The Use of Change Point Detection to Identify Software Performance Regressions in a Continuous Integration System

David Daly
MongoDB Inc
david.daly@mongodb.com

William Brown[**]
Columbia University
w.brown@columbia.edu

Henrik Ingo
MongoDB Inc
henrik.ingo@mongodb.com

Jim O'Leary
MongoDB Inc
jim.oleary@mongodb.com

David Bradford
MongoDB Inc
david.bradford@mongodb.com

ABSTRACT

We describe our process for automatic detection of performance changes for a software product in the presence of noise. A large collection of tests run periodically as changes to our software product are committed to our source repository, and we would like to identify the commits responsible for performance regressions. Previously, we relied on manual inspection of time series graphs to identify significant changes. That was later replaced with a threshold-based detection system, but neither system was sufficient for finding changes in performance in a timely manner. This work describes our recent implementation of a change point detection system built upon the E-Divisive means [15] algorithm. The algorithm produces a list of change points representing significant changes from a given history of performance results. A human reviews the list of change points for actionable changes, which are then triaged for further inspection. Using change point detection has had a dramatic impact on our ability to detect performance changes. Quantitatively, it has dramatically dropped our false positive rate for performance changes, while qualitatively it has made the entire performance evaluation process easier, more productive (ex. catching smaller regressions), and more timely.

CCS CONCEPTS

• **Software and its engineering** → **Software performance**; • **Information systems** → *Database performance evaluation*; • **Mathematics of computing** → *Time series analysis*.

KEYWORDS

change points, performance, testing, continuous integration

[*]Work done in part as an intern at MongoDB Inc.

1 INTRODUCTION

We work in a software and services company and need to understand the performance of the software we develop and sell. Our continuous integration system runs thousands of benchmarks periodically (most commonly every 2 hours or every 24 hours), which each produce one or more scalar values as a result. It is a challenge to analyze all of those results and historical data to determine whether the test should be considered passed or failed. It is inherent to this type of benchmarking that the results contain some level of noise. In most of our tests the worst case run to run variation level is currently less than 10%, but some sensitive tests can fluctuate as much as 20% or more over the course of a small numbers of runs. In this paper, we detail the results of our experience in deploying automated change point detection software for identifying performance changes in an evolving code base.

1.1 Existing performance results and analysis

Our performance testing infrastructure is built upon our continuous integration (CI) system: Evergreen [4, 12]. Evergreen tracks every commit to our source repository, compiles the software on a number of different build targets, and runs correctness tests.

Evergreen also runs our performance tests. The performance tests take longer to run than the correctness tests, so we run them less frequently, but otherwise the performance and correctness tests are the same from the perspective of Evergreen.

With any performance test, we must

- Setup a system under test
- Run a workload against the system under test
- Report the results from the test
- Decide (and alert) if the performance changed
- Visualize the results

This paper focuses on the last two bullets[1]. When we first setup our performance tests in our CI system, we had a visualization of the performance over time, but no automated alerting. We had a person (called the "Build Baron") looking through the graphs for regressions, and opening Jira tickets[2] for any changes requiring further investigation. The Build Baron would be looking at performance trend graphs like those shown in Figure 1 and Figure 2.

---

[1]Previous work focused on the first two bullets and getting reproducible results [13]
[2]Atlassian's Jira is our ticketing system.

Figure 1 shows a test with some basic noise, but no real changes in performance, while Figure 2 has two very clear changes in performance. The cases in Figures 1 and 2 should be easy for the Build Baron to triage, but many other cases are not, such as the two graphs in Figure 3 and the graph in Figure 4, which have smaller changes and larger run to run variation. We quickly realized that having humans triaging performance changes solely on the trend graphs was not a tenable solution: humans looking through all the graphs lose focus quickly and the problem would get worse as we added more tests and system configurations. Therefore we added an automated detection system.
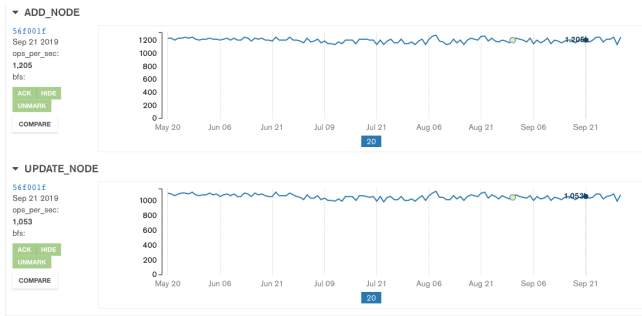

Figure 1: Performance trend graph with noise and no discernible changes.

Figure 2: Example of a clear drop in performance on August 8th and subsequent performance recovery on August 20th.
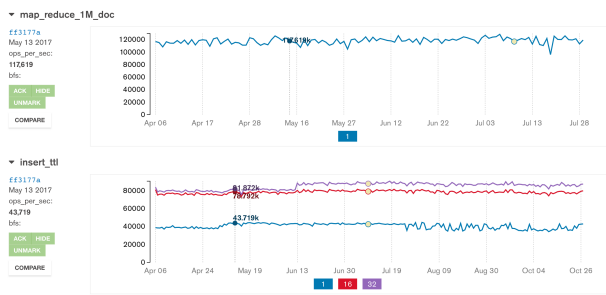

Figure 3: Performance trend graphs with noise. The second graph (insert_ttl) has a change in performance on June 13th smaller than the noise threshold of the first graph. The amount of noise also changes over time for the insert_ttl graph.

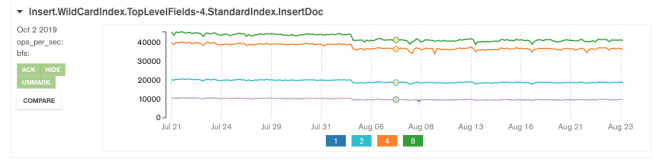The automated detection system was designed to catch

Figure 4: Performance trend graph with a small drop in performance on August 5th. It would be easy for a person to miss that regression when looking through many graphs.

- Sudden drops in performance from a specific software change
- Multiple small changes over time causing a drift down in performance

It did this in a very simple manner by doing three comparisons:

(1) Compare to the immediately preceding run. Did it change more than $X\%$[3]?
(2) Compare to a run from a week ago. Did it change more than $X\%$?
(3) Compare to a run for the last stable release of the software. Did it change more than $X\%$?

This system was a vast improvement over just staring at graphs. It was automated and computers do not get bored or lose focus. However, we soon realized we had a large number of false positives with which to deal (up to 99% as discussed in Section 5 depending on how you count). All the tests have some inherent noise in them. However, different tests and different test configurations have different levels and types of noise. Using a static threshold (e.g., 10%) led to a lot of alerts for changes on noisy tests. At the same time it would either miss real changes that were less than the threshold, or detect the smaller changes at some later time (when the change plus the noise crossed the threshold). For example, in Figure 3 the insert_ttl test shows a clear, but very small change around June 13th, while the map_reduce workload is steady with run to run variation (noise). There is no way we could adjust a common threshold for those two tests that would both catch the change in insert_ttl and not constantly flag the map_reduce test, or constantly flag the insert_ttl 1 thread line after July 19th.

While the system was noisy and painful, it was still useful for us. Over time we added fixes and band aids to the system to reduce the number of false positives (e.g., adjust the thresholds separately for different tests). Even with all those fixes to the previous system, the system:

- Produced a large number of false positive alerts that had to be processed
- Missed smaller regressions (false negatives)
- Flagged failures on commits unrelated to the regression, requiring humans to trace back to find when a regression actually happened
- Only flagged regressions, not improvements
- Consumed a huge amount of time from the team to process all the results

We wanted to improve the process to something fundamentally better. We started by trying to clearly state the problem.

---

[3]Where a human then looked at flagged failures and X is a parameter of the system. We used 10% by default

2 PROBLEM DEFINITION

The software we are testing is updated in discrete commits to our source repository. Our problem is to:

Problem 1. *Detect which commits change the performance of the software (as measured by our performance tests) in the presence of the noise from the testing infrastructure.*

When the problem is phrased like that, it is clear that this is a signal processing problem. The noise of the system makes the performance results a stochastic process. As such, we reviewed the literature on signal processing[4]. In particular, we decided our problem was a change point detection problem. From Matteson and James [15]: "Change point analysis is the process of detecting distributional changes within time-ordered observations." Our system can be represented as the time series:

$$S_t = P_t + N_t \tag{1}$$

Where S_t is the measured signal, P_t is a constant value of performance, and N_t is a random variable representing the noise. It is not clear a priori that the noise component is independent, identically distributed (IID) or not. In fact, we have examples in which the noise is correlated with time.

The noise is really composed of at least three components:

$$N = N_p + N_s + N_w \tag{2}$$

- N_p: The noise from the test platform, including CPU, network, OS, etc.
- N_s: The noise from the software itself. This can be from the compiler as well as other non-determinism or randomness in the software.
- N_w: The noise from the workload generator. The workload generator is also software and may vary in the load it applies to the system.

Reducing N_p and N_w was the focus of previous work [13]. However, as computers and computer networks are inherently very complicated, non-deterministic machines, including things like caches and predictors, we can only hope to reduce those noise components, not remove them. As such, we have to learn to live with noise.

2.1 E-Divisive Means

After reviewing the literature on change point detection, we decided to use the E-Divisive means algorithm [15]. E-Divisive means has the following useful properties for us:

- Does not require any distributional assumptions other than a finite mean
- Can be recursively applied to identify multiple change points
- Does not require training data and works out of the box on a time series
- Optimal sets of change points can be efficiently computed via dynamic programming and permutation sampling

Additionally, the main focus of our work is on retroactive identification of change points, not prediction. Many time series algorithms are intended for forecasting, which is not relevant in the context of commits being pushed by developers.

The algorithm works by hierarchically selecting distributional change points that divide the time series into clusters. For a given cluster, a candidate change point is chosen by computing a statistic $\widehat{Q}$ at every point within the cluster and selecting the point with the largest value. This statistic $\widehat{Q}$ is the discrete analog of the divergence between continuous multivariate distributions. After the first change point is determined, the algorithm is then reapplied to each cluster created by the existing change point. This process repeats until the largest candidate $\widehat{Q}$ value is below a threshold of significance.

2.1.1 $\widehat{Q}$ Statistic Computation. To compute $\widehat{Q}$ for a series of n + m points bisected at the nth point, let:

$$\mathbf{X}_n = \{X_i : i = 1, ..., n\} \tag{3}$$
$$\mathbf{Y}_m = \{Y_j : j = 1, ..., m\} \tag{4}$$

Represent the first n points and the last m points, respectively. An empirical divergence metric can be computed as follows:

$$
\begin{aligned}
\widehat{\mathcal{E}}(\mathbf{X}_n, \mathbf{Y}_m; \alpha) = & \frac{2}{mn} \sum_{i=1}^{n} \sum_{j=1}^{m} |X_i - Y_i|^\alpha \\
& - \binom{n}{2}^{-1} \sum_{1 \le i < k \le n} |X_i - X_k|^\alpha \\
& - \binom{m}{2}^{-1} \sum_{1 \le j < k \le m} |Y_j - Y_k|^\alpha
\end{aligned} \tag{5}
$$

The α parameter can be any value between 0 and 2. We use 1 for simplicity[5]. We then attain $\widehat{Q}$ by weighting the previous result by the size of our clusters:

$$\widehat{Q}(\mathbf{X}_n, \mathbf{Y}_m; \alpha) = \frac{mn}{m + n} \widehat{\mathcal{E}}(\mathbf{X}_n, \mathbf{Y}_m; \alpha) \tag{6}$$

2.1.2 Termination. The procedure for termination in the paper involves numerous random permutations of the values in each cluster to determine the statistical significance. Intuitively, if rearranging the order of data in a cluster does not significantly affect the maximum $\widehat{Q}$, there is likely no meaningful change point.

3 IMPLEMENTATION

Selecting an algorithm for detecting change points was just the beginning. We then needed to:

- Implement (or find an implementation of the algorithm)
- Automate processing our data with the algorithm
- Present the data after processing
- Enable users to act on the presented data

The first step of implementing the algorithm was straightforward as there is an R package [14] for the algorithm. We also wrote a Python implementation[6] to go with our existing Python code base. Both implementations will iteratively find the most significant change point in the series, and then analyze the clusters on either side of the change point, until reaching the stopping criterion. Even

---

[4]See Section 7 for references on signal processing and change point detection

[5]Matteson and James [?] observed similar results with α of 0.5 and 1.5, and presented their simulation studies with α of 1.
[6]The Python implementation of e-divisive means is available here: https://github.com/mongodb/signal-processing-algorithms

though we used our Python implementation, we benefited from the existence of the R implementation. We were able to compare the results from the two implementations to gain confidence in the correctness of our implementation. It also gave us a reference point for the performance of the algorithm.

It took a little more work to transform our existing data into a form that could be consumed by the algorithm. We store results in a MongoDB database, with all the results for a given task build stored together. We unbundled the data so we could easily generate a vector of results for each individual combination of (system variant, task, test, test configurations). We then saved the computed change points in the database in a separate collection. Each change point has the same identifying fields as the results data, plus data from the algorithm itself, such as the $\widehat{Q}$ and the order of the change point (was it the first or N^{th} change point for the series when recursively identifying change points). After identifying the change points for a series, each series is partitioned into clusters in which the results are unchanging. We call those clusters "stable regions". We calculate common statistics for each stable region, and include that in the entry for each change point. Statistics include the min, max, median, mean, variance for the given measure, as well as how many data points are in the stable region.

The algorithm is run at the completion of every set of performance tests. The result data is uploaded and transformed before the algorithm is run. We replace any existing change points with the newly generated change points. Rerunning the analysis allows the change points to be updated as more data becomes available. The data is statistical in nature, so having more data allows better analysis. As such, it is expected that on occasion a previously computed change point will "move" to a new commit or even completely go away as more data becomes available. Additionally, the change point algorithm will never find that newest result to be a change point. If a regression was introduced, the algorithm will only detect a change point after several more test results are collected. A three to five day delay is common for our tests. This contrasts with processing results for correctness tests in which the pass/fail result is determined at run time and will not change at a later time.

3.1 Display

We then displayed the data in two distinct ways. First, we updated the trend graphs discussed in Section 1.1. We added annotations to show the change points as well as any Jira tickets that match the test and revision. Figures 5 and 6 show two trend graphs annotated with both change points and Jira tickets. The green diamonds are the Jira tickets, and the highlighted segments are change points. You can clearly see two tickets matching two green highlighted sections in Figure 5, while there is a small drop with change point and ticket highlighted in Figure 6. That figure is particularly interesting as the change in mean is on the same scale as the level of the noise before and after the change, and a version of that graph without annotations appeared earlier in Figure 4. As we will discuss in Section 3.2, it is a goal of our process that each change point is matched to a Jira ticket or hidden from view if it is clear noise.

We added necessary meta data fields to our Jira tickets so that they can be mapped to their originating Evergreen project, task, test and git revisions. We track both "First failing revision" and "Fix

Figure 5: Trend graph for index build background test. Notice two change points (green highlights) and two JIRA tickets (diamonds).
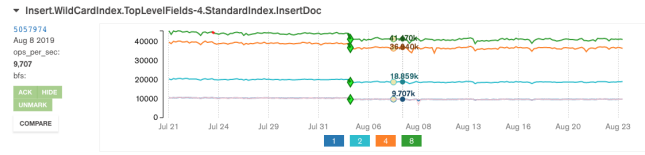

Figure 6: Trend graph for the wildcard index test with annotations for Jira tickets and change points. This figure can be compared to Figure 4 which does not have the annotations.

revision". Additionally, we needed to query this data every time we drew the displays. For performance reasons, we setup a regular batch job to query Jira and load the data into our database.

Note that we do not run the performance tests on every commit to the source repository. That adds complexity to the visualization above. A change point is identified for commits that were built and run. The algorithm cannot determine which commit caused a performance change if the commits have not run. Instead, we can identify the range of commits between the change point and the previous test result as suspect commits. When displaying these results and annotations, we must be able to match up change point, Jira ticket, and builds that have actually run, even if they point to close but different commits.

Those annotations are great for examining how one test performs over time, and identifying the location of significant changes. It is not useful though for tracking all tests, triaging the most important regressions, and grouping tests that change at the same time. For that, we built a new page for viewing the change points. This new page shows all the change points on one page, grouped by commit and is shown in Figures 7 and 8. Change points are initially "unprocessed" and through triage become "processed".

Figure 7 shows the list of software revisions (identified by their git hashes) associated with "unprocessed" change points. All the change points for a single git hash are grouped together. The display has the git hash, test name, task name, date of the build, and the hazard level of the change. Hazard level is computed as the log of the ratio of mean before to the mean after the change point. The hazard level has proven extremely useful in focusing triage on the major regressions and it has the property that two changes that cancel out (e.g., 33% drop followed by a 50% rise) have the same magnitude. The display has the ability to filter or sort on any of the fields. There are two buttons along the top to acknowledge or hide the change points. Hiding or acknowledging makes a copy of the changed point with the addition of field to indicate that the change point has been processed. The process of acknowledging or hiding

Figure 7: Triage view of unprocessed change points.

Figure 8: View of processed change points[7].

a change point also removes it from the unprocessed change point page. It will also be removed from this page if there is already a Jira ticket for the same test at that commit. The list is short in the figure because the remaining change points had already been "processed", with only two less interesting change points remaining.

Once a commit is acknowledged or hidden, it shows up on the processed tab as shown in Figure 8. There we have shown one change point with a significant improvement in performance. We will isolate that change to a single commit, and open a Jira ticket to document it. Once that ticket is opened, the change point will also be removed from this view.

3.2 Triage

We have a rotating role of a dedicated person to check for performance changes. We call this person the "Build Baron". The Build Baron uses all the views of the change points shown in Figures 7 and 8 along with the trend graphs as shown in Figures 5 and 6. Ideally, the Build Baron reviews every change point, deciding if it is a real change worthy of investigation, if it is insignificant, or if it is just noise.

The Build Baron starts with the list of the unprocessed change points as shown in Figure 7. By default we have the list set to filter out some things (lower priority, canary[8] tests, or some results that are merely informational). The Build Baron goes down this list investigating each git hash in the list. Each git hash can have multiple change points (multiples tests, variants, etc). The test column includes links to the task page with the matching trend graph for the Build Baron to decide if the change point is worth further investigation.

They may determine from inspection that there is a clear change. If so, they will try to isolate the offending commit. Recall that we do

not run every commit, so once a performance change is detected, we need to run the tests on the missing commits to determine the precise commit that introduced the performance change. The Build Baron will schedule the builds appropriately, and "acknowledge" the change point by pressing the "acknowledge" button on the change point page. The Build Baron will need to wait for the tests to finish. When the tests finish, the system will rerun the change point detection algorithm in order to update the change point to the correct commit. If the change point represents a new issue, the Build Baron will open a Jira ticket for the change and assign it to the appropriate team to investigate. If there is an existing Jira ticket (e.g., this is the fix to a previous regression) the Build Baron will update the existing Jira ticket. If the performance change is a significant enough regression, we may also revert the commit.

Not all change points lead to Jira tickets. Occasionally we will have spurious drops in performance, with the performance changing for one build and then returning to previous performance. Drops such as these may go away if we rerun the task. If the performance change is large enough it can lead to a change point being created. We start by checking these against our canary workloads. As part of previous noise reduction work we have a number of canary workloads. The canary workloads tell us nothing about the software we are testing, but a lot about the test system. The results should be constant. If we detect a statistically significant change for a canary or otherwise suspect a spurious drop, we treat all the results from that test run as suspect and rerun the test.

Alternatively, the change point may be due to noise. The algorithm is very good at accounting for IID noise, but not for correlated noise. Unfortunately, we have some correlated noise in our system – some from time varying behavior on our cloud providers, other due to things as finicky as code alignment issues in compilation and its impact on processor caches. Either case usually can be confirmed visually, and the Build Baron will "Hide" the change point by using the Hide button. That will remove the change point from the change point page and change the color of the change point on the trend

---

[7]Our Best Buy tests are based off of the open data set from Best Buy [1]

[8]Canary tests are tests run to detect changes to the test system itself. Changes in performance for canary tests indicate a problem with the test itself. See [13] for more on our use of canary tests.

graph to blue. If we have a sustained change in performance due to cloud changes, we open a Jira ticket for that in order to document it. That way, anyone looking at the trend graphs will be able to see why the performance changed.

All change points that are processed (hidden or acknowledged) and all change points that match a Jira ticket are removed from the change point list. The list is meant for triage and those points have already been triaged. We have the history of the points however and can use it to compile effectiveness data about our system.

Since we group change points by git hash in the display, we may have the case in which we want to acknowledge some of the change points in the group and hide others. This is fully supported.

4 OPTIMIZATIONS

The computational complexity of a naive implementation of E-Divisive means is $O(\kappa T^3)$, where κ is the number of change points and T the number of observations in the series. Matteson and James [15] point out that it is possible to compute the series $\widehat{Q}$ in $O(T^2)$ (per change point) as each value in the series can be derived from the previous with linear time. We implemented this optimization in our POC.

The major generations of the implementation were:

- Naive $O(\kappa T^3)$ implementation
- $O(\kappa T^2)$ implementation (row/column differences)
- Use NumPy arrays
- Use native C implementation (Python module)

We also tested compiling the first 2 versions with Cython, but this yielded only about 10% improvements and was discarded as an option.

With a test data series of 173 values, using an Intel(R) Core(TM) i7-6560U CPU @ 2.20GHz, single threaded executions yielded the following results:

Implementation	Time (s)	Ratio
Naive	0.622551	6103.44
Row/column differences	0.011567	113.40
NumPy	0.001282	12.57
Native	0.000102	1.00

At this modest problem size, the native C implementation is over 6,000x faster than our original naive implementation, 113x faster than a non-naive implementation, and 13x faster than our best Python only implementation.

At the end of a CI task there are 20-60 new results. The historical time series for those is fetched and the algorithm is used to re-compute change points. The computation is easy to parallelize as there are multiple series to be processed. Using the native C implementation and parallelizing the computation, we are able to re-compute all the change points in under 10 seconds. The majority of that time is spent waiting for network communication to the database, with the E-Divisive mean computation taking only milliseconds. The computation performance is a combination of the performance enhancements discussed above, and a conscious decision to limit our analysis to results from N recent months, by using up to 500 data points[9]. We do this since we are interested in

---

[9]This is approximate. If more than 500 data points are available, we will select data points going back to an already computed change beyond that range. As such, the computation will use more than 500 points.

catching regressions in new commits and new releases. Due to this, our time series typically have 100 to 500 values and 1 to 10 change points, and the computation time is minimal.

In our original project plan we were concerned about the $O(\kappa T^3)$ complexity and had envisioned running change point detection in a separate asynchronous job so as to not impact test execution time. Due to the above optimizations we are able to run the change point detection immediately after each test execution in a simpler process. When the Build Baron is running additional tests to isolate the exact commit that changed performance, having the change point detection immediately refreshed reduces the time to isolate the change.

Even though the computation time is minimal in production, we still care about the execution time. Occasionally we need (or want) to recompute all the change points in our system (e.g., we updated something related to the computation or we fixed a bug). Based on the current performance of the algorithm, we believe we could recompute all the change points across our system in under an hour. Previously we had a need to regenerate change points for a section of the data and we tried to regenerate the points over the weekend. The job was still running Monday morning when we checked on it. When we tried to regenerate the data more recently for the same subset, it took approximately 30 minutes. Finally, we may want to extend the length of history we use in the computation in the future, for example if we add a project that runs much more frequently.

5 IMPACT

The point of this work, as discussed in Section 2, was to more efficiently detect which commits change the performance of our software. It has succeeded. Triaging performance changes has gone from being a full time job that one person could not completely keep up with, to one that the Build Baron can process all of the change points and have time left over. There is a morale improvement, as the Build Baron is working on real performance changes that make a difference to the company. Additionally, we are able to detect smaller changes in performance than we could previously.

We are also tracking performance improvements now in addition to performance regressions. This has had a few impacts:

- We have caught correctness bugs (e.g., important check suddenly skipped).
- We have been able to provide positive feedback to developers for making things faster. Someone notices the impact of their work now.
- We can share improvements with marketing [19].

While we clearly and strongly know that we have qualitatively improved our system, we have also quantitatively improved it. We have two sets of data: First, we performed a proof of concept (POC) before fully implementing the system. During the POC we looked at false negative and false positive rates, as well as general feasibility related issues. Since then we have implemented the system in full, including collecting the labeling data and the information in Jira tickets. The combination of that data allows us also to measure the effectiveness of the system in production.

5.1 Learnings from the POC

For the POC we implemented the calculation of change points as a batch job. We then compared the changes points to the results from our existing system and in Jira tickets. We focused on a 5 month period of results. At that time our automated system used two level of Jira tickets. The first level was automatically generated by the analysis code. The Build Baron then would go through that list. The Build Baron would move representative failures to the second Jira project, and link the rest. For the 5 month period we had 2393 tickets in the first project, and 160 in the second. Of the 160 tickets in the second project, 24 of them were "useful."[10] So, on average 100 tickets in the first project reduced down to 1 useful ticket in the second project, with a human going through all of those. In other words, the old system generated a huge number of Jira tickets, of which the vast majority were false positives, while also suffering from false negatives.

The change point algorithm found all the regressions we were tracking in Jira tickets. The change point algorithm may have missed performance changes, but if it did those were performance changes we were already missing. From a false negative rate, the new system was strictly better than the old system.

We also looked at the change points that did not match existing tickets. We started by sorting on the $\widehat{Q}$ value. The first three that we investigated led to two new Jira tickets. The first element was one ticket, while the second and third items shared a ticket. The second item corresponded to a drop in performance and and the third item was a later fix for the earlier drop.

Since we did not have complete reference data for all the performance data, we instead sampled the change points to estimate a false positive rate. We considered only the change points that did not match existing Jira tickets, as those are the potential false positives. We looked at 10 of those change points. Of those:

- 2 were clear changes in performance we had not caught previously.
- 4 were clearly noise.
- 4 were not clear cut. They included
  - One clear change in distribution (the test got noisier)
  - Two cases in which the distribution changed slightly. We think this was likely correlated noise in our system.
  - One large outlier value got the system to create a change point.

Of the cases that were noise we used the order of change point statistic to search for more significant regressions. In all cases there was a point in which change points went from noise or maybe regressions, to clear cut changes.

Depending on how you count the four "not clear cut" cases, we had a false positive rate between 40% and 80%, with the possibility to tune the system. Given the previous system required looking at 100 things for each useful one, this was a huge improvement and we knew we wanted to move forward.

5.2 Learnings from Production

After the POC we put the change point detection algorithm into production and have been collecting data. Looking at a 3 month

---

period from summer of 2019 for our primary performance project we have:

- 12,321 distinct change points
- Corresponding to 178 unique git hash values
- 126 of those 178 correspond to a Jira ticket
- 122 of those 178 were acknowledged or hidden
- there were 79 unique Jira tickets for those change points

Therefore, the Build Baron had 178 things to look at during this period, and more than 2/3 of all of them were issues we felt were worth tracking and possibly working on. We went from 100 notifications leading to 1 tracked issue, to 9 notifications leading to 6 worth follow-up and 4 tracked issues. Given that we want the process to be slightly biased toward false positives, this result is close to optimal.

The process still is not perfect though. We also have:

- Acknowledged change points that do not match a raw change point
- Acknowledged change points without a matching Jira ticket
- Hidden change points associated with Jira tickets

The two cases of acknowledged change points not matching a Jira ticket or a raw change point appear to come from a few cases related to triaging: aggressive triaging and mistaken triaging. The change points are grouped by git hash. It is easy to check some of the change points for the git hash and observe a distinct performance change, and aggressively acknowledge all the points. Alternatively, sometimes someone hits the wrong button and acknowledges when they probably should have hidden a point. We observe both of those cases when inspecting points for these two cases. We are adding more auditing of how we label change points.

Hidden change points that do not match a raw change point mainly come from a combination of marginal change points aggressively triaged, in which the change point algorithm is able to exclude once it has additional data.

6 OPEN PROBLEMS

This work is part of a larger effort to understand the performance of our software, and to know when the fundamental performance changes. This builds on work to

- Build tools for making performance workloads
- Automate running of performance tests
- Include external benchmarks in our system
- Reduce the noise in our test bed

There are a number of areas we continue to work on in this space

- Better workload generation tools
- More flexible result types
- Better signaling of performance changes
- Determine if a "patch" has changed performance
- Identify, exclude, and rerun test runs in which "canaries" have failed
- Use the data from processing change points to make the algorithm more effective

Better workload generation tools is about making it easier for our engineers to craft interesting workloads, while the more flexible result types deal with handling latencies, percentiles, throughputs, etc.

The next three items show the limits of the change point detection algorithm. E-Divisive means is great at detecting changes in distribution, but it needs multiple data points to compute that distribution. It also has a bias towards the center of the longest clusters. As such, it cannot tell you anything about the most recent test run, and it will not detect two distinct changes if they are too close together, such as in Figure 9. In Figure 9 a performance impacting code change went in and was reverted soon afterwards. There are not enough commits for the algorithm to detect that two distinct changes happened.
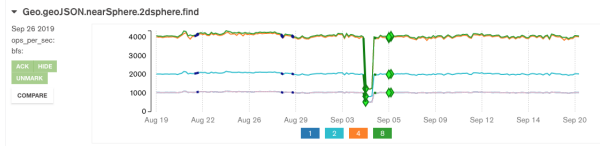

Figure 9: Example of a drop in performance followed quickly by a fix. The regression and fix are too close together for the E-Divisive means algorithm to identify both changes.

The bias towards the center of clusters is a result of the coefficient $mn/(m + n)$ in Equation 6. The coefficient was an integral assumption in the consistency proof in [15]. Since $m + n = T$, where T is the number of observations in the series, we get $mn/(m + n) = m(T - m)/T = m - m2/T$. The first and last values of this coefficient are therefore $1 - 1/T$ when either m or n equals 1. All other values of this coefficient are greater than one (e.g. $2 - 4/T$ for $m = 2$) and the coefficient achieves the maximum value of $T/4$ when $m = n = T/2$. Even if it would lead to possibly invalidating the mathematical proof behind E-Divisive, in practice it is tempting to explore whether these limitations could be removed.

The algorithm has a parameter for the minimum cluster size of observations: N. The tests used later in [15] used minimum cluster sizes N=30 and N=10, while we have operated with N=3 to have faster notification. With N=3 we require at least 3 data points after the change (including the change) to detect the change. The combined effect of these limitations is that in our primary performance project we typically get alerts of a new regression 5-6 days after it was committed.

Ultimately change point detection is not suitable for determining whether a single new test result is a regression or not. We can use outlier detection to answer that question. The change point detection algorithm can identify a stable region, and then we can use outlier detection to determine if the patch build is statistically the same as the stable region or not. This will allow us to react early to very clear regressions, while change point detection can catch smaller regressions after a delay. We also have a use case in which developers can test (patch test) their code changes (patches) before pushing them to the source repository. We can use the same outlier detection to detect performance changes in patch tests. Finally, with our canary tests the goal is reversed: Rather than ignoring outliers, we want to be alerted about outliers so that we can rerun the tests or quarantine the results.

As mentioned in Section 3.2, we run a number of canary tests to detect changes to the test system itself. If we have a sudden change in the canary workloads we would like to flag the results and possibly rerun the tests. The detection of performance chances in the canary workloads, flagging the results, and rerunning suspect executions are all amenable to automation.

Finally, we are collecting a large corpus of training data related to the change point algorithm, through the acknowledged and hidden change points. This will allow us to compare the existing algorithm to other change point detection algorithms. Additionally, we expect we can use this data to either manually adjust the algorithm, or train more advanced machine learning algorithms.

7 RELATED WORK

Previous work has focused on reducing the noise in our test environment [13]. Our test environment is built on the Evergreen [12] CI system. There are many other continuous integration systems, such as Jenkins [17], BuildBot [9], Travis CI [8], Gitlab CI [5], and Circle CI [3]. The other CI systems have various supports for graphing performance results. For instance, Jenkins has a performance plugin [6] to support running common benchmarking tools, reporting results, and plotting trend graphs.

The Chromium project has implemented its own CI system called LUCI [7] and implemented regression detection [2] within it. Each test requires setting explicit threshold values for alerting a regression. As such, it is very similar in concept to our previous analysis system.

Google has a patent filing [18] for an algorithm called "Window deviation analyzer" (WDA). It describes identifying windows of data in a time series, computing median and mean for each window, and estimating if the most recent sample differs more than some threshold from those values. Like our system, Google's WDA will fail a test some number of commits after the commit that introduced a regression. Unlike our system, WDA requires an explicit threshold and it is does not identify a specific offending commit (it identifies a window of commits within which a regression has occurred).

For algorithms for dealing with noise in time series data, we found the NIST/Sematech e-handbook on Statistical Methods [16] a good overview, particularly chapter 6 "Process and Product Monitoring and Control". We looked into using techniques such as Moving Averages or Box Jenkins [11] for the time series data. That would involve building a model of the process and detecting when samples diverge from the existing model.

Additionally, see [10] for an extensive survey on change point detection techniques in time series data, covering supervised methods such as multi-class, binary, and virtual classifiers, and unsupervised methods such as likelihood ratio, subspace models, and probabilistic models. There are frequentist/Bayesian and parametric/non-parametric methods algorithms for both the supervised and unsupervised classes.

8 CONCLUSION

In this paper we describe how we detect performance changes in our software using our continuous integration system and change point detection. We run a large collection of tests periodically as changes to our software product are committed to our source repository. We would like to know when the performance changes for those tests. Originally we used humans looking and graphs, which was later replaced with a threshold based automatic detection system, but

neither system was sufficient for finding changes in performance in a timely manner. This paper describes our move to using change point detection (specifically E-Divisive means) to detect and notify when the performance changes.

In addition to implementing the change point algorithm, we also had to integrate it into our existing performance testing system and visualize the results so that engineers could triage and use the information. We describe what was required to get change point detection working in our system and how we triage and process the results. We also describe our efforts to speed up the algorithm so that we could use it synchronously within our workflow.

The net impact of this work was large for us: Quantitatively, it dramatically dropped our false positive rate for performance changes, while qualitatively it made the entire process easier, more productive (ex. catching smaller regressions), and more timely.

There is more work to be done to continue to improve our ability to detect when the performance of our software changes. We discuss a number of open problems for performance testing in general and for using change point detection in particular.

REFERENCES

[1] [n.d.]. Best Buy APIs open data set. https://bestbuyapis.github.io/api-documentation/#overview
[2] [n.d.]. (Chromium) Regression Detection for Performance Tests. wiki. https://www.chromium.org/chromium-os/testing/perf-regression-detection
[3] [n.d.]. Circle CI. https://circleci.com/docs/
[4] [n.d.]. Evergreen Continuous Integration. https://github.com/evergreen-ci/evergreen/wiki
[5] [n.d.]. Gitlab CI/CD. https://docs.gitlab.com/ee/ci/
[6] [n.d.]. Jenkins Performance Plugin. wiki. https://wiki.jenkins.io/display/JENKINS/Performance+Plugin
[7] [n.d.]. (LUCI) A Tour of Continuous Integration UI. https://chromium.googlesource.com/chromium/src.git/+/master/docs/tour_of_luci_ui.md
[8] [n.d.]. Travis CI. https://docs.travis-ci.com
[9] Ruth Ablett, Ehud Sharlin, Frank Maurer, Jorg Denzinger, and Craig Schock. 2007. Buildbot: Robotic monitoring of agile software development teams. In *RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 931–936.
[10] Samaneh Aminikhanghahi and Diane J Cook. 2017. A survey of methods for time series change point detection. *Knowledge and information systems* 51, 2 (2017), 339–367.
[11] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. 2015. *Time series analysis: forecasting and control.* John Wiley & Sons.
[12] Kyle Erf. 2016. Evergreen Continuous Integration: Why We Reinvented The Wheel. Blog Post. https://engineering.mongodb.com/post/evergreen-continuous-integration-why-we-reinvented-the-wheel
[13] Henrik Ingo and David Daly. 2019. Reducing variability in performance tests on EC2: Setup and Key Results. Blog Post. https://engineering.mongodb.com/post/reducing-variability-in-performance-tests-on-ec2-setup-and-key-results
[14] Nicholas James and David Matteson. 2015. ecp: An R Package for Nonparametric Multiple Change Point Analysis of Multivariate Data. *Journal of Statistical Software, Articles* 62, 7 (2015), 1–25. https://doi.org/10.18637/jss.v062.i07
[15] David S. Matteson and Nicholas A. James. 2014. A Nonparametric Approach for Multiple Change Point Analysis of Multivariate Data. *J. Amer. Statist. Assoc.* 109, 505 (2014), 334–345. https://doi.org/10.1080/01621459.2013.849605 arXiv:https://doi.org/10.1080/01621459.2013.849605
[16] NIST/SEMATECH. 2012. e-Handbook of Statistical Methods. http://www.itl.nist.gov/div898/handbook/ note.
[17] John Ferguson Smart. 2011. *Jenkins: The Definitive Guide: Continuous Integration for the Masses.* " O'Reilly Media, Inc.".
[18] Anthony Vallone. 2018. Window Deviation Analyzer. Patent Publication 2018/0150373 A1, Filed Nov. 28, 2016, Published May 31, 2018.
[19] DJ Walker-Morgan. 2019. MongoDB 4.2 Performance Boosts. Blog. https://www.mongodb.com/blog/post/mongodb-42-performance-boosts