements and execution time. It can be seen that proper clusters are not formed based on five clusters with Kmeans. This is because clustering in KMeans is an iterative process; user has to update the number of clusters iteratively. Furthermore, the data in figure 7 shows that the more appropriate number of clusters can be either 2 or 3. We will now apply different clustering
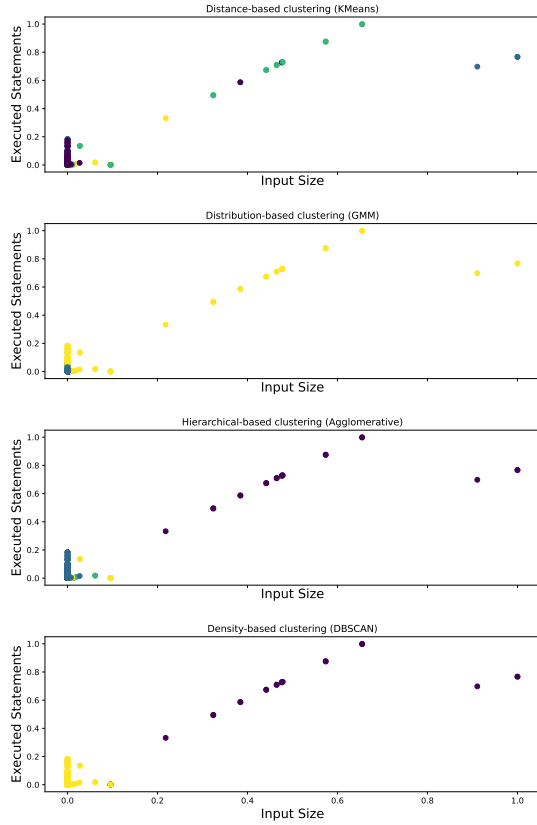
Fig. 8. Different Clustering algorithms demonstrating the number of clusters

algorithms and check which of the algorithm provides a better clustering.

C. Clustering based on different algorithms

Figure 8 shows clustering applied based on different types of clustering algorithm. This is shown with a scattered plot based on two features i.e., input size on x-axis and number of executed statements on y-axis.

We have applied 4 different clustering algorithms such as density-based (K-means), distribution-based (Gaussian mixture model), hierarchical (Agglomerative), message passing based (Affinity Propagation) and Density-based (DBSCAN). We also used Affinity propagation but we realized that it was not able to create clusters; it was reporting convergence issue. We found out that this is due to a bug which is reported in the issue page of scikit-learn repository [1].

Figure 8 gives a comparison of different clustering algorithms. It is evident that kmeans, GMM and agglomerative shows that clusters are being overlapped. The clustering done by DBSCAN algorithm seems more appropriate. It divides the clusters into two group.

1) Significance of cluster inspection: Our clustering analysis gives a good indication why we need to have an inspection in our approach. An incorrect number of cluster would not

[1] https://github.com/scikit-learn/scikit-learn/issues/17657

only affect our approach significantly but it will also not give a useful insight into our data.

With DBSCAN, the clustering behavior remains consistent compared to other clustering techniques. Furthermore, clustering done by DBSCAN shows that test inputs containing very low input size will have a program behavior where fewer statement are executed, fewer iteration will be made resulting in lower execution time. On the other hand, large input size will lead to high number of statements being executed resulting in higher execution time.

We will now demonstrate our decision making algorithm by selecting data points from each cluster and apply changes to the code.

V. RECOMMENDATION ON CODE CHANGE

In this section we demonstrate the feasibility of our decision making part of the approach which comprises of Algorithm 1 (decision to perform tests) and Algorithm 2 (cluster data point evaluation). We have described both algorithms and their working principle in section III.

We apply our algorithms on two code changes 1) A known bug in test suite which has been taken from previous study [14] and 2) code review based on code analysis tool — PyLint [43] which is used for error checking in the code.

This provides a reasonable evaluation of our decision making algorithm since one update already has a bug and the other update requires changes in different source files.

A. Code update 1: Bottleneck in test suite

There is a bottleneck in a constructor which occupies 99% of the total time taken by the unit test execution. This bottleneck slows down the overall total CI execution time by an average of 38%. More information about the performance bug can be found in the study [14].

The idea of our algorithm is to obtain minimal information to provide an informed decision about initiation of performance tests. Therefore, three data points are selected from each cluster. We also keep the acceptable limit of slow-down to 38% since this was the overall slow-down observed due to the performance bug in test suite.

TABLE I
DEMONSTRATION OF THE DECISION CHECKS FOR CODE UPDATE 1

Checks	C11	C12	C13	C21	C22	C23
Timethreshold	True	False	False	True	True	False
Gradient	x	True	False	x	x	True

Table I shows how the two decision checks provides their verdict on the slow down. In table I, C_{ij} refers to the cluster number and the selected inputs respectively. For example, C11 corresponds to the first cluster and first input. The cross (x) indicates that Gradient check was not required since the condition of time threshold holds true. It can be seen that 50% of updated data points (i.e., C12, C13 and C23) did not pass the conditions for time threshold. Gradient check adds a valuable asset in the decision making since it detects the rate

of slow down between previous data point and its updated one. For example, C12 and C23 has a rate of slow down of >50% (i.e., 66% for C12, and 99% for C23) which exceeds the acceptable limit of 38% that has been set. Hence, our algorithm gives positive decision to initiation a performance tests based on the decision checks of 5 out of 6 data points.

In case of C12, both checks did not pass, even though there was rate of slow down of 20%. Developers can tune the acceptable limit based on the need of the system. For example, in some applications, it may not be feasible to have a slow-down greater than 20%.

The recommendation for code update 1 is correct since there is a bottleneck in the code.

B. Code update 2: Code review with PyLint

PyLint is a static code analysis tool for used error checking. It is a popular tool for assuring code quality that gives an overall rating to the code. Furthermore, it is also used by many developers as well as big tech companies like Google [44]. Therefore, it is a feasible tool to use for demonstrating a code update in our case study.

After analyzing the project, PyLint reported around 78% of errors and warnings. Most of these errors can be referred to as cosmetic bugs e.g., errors related to the use of imports (Import checker messages [2]) or too many boolean expressions (Design Checker Messages). Therefore, we identified only those errors that are related to code refactoring. For example, use of enumerate instead of iterating with range and length functions.

We identified 7 places in the project were this error occurs and we updated the code by replacing range() and len() with enumerate(). The reason for selecting this type of ”error” is two-fold; firstly, we don't want to refactor the code too much such that it introduce additional errors, and secondly enumerate has an impact on the execution time.

Table II shows that none of the conditions for TimeThreshold and Gradient are met i.e., all the conditions are False. Furthermore, we also observed the rate of slow-down remained under 0.5%.

TABLE II
DEMONSTRATION OF THE DECISION CHECKS FOR CODE UPDATE 2

Checks	C11	C12	C13	C21	C22	C23
Timethreshold	False	False	False	False	False	False
Gradient	False	False	False	False	False	False

For code update 2, the recommendation is to skip performance testing.

C. Importance of data point sampling

Developers can be skeptical about the use of recommendation for bug identification since a miss opportunity of a performance bug detection can lead to dire consequences and their fix can be quiet time-consuming and costly [45]. In this

[2]https://vald-phoenix.github.io/pylint-errors/

regard, developers can face a situation where decision is 45-55% i.e., 45% of the data points are in favor of initiating performance test and 55% are in favor of skipping the tests.

To increase developers confidence, our approach allows developers to increase the sample size of the data points during sampling stage (section III-D), and see if the decision changes i.e., more data points are in favor of initiating performance tests. However, there is a trade-off in increasing the sample size. Every time a new sample is used; a quick test is need to be run. Therefore, it is important to select sample size wisely when there is a doubt in recommendation.

VI. CONCLUSION

Software performance testing is a crucial software quality assurance mechanism. However, performance testing is a time-consuming process since it requires a set of inputs to exercise different code sections. Therefore, running performance tests on every code update makes it challenging for developers because the time it takes for testing to complete. Furthermore, this becomes impractical in organizations that deals with (ever) growing data i.e., more test inputs are added to test the code. To overcome these difficulties, we propose an approach that recommends developers when to run or skip performance tests by using minimal information from the code. We demonstrate the feasibility of our approach by applying it to two code updates and we found that our approach makes sound recommendations.
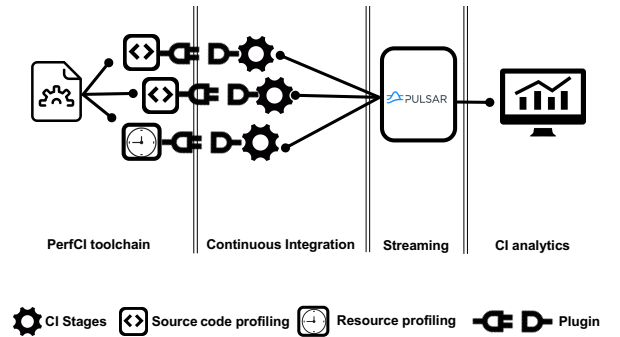
APPENDIX A
CI ANALYTICS SYSTEM

Fig. 9. Overview of CI analytic system

Figure 9 shows the overall architecture of our CI analytic system. There are four key stages of the architecture such as PerfCI toolchain, Continuous Integration, Streaming and Analytic. We briefly explain each of these four stages.

Stage 1 (PerfCI toolchain): We built our system on top of our previous work — PerfCI [14] which allows developers to write pluggable user-defined analysis for Continuous Integration. The structure of PerfCI analysis is shown in the following code fragment. Developer implement three methods which are in the abstract class **InstrumentEvents**. The analysis is added to the test suite by instrumenting test cases and rewriting the CI configuration file.

```
class ResourceCollector(IntrumentEvents):
    def start_measurement(self):
        ....

    def end_measurement(self):
        ....

    def record_data(self, *args):
        ....
```

Stage 2 (Continuous Integration): When the CI process runs, the $start\_measurement()$ and $end\_measurement()$ methods are executed on per unit test basis which gathers the resource information such as execution time and memory utilization. Furthermore, code-level metrics such as number of function calls are also gathered based on the PerfCI analysis defined by the user.

Stage 3 (Streaming): One of the problems with Continuous Integration services is that it does not allow the storage of large volume of data (i.e., maximum storage size is 1GB). This is only allowed for limited time duration i.e., it will delete the storage information after a week. To overcome this issue, we use streaming framework such as Apache pulsar [46] which streams our profiling data and offloads it out of the continous integration service to a remote server. This provides two advantages 1) one can perform real-time analysis and 2) large volumes of data can be handled. Furthermore, we use Apache pulsar because of its multi-tenancy feature which allows for increase scalability, cost reduction and improve security. This is quite crucial for organization like CERN.

Stage 4 (Analytics): The analytic system has a front-end which is implemented in plotly dash []. This allows the development of interactive visualization. The main advantage of dashboard is that the developers can visualize the clustering data i.e., inspection. Once the developer is satisfied with the clustering information. S/he can apply sampling, sample profiling and slow down detection to evaluate whether performance testing is required. This is one of the main advantage of our CI analytic system.

We now highlight some important features of our CI analytic system's dashboard. These are labeled in numbers in figure 10. We briefly explain each of these numeric labels as follows:

1. Different code updates. These are displayed as commit hash. This allows developers to observe changes in the test data on each update.

2. Execution time of each unit tests. Different unit test's execution time can be compared.

3. Total execution time of test suite. This gives a quick overview of the performance of the test at each update.

4. Memory utilized during the execution of each test case.

5. To identify correlation between different variables.
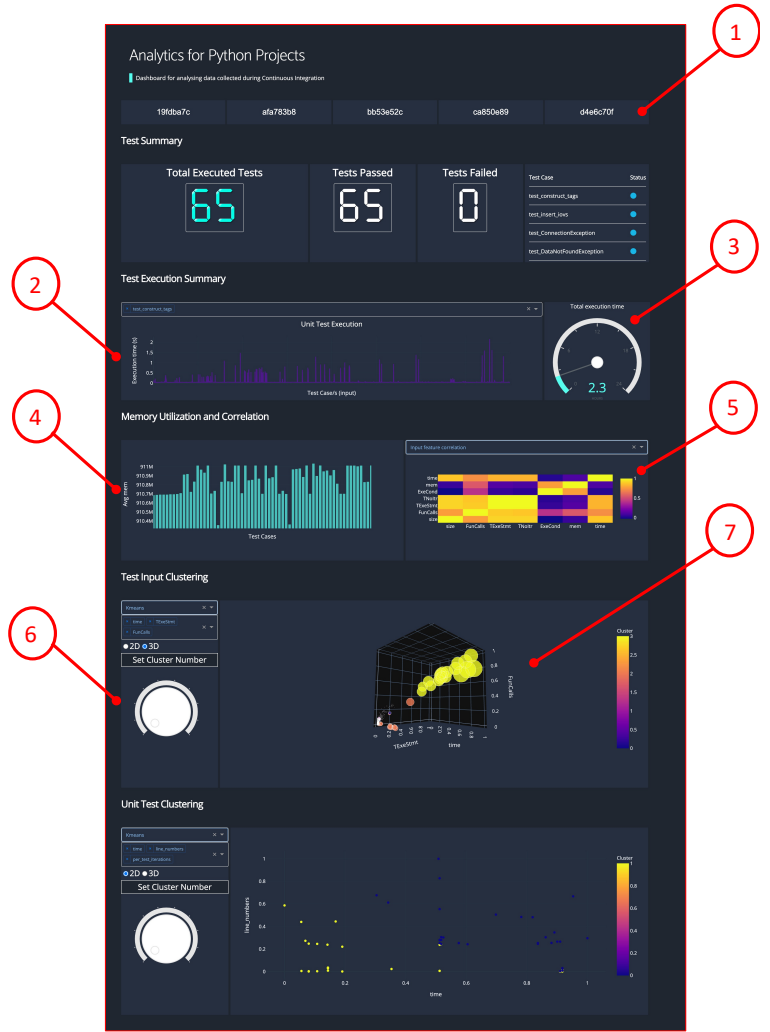

Fig. 10. Analytic Dashboard

6. Cluster Panel which provides different options such as selection of different clustering algorithm, changing visualization between 2D and 3D. This is an important feature of our dashboard that is used for inspection.

7. Three-dimensional view of clustering data. This can be used by developers if there is difficulty in identifying clear clusters.

APPENDIX B
P-VALUE TEST

Figure 11 shows the trend of P-values over 5 different updates to test inputs. This test is applied on the variables whose correlation is presented in section IV-A. On x-axis are the five commits i.e., updates to inputs and Y-axis shows the P-values. The title represents the two variables whose p-values are being calculated. The trends in diagonal (starting from top-left) should be ignored as they represent statistical significance between same two variables or attributes. Variables which demonstrate high correlation also has high stable statistical
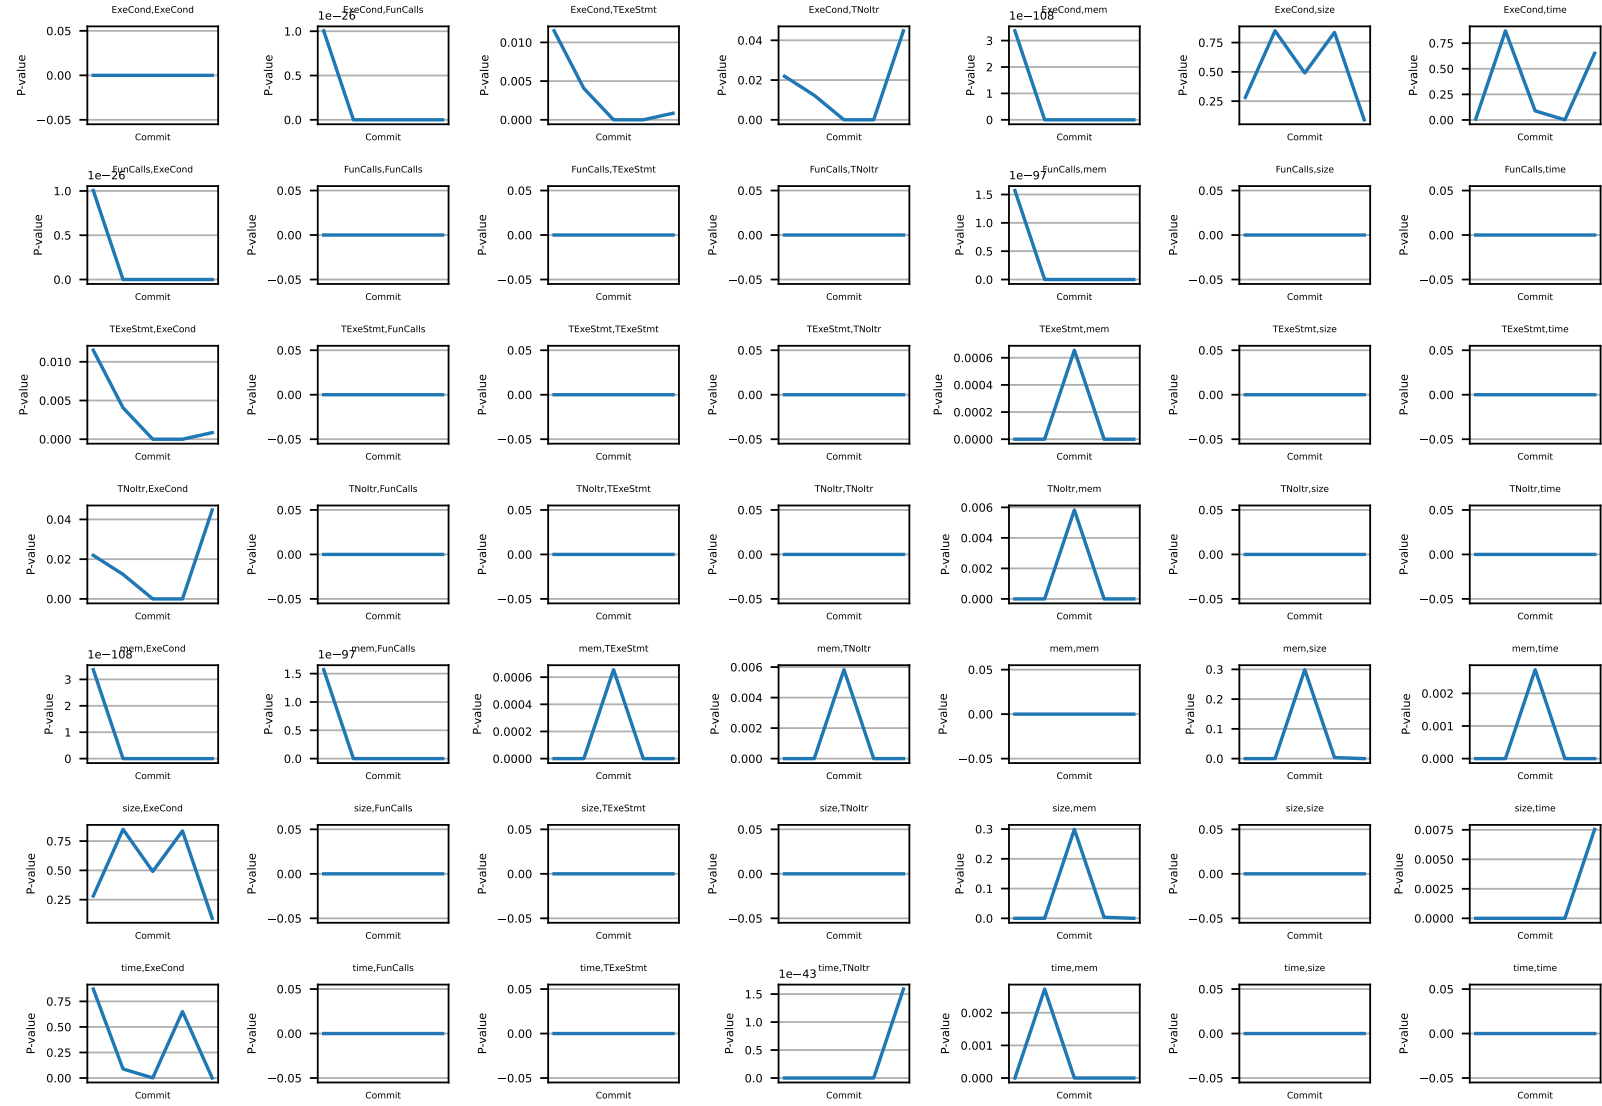
Fig. 11. Statistical signficant test (P-value) for variables in evolving inputs.

significance. Furthermore, significance of time and number of iterations change in the last commit, but the value is still well close to zero (i.e., the value is in the order of $1.5e-43$). Therefore, there is high statistical significance for the variables representing program behavior, even with increasing the number of test inputs. This analysis provides confidence about the variable or attributes that we have used. We will now minimize the number of test inputs by grouping attributes that represents program behavior based on test input.

REFERENCES

[1] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, "Understanding and detecting real-world performance bugs," in *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 77–88. [Online]. Available: https://doi.org/10.1145/2254064.2254075

[2] T. Register, https://www.theregister.co.uk/2002/07/03/1901_census_site_still_down/, 2002, (Accessed on 05/25/2020).

[3] P. Stefan, V. Horky, L. Bulej, and P. Tuma, "Unit testing performance in java projects: Are we there yet?" in *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*, ser. ICPE '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 401–412. [Online]. Available: https://doi.org/10.1145/3030207.3030226

[4] X. Han, T. Yu, and D. Lo, "Perflearner: Learning from bug reports to understand and generate performance test frames," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, ser. ASE 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 17–28. [Online]. Available: https://doi.org/10.1145/3238147.3238204

[5] "Storage — cern," https://home.cern/science/computing/storage, (Accessed on 10/07/2021).

[6] A. B. de Oliveira, S. Fischmeister, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Perphecy: Performance regression test selection made simple but effective," in *2017 IEEE International Conference on Software Testing, Verification and Validation, ICST 2017, Tokyo, Japan, March 13-17, 2017*. IEEE Computer Society, 2017, pp. 103–113. [Online]. Available: https://doi.org/10.1109/ICST.2017.17

[7] D. Alshoaibi, K. Hannigan, H. Gupta, and M. W. Mkaouer, "Price: Detection of performance regression introducing code changes using static and dynamic metrics," in *Search-Based Software Engineering*, S. Nejati and G. Gay, Eds. Cham: Springer International Publishing, 2019, pp. 75–88.

[8] J. Dawes, "A python object-oriented framework for the cms alignment and calibration data," *Journal of Physics: Conference Series*, vol. 898, p. 042059, 10 2017.

[9] A. B. De Oliveira, S. Fischmeister, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Perphecy: Performance regression test selection made simple but effective," in *2017 IEEE International Conference on Software Testing, Verification and Validation (ICST)*, 2017, pp. 103–113.

[10] "Jesperstromblad/slowdowndetection: Implementation of slowdown detection algorithm," https://github.com/JesperStromblad/Slowdowndetection, (Accessed on 11/09/2021).

[11] "Jesperstromblad/ci-analytics-system: Continuous integration analytics system for python-based projects," https://github.com/JesperStromblad/ci-analytics-system, (Accessed on 11/09/2021).

[12] C. Collaboration, S. Chatrchyan, G. Hmayakyan, V. Khachatryan, A. Sirunyan, W. Adam, T. Bauer, T. Bergauer, H. Bergauer, M. Dragicevic *et al.*, "The cms experiment at the cern lhc," *JInst*, vol. 3, p. S08004, 2008.

[13] J. H. Dawes, G. Reger, G. Franzoni, A. Pfeiffer, and G. Govi, "Vypr2: A framework for runtime verification of python web services," in *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part II*, ser. Lecture Notes in Computer Science, T. Vojnar and L. Zhang, Eds., vol. 11428. Springer, 2019, pp. 98–114. [Online]. Available: https://doi.org/10.1007/978-3-030-17465-1_6

[14] O. Javed, J. H. Dawes, M. Han, G. Franzoni, A. Pfeiffer, G. Reger, and W. Binder, "Perfci: A toolchain for automated performance testing during continuous integration of python projects," in *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 1344–1348.

[15] J. Chen, W. Shang, and E. Shihab, "Perfjit: Test-level just-in-time prediction for performance regression introducing commits," *IEEE Transactions on Software Engineering*, pp. 1–1, 2020.

[16] J. P. Sandoval Alcocer, A. Bergel, and M. T. Valente, "Learning from source code history to identify performance failures," in *Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering*, ser. ICPE '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 37–48. [Online]. Available: https://doi.org/10.1145/2851553.2851571

[17] S. Mostafa, X. Wang, and T. Xie, "Perfranker: Prioritization of performance regression tests for collection-intensive software," in *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 23–34. [Online]. Available: https://doi.org/10.1145/3092703.3092725

[18] R. Hashemian, N. Carlsson, D. Krishnamurthy, and M. Arlitt, "Iris: Iterative and intelligent experiment selection," in *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*, ser. ICPE '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 143–154. [Online]. Available: https://doi.org/10.1145/3030207.3030225

[19] D. Westermann, R. Krebs, and J. Happe, "Efficient experiment selection in automated software performance evaluations," in *Computer Performance Engineering*, N. Thomas, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 325–339.

[20] Q. Luo, "Input-sensitive performance testing," in *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. FSE 2016. New York, NY, USA: Association for Computing Machinery, 2016, p. 1085–1087. [Online]. Available: https://doi.org/10.1145/2950290.2983953

[21] Y. Feng, Q. Shi, X. Gao, J. Wan, C. Fang, and Z. Chen, "Deepgini: prioritizing massive tests to enhance the robustness of deep neural networks," in *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, USA, July 18-22, 2020*, S. Khurshid and C. S. Pasareanu, Eds. ACM, 2020, pp. 177–188. [Online]. Available: https://doi.org/10.1145/3395363.3397357

[22] Z. Li, X. Ma, C. Xu, C. Cao, J. Xu, and J. Lü, "Boosting operational dnn testing efficiency through conditioning," in *Proceedings*

of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ser. ESEC/FSE 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 499–509. [Online]. Available: https://doi.org/10.1145/3338906.3338930

[23] "Start a free trial with circleci - circleci," https://circleci.com/enterprise-trial-install/?utm_source=google&utm_medium=sem&utm_campaign=sem-google-dg--emea-en-brandAuth-maxConv-auth-brand&utm_term=g_e-circleci_c__linux_20210604&utm_content=sem-google-dg--emea-en-brandAuth-maxConv-auth-brand_keyword-text_eta-circleCI_exact-&gclid=Cj0KCQjwtrSLBhCLARIsACh6RmglTWkvMbNvM5N3i1pNo09_lBGYnOBVPvJPz_GUeYwjfHGxVaOafjsaAn4TEALw_wcB, (Accessed on 10/18/2021).

[24] "Travis ci — test and deploy with confidence," https://www.travis-ci.com/, (Accessed on 10/18/2021).

[25] "Cloudbees codeship — software as a service (saas) ci/cd solution," https://www.cloudbees.com/products/codeship, (Accessed on 10/18/2021).

[26] "Bitbucket pipelines - continuous delivery — bitbucket," https://bitbucket.org/product/features/pipelines, (Accessed on 10/18/2021).

[27] "Continuous integration & delivery - semaphore," https://semaphoreci.com/, (Accessed on 10/18/2021).

[28] "performance-plugin — performance test running and reporting for jenkins ci," http://jenkinsci.github.io/performance-plugin/RunTests.html, (Accessed on 10/18/2021).

[29] D. G. Reichelt, S. Kühne, and W. Hasselbring, "Peass: A tool for identifying performance changes at code level," in *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2019, pp. 1146–1149.

[30] R. Singh and M. Santosh, "Test case minimization techniques: a review," *International Journal of Engineering Research & Technology (IJERT)*, vol. 2, no. 12, 2013.

[31] M. Hauswirth, P. F. Sweeney, A. Diwan, and M. Hind, "Vertical profiling: Understanding the behavior of object-priented applications," *SIGPLAN Not.*, vol. 39, no. 10, p. 251–269, Oct. 2004. [Online]. Available: https://doi.org/10.1145/1035292.1028998

[32] "Ci/cd pipelines — gitlab," https://docs.gitlab.com/ee/ci/pipelines/, (Accessed on 10/19/2021).

[33] C. Xia, Y. Zhang, and Z. Hui, "Test suite reduction via evolutionary clustering," *IEEE Access*, vol. 9, pp. 28 111–28 121, 2021.

[34] N. Chetouane, F. Wotawa, H. Felbinger, and M. Nica, "On using k-means clustering for test suite reduction," in *2020 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*, 2020, pp. 380–385.

[35] A. Vescan and C. Şerban, "Towards a new test case prioritization approach based on fuzzy clustering analysis," in *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2020, pp. 786–788.

[36] P. Agarwal, M. A. Alam, and R. Biswas, "Issues, challenges and tools of clustering algorithms," *arXiv preprint arXiv:1110.2610*, 2011.

[37] J. Erikson, P. Funk, J. Gustafsson, and B. Lisper, "A tool concept for execution time analysis of legacy systems," in *14th Euromicro Conference on Real-Time Systems, Work-in-Progress*. Citeseer, 2002.

[38] "Git - git-diff documentation," https://git-scm.com/docs/git-diff, (Accessed on 11/04/2021).

[39] P. number SNIC 2021/18-7, "Uppsala – essence," http://essenceofescience.se/tag/uppsala/, (Accessed on 10/20/2021).

[40] CERN, "Gitlab pipeline," http://gitlab.cern.ch/, (Accessed on 10/20/2021).

[41] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *Journal of the royal statistical society. series c (applied statistics)*, vol. 28, no. 1, pp. 100–108, 1979.

[42] P. Bholowalia and A. Kumar, "Ebk-means: A clustering technique based on elbow method and k-means in wsn," *International Journal of Computer Applications*, vol. 105, no. 9, 2014.

[43] "Pylint - code analysis for python — www.pylint.org," https://pylint.org/, (Accessed on 10/21/2021).

[44] "Interview: Claudiu popa - using pylint for python static analysis - sqreen blog," https://blog.sqreen.com/interview-pylint-for-python-static-analysis/, (Accessed on 10/21/2021).

[45] S. Zaman, B. Adams, and A. E. Hassan, "Security versus performance bugs: A case study on firefox," in *Proceedings of the 8th Working*

Conference on Mining Software Repositories, ser. MSR '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 93–102. [Online]. Available: https://doi.org/10.1145/1985441.1985457

[46] "Apache pulsar," https://pulsar.apache.org/en/, (Accessed on 11/09/2021).