

Dynamically Reconfiguring Software Microbenchmarks: Reducing Execution Time without Sacrificing Result Quality

Christoph Laaber
University of Zurich
Zurich, Switzerland
laaber@ifi.uzh.ch

Harald C. Gall
University of Zurich
Zurich, Switzerland
gall@ifi.uzh.ch

Stefan Würsten
University of Zurich
Zurich, Switzerland
stefan.wuersten@uzh.ch

Philipp Leitner
Chalmers | University of Gothenburg
Gothenburg, Sweden
philipp.leitner@chalmers.se

ABSTRACT

Executing software microbenchmarks, a form of small-scale performance tests predominantly used for libraries and frameworks, is a costly endeavor. Full benchmark suites take up to multiple hours or days to execute, rendering frequent checks, e.g., as part of continuous integration (CI), infeasible. However, altering benchmark configurations to reduce execution time without considering the impact on result quality can lead to benchmark results that are not representative of the software’s true performance.

We propose the first technique to dynamically stop software microbenchmark executions when their results are sufficiently stable. Our approach implements three statistical stoppage criteria and is capable of reducing Java Microbenchmark Harness (JMH) suite execution times by 48.4% to 86.0%. At the same time it retains the same result quality for 78.8% to 87.6% of the benchmarks, compared to executing the suite for the default duration.

The proposed approach does not require developers to manually craft custom benchmark configurations; instead, it provides automated mechanisms for dynamic reconfiguration. Hence, making dynamic reconfiguration highly effective and efficient, potentially paving the way to inclusion of JMH microbenchmarks in CI.

CCS CONCEPTS

• **General and reference** → **Measurement; Performance; • Software and its engineering** → **Software performance; Software testing and debugging.**

KEYWORDS

performance testing, software benchmarking, JMH, configuration

ACM Reference Format:

Christoph Laaber, Stefan Würsten, Harald C. Gall, and Philipp Leitner. 2020. Dynamically Reconfiguring Software Microbenchmarks: Reducing Execution Time without Sacrificing Result Quality. In *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium*

ESEC/FSE '20, November 8–13, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '20), November 8–13, 2020, Virtual Event, USA*, <https://doi.org/10.1145/3368089.3409683>.

on the Foundations of Software Engineering (ESEC/FSE '20), November 8–13, 2020, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3368089.3409683>

1 INTRODUCTION

Performance testing enables automated assessment of software performance in the hope of catching degradations, such as slowdowns, in a timely manner. A variety of techniques exist, spanning from system-scale (e.g., load testing) to method or statement level, such as software microbenchmarking. For functional testing, CI has been a revelation, where (unit) tests are regularly executed to detect functional regressions as early as possible [22]. However, performance testing is not yet standard CI practice, although there would be a need for it [6, 36]. A major reason for not running performance tests on every commit is their long runtimes, often consuming multiple hours to days [24, 26, 32].

To lower the time spent in performance testing activities, previous research applied techniques to select which commits to test [24, 45] or which tests to run [3, 14], to prioritize tests that are more likely to expose slowdowns [39], and to stop load tests once they become repetitive [1, 2] or do not improve result accuracy [20]. However, none of these approaches are tailored to and consider characteristics of software microbenchmarks *and* enable running full benchmark suites, reduce the overall runtime, while still maintaining the same result quality.

In this paper, we present the first approach to dynamically, i.e., during execution, decide when to stop the execution of software microbenchmarks. Our approach—dynamic reconfiguration—determines at different checkpoints whether a benchmark execution is stable and if more executions are unlikely to improve the result accuracy. It builds on the concepts introduced by He et al. [20], applies them to software microbenchmarks, and generalizes the approach for any kind of stoppage criteria.

To evaluate whether dynamic reconfiguration enables reducing execution time without sacrificing quality, we perform an experimental evaluation on ten Java open-source software (OSS) projects with benchmark suite sizes between 16 and 995 individual benchmarks, ranging from 4.31 to 191.81 hours. Our empirical evaluation comprises of three different stoppage criteria, including the one from He et al. [20]. It assesses whether benchmarks executed with dynamic reconfiguration in controlled, bare-metal environments

(1) maintain their result quality and (2) have shorter execution times, compared to being executed with the default JMH configuration.

We find that for the majority of studied benchmarks the result quality remains the same after applying our approach. Depending on the stoppage criteria, between 78.8% and 87.6% of the benchmarks *do not* produce different results, with an average performance change rate between 1.4% and 3.1%. Even though computation of the stoppage criteria introduces an overhead between <1% and ~11%, dynamic reconfiguration still enables saving a total of 66.2% to 82% of the execution time across all projects. For individual projects, benchmark suites take 48.4% to 86.0% less time to execute. Our empirical results support that dynamic reconfiguration of software microbenchmarks is highly effective and efficient in reducing execution time without sacrificing result quality.

Contributions. The main contributions of our study are:

- We present the first approach to dynamically stop the execution of software microbenchmark using three different stoppage criteria.
- We provide empirical evidence that demonstrates the effectiveness and efficiency of dynamic reconfiguration for ten OSS applications.
- We provide a fork of JMH that implements our approach on Github [34] and as part of our replication package [35].
- To investigate whether real-world benchmark suites could benefit from our approach to save time, we collect the largest data set of JMH OSS projects (753 projects with 13,387 benchmarks) including extracted source code properties such as benchmark configurations and parameters.

2 JAVA MICROBENCHMARK HARNESS (JMH)

JMH¹ is the de-facto standard framework for writing and executing software microbenchmarks (in the following simply called benchmarks) for Java. Benchmarks operate on the same level of granularity as unit tests, i.e., statement/method level, and are similarly defined in code and configured through annotations. Different from unit tests where the outcome is binary, i.e., a test passes or fails (disregarding flakiness), benchmarks produce outputs for a certain performance metric, such as execution time or throughput. As these performance metrics are easily affected by confounding factors, such as the computer’s hardware and software, background process, or even temperature, one must execute benchmarks repeatedly to obtain rigorous results that are representative of the software’s true performance [18].

Figure 1 depicts a standard execution of a JMH benchmark suite B , where benchmarks b are sequentially scheduled. Every benchmark execution starts with a number of warmup forks wf , to bring the system into a steady state, whose results are discarded. A fork is JMH parlance for running a set of measurements in a fresh Java Virtual Machine (JVM). The warmup forks are followed by a number of measurement forks f (often simply called forks). Due to dynamic compilation, every fork is brought into steady state by running a series of warmup iterations wi , after which a series of measurement iterations mi are executed. An iteration has a specific duration – wt or mt for warmup time and measurement time, respectively–

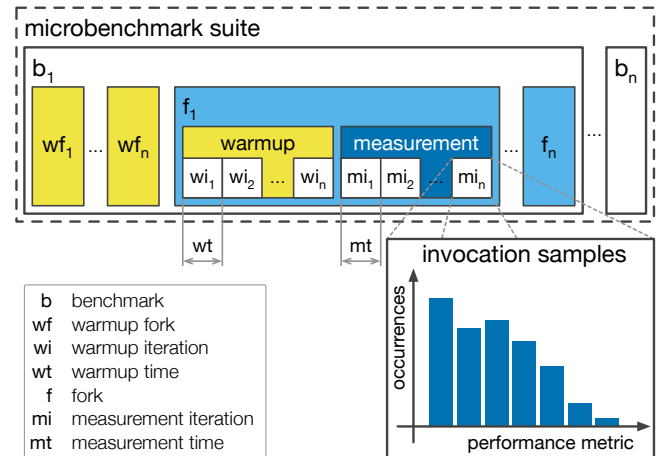


Figure 1: JMH Execution

for which the benchmark is executed as often as possible, and the performance metrics for a sample of the invocations is reported. Performance metrics from warmup iterations are discarded, and the union of the measurement iterations across all forks form the benchmark’s result. All these values can be configured by the developer through JMH annotations or the command line interface (CLI), otherwise default values are used.

JMH supports benchmark fixtures, i.e., setup and teardown methods, as well as parameterization of benchmarks. A parameterized benchmark has a number of parameters with (potentially) multiple values; JMH then runs the benchmark once for every parameter combination, which are formed as the cartesian product of the individual parameters.

JMH uses different sets of default configuration values, depending on the version: ≤ 1.20 and ≥ 1.21 . Versions until 1.20 use 10 forks (f) running 40 iterations (20 wi and mi each) with an iteration time (wt and mt) of 1s; starting from 1.21, defaults are 5 forks (f), 5 iterations (both wi and mi), and 10s iteration time (both wt and mt) [47, 48]. JMH does not use warmup forks (wf) by default.

Consequently, and as Fig. 1 depicts, we can define the overall warmup time as $t_w^b = wf * (wi * wt + mi * mt) + f * wi * wt$, the overall measurement time as $t_m^b = f * mi * mt$, and the benchmark execution time as $t^b = t_w^b + t_m^b + t_{fix}^b$, where t_{fix}^b is the time spent in benchmark fixtures. Finally, the full microbenchmark suite execution time T is the sum of all benchmark parameter combinations, defined as $T = \sum_{b \in B} \sum_{p \in P^b} t^{b,p}$, where P^b the set of parameter combinations for a benchmark b . These definitions will be used in the remainder of the paper.

3 PRE-STUDY

To motivate our work, we conduct a pre-study investigating whether benchmark execution times are in fact a problem in real-world OSS projects using JMH.

3.1 Data Collection

We create, to the best of our knowledge, the most extensive OSS JMH data set to date from Github, by querying and combining

¹<https://openjdk.java.net/projects/code-tools/jmh> (incl. source code examples)

three sources: (1) Google BigQuery’s most recent Github snapshot², queried for `org.openjdk.jmh import` statements [12, 32]; (2) Github’s search application programming interface (API) with an approach as outlined by Stefan et al. [49]; and (3) Maven Central searching for projects with JMH as dependency. Our final dataset consists of 753 pre-study subjects after removing duplicate entries, repositories that do not exist anymore, projects without benchmarks in the most recent commit, and forked projects.

For each project, we apply the tool *bencher* [30] to construct abstract syntax trees (ASTs) for Java source-code files and extract information related to (1) execution configuration (@Fork, @Warmup, @Measurement, and @BenchmarkMode) and (2) benchmark parameterization (@Param). In addition, (3) we extract the JMH version from the build script (*Maven* and *gradle*).

3.2 Summary of Pre-Study Subjects


The 753 projects have in total 13,387 benchmark methods and 48,107 parameter combinations. 400 (53.1%) projects feature fewer than 10 benchmarks, and 52 (6.9%) projects contain 50 benchmarks or more. On average, a project has 19.7 ± 44.9 benchmarks, with a median of 7. The project with the largest benchmark suite is *eclipse/eclipse-collections* with 515 benchmarks.

Benchmark parameterization is quite common with projects having, on average, 70.6 ± 303.3 parameterized benchmarks, with a median of 9. 76.9% of the benchmarks have 10 parameter combinations or fewer. We find the largest number of parameter combinations in the project *msteindorfer/criterion*, with 4,132 combinations; and the project with the the most parameter combinations for a single benchmark is *apache/hive*, which contains an individual benchmark³ with an astounding 2,304 different combinations. However, the majority of the benchmarks are not parameterized, i.e., 10,394 (77.6%).

Extracting the JMH version is crucial for our analysis, as the default values of the execution configurations have been changed with version 1.21 (see also Section 2). However, automatically extracting the JMH version is not possible for each project. We are able to successfully extract the JMH version from build scripts for 573 (76%) of our pre-study subjects, containing 10,816 (80.8%) benchmarks. About 20% of the projects (containing 4,115 (38.0%) benchmarks) already use the most recent JMH version.

3.3 Results

We use this data to analyze how much time benchmark suites in the wild take to execute. Figure 2a depicts a summary of benchmark suite execution times T for the 573 studied projects where JMH version extraction was successful. The runtimes vary greatly, ranging from 143 milliseconds for *protobufel/protobuf-el* to no less than 7.4 years for *kiegroup/kie-benchmarks* (clearly, this project does not intend to execute all benchmarks at once), with a median of 26.7 minutes. 364 (49%) benchmark suites run for an hour or less, which is probably acceptable, even in CI environments. However, 110 (15%) suites take longer than 3 hours, with 22 projects (3%)



(a) Benchmark suite execution (b) Decreased t^b compared to times T

Figure 2: Impact of custom configurations on the execution times of (a) benchmark suites and (b) benchmarks.

exceeding 12 hours runtime. For example, the popular collections library *eclipse/eclipse-collections* has a total benchmark suite runtime of over 16 days, executing 515 benchmarks with 2,575 parameter combinations. We conclude that at least 15% of the pre-study subjects would greatly benefit from an approach to reduce benchmark execution times given their current configuration.

The benchmark suite execution time is based on the extracted JMH configurations from the projects. We speculate that developers specifically apply custom configurations to reduce the default settings of JMH. Indeed, 4,836 (36%) benchmarks have a configuration change that affects its runtime, of which 4,576 (34%) have a decreased benchmark time t^b with respect to JMH defaults (see Fig. 2b). We observe that for the majority of the benchmarks the execution time is in fact drastically reduced: for 3,735 (28%) and 2,379 (18%) by a factor ≥ 5 and ≥ 10 , respectively. Still 374 (3%) benchmarks are reduced by a factor ≥ 50 . While only a minority of 250 (2%) of the benchmarks belonging to just 17 (3.0%) of projects are configured to increase execution time compared to the defaults.

Pre-Study Summary. OSS developers extensively customize benchmark configurations, often setting their values considerably lower than the JMH default. Despite these changes, 15% of the projects still have a benchmark suite execution time of over 3 hours. These findings indicate that developers of many projects could be supported by a data-driven way to reduce the execution time of JMH benchmarks.

4 DYNAMIC RECONFIGURATION

In Section 3, we found that real-world OSS benchmark suites often are configured to considerably reduce runtime, with respect to JMH’s defaults; still many run for multiple hours, making it effectively impossible to assess performance on every software change. We hypothesize that this time reduction is an effort by developers to keep benchmark suite runtimes reasonable without confirming that benchmark results remain stable (accurate).

This section introduces an approach to dynamically stop benchmarks when their result is stable, with the goal of saving execution time without sacrificing quality.

²https://console.cloud.google.com/bigquery?project=fh-bigquery&page=dataset&d=github_extracts&p=fh-bigquery

³`VectorGroupByOperatorBench.testAggCount`

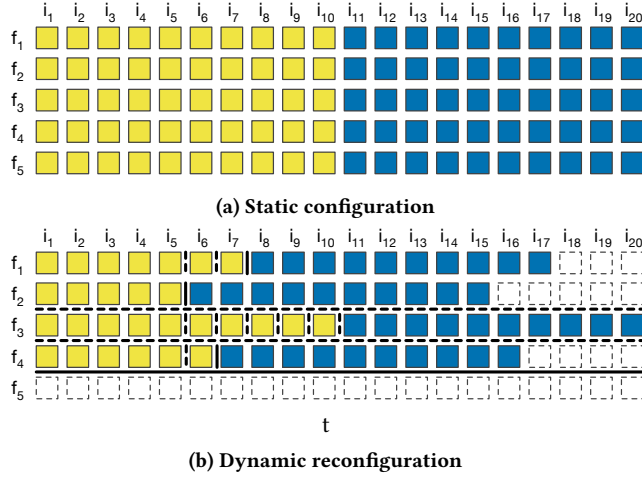


Figure 3: Standard JMH execution with static configuration vs. dynamic reconfiguration approach. A yellow box is a warmup iteration, a blue box is a measurement iteration, and a dashed box is a skipped iteration. A solid line indicates that the stoppage criterion is met, and a dashed line indicates the opposite.

4.1 Approach

JMH allows developers to define benchmark configurations before execution, either through annotations or CLI parameters, and then executes all benchmarks according to this configuration (see Section 2). We call this the “static configuration” of a benchmark execution. Figure 3a shows the static configuration where every row indicates a JMH fork (f_1 – f_5) and every column/rectangle an iteration (i_1 – i_{20}) of the corresponding fork. Yellow rectangles (i_1 – i_{10}) indicate warmup iterations, and blue rectangles (i_{11} – i_{20}) indicate measurement iterations. This static configuration bears the problem that all forks are executed with the same configuration, irrespective of the accuracy of the results, potentially wasting precious runtime.

In order to stop benchmark executions when their result is accurate enough, we propose dynamic benchmark reconfiguration, i.e., an approach that dynamically decides, at certain checkpoints, when the benchmark results are unlikely to change with more executions. This happens at two points: (1) within a fork, when the execution reaches a steady state, i.e., the warmup phase was executed long enough, and (2) after a fork’s execution, when it is unlikely that more forks will lead to different results. Figure 3b illustrates dynamic reconfiguration. Vertical bars indicate checkpoints after iterations (line 7), horizontal bars indicate checkpoints after forks (line 10), and white, dashed boxes represent iterations that are skipped.

Algorithm 1 depicts the pseudo code for our dynamic reconfiguration approach. The algorithm takes the benchmark b to execute, its extended JMH execution configuration C^b , a stability function $stable$ that is executed at every checkpoint, and a threshold t for deciding what is considered stable. C^b is a tuple of configuration values defined as $C^b = \langle wi^{min}, wi^{max}, mi, f^{min}, f^{max}, wf, wt, mt \rangle$ (see also Section 2). Note that checkpoints only happen after i_5 and

Algorithm 1: Dynamic reconfiguration algorithm

Input: $b \in B$: the benchmark to execute
 $C^b = \langle wi^{min}, wi^{max}, mi, f^{min}, f^{max}, \dots \rangle$: execution configuration for b
 $stable : M \times T \mapsto \{true, false\}$: stability function at threshold $t \in T$ for a set of measurements $M' \in M$
 t : stability threshold (specific for $stable$)
Data: $execute : B \mapsto M$: executes a benchmark iteration
Result: Measurements M^b of the benchmark b

```

1 begin
2    $M^b \leftarrow \emptyset$ 
3   for  $f \leftarrow 1$  to  $f^{max}$  do
4      $M_w \leftarrow \emptyset$ 
5     // dynamic warmup phase
6     for  $wi \leftarrow 1$  to  $wi^{max}$  do
7        $M_w \leftarrow M_w \cup execute(b, C^b)$ 
8       // warmup stoppage
9       if  $wi \geq wi^{min} \wedge stable(M_w, t)$  then break
10    // measurement phase
11    for 1 to  $mi$  do
12       $M^b \leftarrow M^b \cup execute(b, C^b)$ 
13      // fork stoppage
14      if  $f \geq f^{min} \wedge stable(M^b, t)$  then break
15  return  $M^b$ 

```

f_2 in the example, defined as wi^{min} and f^{min} . If a benchmark is not stable at a checkpoint, the bar is dashed (solid otherwise) and the warmup phase continues or another fork is spawned.

To circumvent the situation where a benchmark’s warmup phase never reaches a steady state or the overall measurements are never accurate enough, our approach takes a maximum number of warmup iterations (wi^{max}) and forks (f^{max}), e.g., f_3 has a dashed bar after the last warmup iteration. This guarantees that a single benchmark execution never exceeds a configurable time budget, which defaults to JMH’s warmup iterations (wi) and forks (f).

Benchmarks often exhibit multiple steady states resulting in multi-modal distributions, and outliers, due to non-deterministic behavior, might still occur even after $stable$ considered a fork to be in a steady state [19]. Therefore, our approach uses a fixed number of multiple measurement iterations mi (lines 8–9), as a single measurement iteration would not accurately represent a fork’s performance.

4.2 Stoppage Criteria

To decide whether a fork reached a steady state (line 7) or the gathered measurements are stable (line 10), our approach needs to decide whether more measurements provide significantly more accurate results. For this, we rely on statistical procedures on the performance measurement distributions. That is, if more measurements (i.e., data points) are unlikely to change the result distribution, we consider the measurement stable. There are three key aspects to consider: (1) a stability criteria $sc : M \mapsto \mathbb{R}^+$ that assigns a stability value $s \in \mathbb{R}$ to a set of measurements $M' \in M$; (2) a threshold $t \in T$ that indicates whether a stability value s is considered stable; and (3) a stability function $stable : M \times T \mapsto \{true, false\}$ that, based on a set of stability values (extracted from a set of measurements $M' \in M$) and a threshold $t \in T$, decides whether a set of performance measurements is stable or not.

4.2.1 Stopping Warmup Phase. The first stoppage point (line 7) decides whether a fork is in a steady state, which indicates the end of the warmup phase and the start of the measurement phase. For this, the dynamic reconfiguration approach uses a sliding-window technique where the measurement distributions at the last iterations are compared along a stability criteria. Let us consider the set of warmup measurements M_w (across multiple warmup iterations) such that $m_i \in M_w$ is the measurement at the i^{th} iteration. We then define the sliding-window warmup vector $W_{i''}$ after a current iteration i'' , a sliding-window size s^W , and the resulting start iteration of the window $i' = i'' - s^W : i' \geq 1$ in Eq. (1).

$$W_{i''} = \left\langle sc \left(\bigcup_{i=i'}^x m_i \right) \mid i' \leq x \leq i'' \right\rangle \quad (1)$$

4.2.2 Stopping Forks. The second stoppage point (line 10) decides whether the benchmark measurement results M^b are sufficiently stable and no additional fork needs to be spawned, therefore stopping the execution of benchmark b . Let us consider the set of measurements M^b (across multiple forks) such that $M_f^b \subseteq M^b$ is the subset of measurements at fork number f . We then define the fork vector $F_{f''}$ after a current fork f'' in Eq. (2).

$$F_{f''} = \left\langle sc \left(\bigcup_{f=1}^x M_f^b \right) \mid 1 \leq x \leq f'' \right\rangle \quad (2)$$

4.2.3 Stability Criteria and Function. The dynamic reconfiguration approach allows for different stability criteria (sc) and functions ($stable$), and we identified and evaluated three:

Coefficient of variation (CV): coefficient of variation (CV) is a measure of variability under the assumption that the distribution is normal. However, performance distributions are usually non-normal, e.g., multi-modal or long-tailed [11, 37]. As a stability criteria sc , CV might still be a “good enough” proxy to estimate a benchmark’s stability, especially due to its low computational overhead. Depending on the benchmark, the stability values in the vector $v \in \{W_{i''}, F_{f''}\}$ converge towards different values, making a global threshold t for all benchmarks unrealistic. Instead, we compare all stability values from v such that the delta between the largest and the smallest is at most t . Formally, $stable^{var}(M', t) = true \iff \max(v) - \min(v) \leq t$.

Relative confidence interval width (RCIW): The second stability criteria sc —relative confidence interval width (RCIW)—is similar to CV, as it estimates a benchmark’s variability, hence $stable^{var}$ also applies here. Different from CV, we employ a technique based on Monte Carlo simulation called bootstrap [13, 21] to estimate the RCIW for the mean. For this, we utilize the tool *pa* [31] that implements a technique by Kalibera and Jones [27]. It uses hierarchical random resampling [43] with replacement, which is tailored to performance evaluation. The hierarchical levels are (1) invocations, (2) iterations, and (3) forks (we refer to *pa* [31] and Kalibera and Jones [27] for details).

Kullback-Leibler divergence (KLD): The third stability criteria sc uses a technique outlined by He et al. [20] that constructs a probability that two distributions d_1 and d_2 are similar based on the Kullback-Leibler divergence (KLD) [29]. sc computes this

probability (for every element of the vector v) where d_1 is the measurement distribution *excluding* the last measurement (warmup iteration i or fork f) and d_2 is the measurement distribution *including* the last measurement. Consequently and different from a variability-based stability criteria, the vector v consists of probabilities rather than variabilities. The stability function $stable$ checks whether the mean probability of the stability values from v are above the threshold t . Formally, $stable^{prob}(M', t) = true \iff \text{mean}(v) > t$.

4.3 Modified JMH Implementation

We implemented the dynamic reconfiguration approach with the three stoppage criteria for JMH version 1.21, by adding a reconfiguration benchmark mode with stoppage criteria (sc and $stable$) and threshold (t) properties, annotation properties for wi^{min} and f^{min} , and corresponding CLI flags. Additionally, we adapted JMH’s console and JavaScript Object Notation (JSON) result file output to include the new configuration options and added a warning if the stability criterion has not been met for a benchmark. The modified fork of JMH is available on Github [34] and part of our replication package [35].

5 EMPIRICAL EVALUATION

To assess whether dynamic reconfiguration is effective and efficient, we conduct an experimental evaluation on a subset of the Java OSS projects identified in our pre-study (see Section 2). Our evaluation compares three dynamic reconfiguration approaches (one for every stoppage criterion). As a baseline for comparison, we use standard JMH with static configuration and the default values.

To support open science, we provide all evaluation data and scripts in a replication package [35].

5.1 Research Questions

First, we want to ensure that dynamic reconfiguration *does not* change the results compared to static configuration. If the results of the same benchmark executed with static configuration and with dynamic reconfiguration are equal, we conclude that dynamic reconfiguration is effective in preserving result quality. For this, we formulate RQ 1:

RQ 1 How does dynamic reconfiguration of software microbenchmarks affect their execution result?

Second, we want to evaluate if dynamic reconfiguration improves the overall runtime of a benchmark suite, compared to static configuration, including the overhead imposed by the stoppage criteria computation. For this, we formulate RQ 2:

RQ 2 How much time can be saved by dynamically reconfiguring software microbenchmarks?

As a benchmark’s result quality (accuracy) and runtime are competing objectives, the combination of the results from RQ 1 and RQ 2 validates whether dynamic reconfiguration enables “reducing execution time without sacrificing result quality”.

5.2 Study Subjects

Evaluating the dynamic reconfiguration approach on all 753 pre-study subjects (see Section 3) is infeasible as executing benchmark

Table 1: Selected study subjects. All projects are hosted on Github except the ones indicated

Name	Project	Version	# Benchs. ¹	# Param. Benchs. ¹	Exec. Time	Domain
<i>byte-buddy</i>	<i>raphw/byte-buddy</i>	c24319a	39	39	5.42h	Bytecode manipulation
<i>JCTools</i>	<i>JCTools/JCTools</i>	19cbaae	60	148	20.56h	Concurrent data structures
<i>jdk</i>	<i>jmh-jdk-microbenchmarks</i> ²	d0fab23	994	1,381	191.81h	Benchmarks of the JDK
<i>jenetics</i>	<i>jenetics/jenetics</i>	002f969	40	40	5.56h	Genetic algorithms
<i>jmh-core</i>	<i>jmh-core-benchmarks</i> ³	a07e914	110	110	15.28h	Benchmarks of JMH
<i>log4j2</i>	<i>apache/logging-log4j2</i>	ac121e2	358	510	70.83h	Logging
<i>protostuff</i>	<i>protostuff/protostuff</i>	2865bb4	16	31	4.31h	Serialization
<i>RxJava</i>	<i>ReactiveX/RxJava</i>	17a8eef	217	1,282	178.06h	Asynchronous programming
<i>SquidLib</i>	<i>SquidPony/SquidLib</i>	055f041	269	367	50.97h	Visualization
<i>zipkin</i>	<i>openzipkin/zipkin</i>	43f633d	61	61	8.47h	Distributed tracing

¹ The numbers correspond to succeeding benchmarks and excludes 38 failing parameterizations. See our replication package for a list [35]

² Repository: <http://hg.openjdk.java.net/code-tools/jmh-jdk-microbenchmarks>

³ Module directory in repository: <https://hg.openjdk.java.net/code-tools/jmh>

suites potentially takes a long time. Hence, we perform purposive sampling [5] to select a subset of ten, non-trivial projects from a wide variety of domains with small (16) to large (994) benchmark suites. Our evaluation executes all 3,969 benchmark parameter combinations of the ten study subjects, which are 8.2% of the 48,107 parameter combinations from the pre-study.

Table 1 lists the study subjects with their number of benchmarks (“# Benchs.”) and benchmark parameter combinations (“# Param. Benchs.”), *git* version used for the evaluation (“Version”), and execution time when using JMH default values (“Exec. Time”).

5.3 Study Setup

We execute all benchmarks, retrieve the benchmark results, and afterwards apply dynamic reconfiguration and the stoppage criteria to the obtained data set. This allows us to experiment with thresholds and parameters without having to rerun the full benchmark suites with our modified JMH implementation (with dynamic reconfiguration).

5.3.1 Execution and Data Gathering. As performance measurements are prone to confounding factors [11, 15, 18, 37, 40], we apply the subsequent steps to follow a rigorous methodology in order to increase result reliability.

- (1) All benchmark suites are patched with JMH 1.21.
- (2) We compile and execute all benchmarks with AdoptOpenJDK and Java HotSpot virtual machine (VM) version 1.8.0_222-b10, except *log4j2* which requires a Java Development Kit (JDK) version ≥ 9 , hence we employ version 13+33.
- (3) We run the benchmarks on a bare-metal machine [4, 46] with a 12-core Intel Xeon X5670 @2.93GHz CPU, 70 GiB memory, and a Samsung SSD 860 PRO SATA III disk, running ArchLinux with a kernel version 5.2.9-1-1-ARCH.
- (4) All non-mandatory background processes except *ssh* are disabled, without explicitly disabling software/hardware optimizations.
- (5) Regarding benchmark suite execution, we configure and execute all benchmarks with five forks *f*, 100 measurement iterations *mi*, 1s measurement time *mt*, and JMH’s `sample` mode,

set through JMH’s CLI. This configuration corresponds to the JMH 1.21 defaults, only *mt* changes from 10s to 1s but, at the same time, *mi* increases by a factor of 10, which grants our approach more checkpoints. Note that warmup iterations *wi* are set to zero but *mi* is doubled (from 50 to 100), which is required to obtain results for every iteration to dynamically decide when to stop the warmup phase. The resulting execution configuration is then $C^b = \langle 0, 0, 100, 5, 5, 0, 0s, 1s \rangle$.

- (6) We remove outliers that are a magnitude larger than the median.

5.3.2 Approach. With the obtained performance results from the suite executions, we evaluate dynamic reconfiguration with the following configuration parameters. Recall the configuration definition $C^b = \langle wi^{min}, wi^{max}, mi, f^{min}, f^{max}, wf, wt, mt \rangle$ (see Section 4.1).

Static Configuration (Baseline). The baseline, i.e., JMH with static configuration, uses the JMH 1.21 default configuration for all benchmarks. For this, we remove from the gathered data the first 50 iterations (corresponding to *wi*) from each fork and use the 50 remaining iterations as *mi*. Hence, the baseline has the following configuration: $C^b = \langle 50, 50, 50, 5, 5, 0, 1s, 1s \rangle$.

We consciously decided for the JMH default configuration as baseline and against the developers’ custom benchmark configurations for the following reasons: (1) 36% of the pre-study benchmarks change the benchmark execution time through custom configurations, hence, 64% of the benchmarks still use the JMH default configuration; (2) the majority of these benchmarks (28% of all pre-study benchmarks) only use a single fork *f*, which is considered bad practice as inter-JVM-variability is common [32], basically invalidating developers’ custom configurations for rigorous benchmarking; and (3) a unified benchmark configuration as the baseline enables comparability across our study subjects.

Dynamic Reconfiguration. For the dynamic reconfiguration approaches, we employ the configuration $C^b = \langle 5, 50, 10, 2, 5, 0, 1s, 1s \rangle$ for all benchmarks, which changes the minimum warmup iterations ($wi^{min} = 5$) and minimum forks ($f^{min} = 2$) compared to the

baseline. Note that we also reduce mi to 10 instead of 50, which the baseline uses. Initial experiments showed that an increase in measurement iterations, *after* a steady state is reached, has only a minor effect on result accuracy *but* with considerably longer runtimes.

We use the following parameters for the three dynamic reconfiguration approaches (one per stoppage criterion).

- (1) We draw a weighted sample of 1,000 invocations per iteration to reduce computational overhead at checkpoints.
- (2) The sliding-window size is set to $s^W = 5$.
- (3) **CV** uses a threshold $t = 0.01$, which corresponds to a maximum variability difference in the sliding window of 1%.
- (4) **RCIW** uses a 99% confidence level, 1,000 bootstrap iterations (which is a good tradeoff between runtime overhead and estimation accuracy), and a threshold $t = 0.03$ following best practice [18].
- (5) **KLD** partitions the distributions d_1 and d_2 into 1,000 strips for the KLD calculation [20]; removes outliers that are more than $1.5 \times IQR$ away from the median; and uses a threshold $t = 0.99$, which corresponds to a mean probability within the sliding window of 99% or larger. More strips would result in longer calculation times for the kernel density estimation and, consequently, in a higher runtime overhead. Without the outlier removal, KLD would not converge above the probability threshold t , and, hence, our approach would not stop the benchmark execution. Note that the outlier removal is *only* performed as part of the stoppage criteria calculation of our approach; for the evaluation, we consider all measurements and *do not* remove any outliers (see Section 5.4).

5.4 Results and Analysis

We now present the results of our empirical evaluation by comparing the benchmark results of the static configuration to the ones of our dynamic reconfiguration approaches with the three stoppage criteria.

5.4.1 RQ 1: Result Quality. To assess whether applying dynamic reconfiguration changes benchmark results and to answer RQ 1, we perform two analyses between the execution results coming from the baseline with static configuration and each of the three dynamic reconfiguration approaches: (1) statistical A/A tests and (2) mean performance change rate.

A/A Tests. An A/A test checks whether results from two distributions are *not* significantly different, where no difference is expected. In our context, this means if an A/A test between static configuration and dynamic reconfiguration (for each stoppage criterion) does not report a difference, we conclude that dynamic reconfiguration does not change the benchmark result. Following performance engineering best practice [9, 10, 27, 33], we estimate the confidence interval for the ratio of means with bootstrap [13], using 10,000 iterations [21], and employing hierarchical random resampling with replacement on (1) invocation, (2) iteration, and (3) fork level [27] (again relying on *pa* [31]). If the confidence interval (of the ratio) straddles 1, there is no statistically significant difference. Note that this procedure is different from the stoppage criteria RCIW (see Section 4); here we compare the results (all measurement iterations mi

Table 2: Result quality differences between static configuration approach and dynamic reconfiguration approaches

	CV	RCIW	KLD
A/A tests <i>not</i> different	78.8%	87.6%	79.6%
Mean change rate	3.1% \pm 8.1%	1.4% \pm 3.8%	2.4% \pm 7.4%
# benches < 1%	57.4%	73.2%	62.3%
# benches < 2%	72.4%	87.0%	78.2%
# benches < 3%	79.6%	91.9%	84.6%

from all forks f) of two techniques, whereas RCIW uses confidence interval widths as a variability measure of a single technique.

The first row of Table 2 shows the A/A results. For a majority of the 3,969 benchmark parameter combinations, applying dynamic reconfiguration *does not* result in significantly different distributions. About 80% or more of the benchmarks have similar result distributions compared to the static configuration. RCIW achieves the best result with 87.6%, while CV and KLD perform similarly well with 78.8% and 79.6%, respectively. Note that the static approach uses 50 measurement iterations (mi) while the dynamic approach “only” runs 10, indicating that if a steady state is reached (which is one goal of dynamic reconfiguration) more measurement iterations have a negligible impact on the overall result.

Change Rate. In addition to A/A tests, we assess the performance change rate between the static configuration approach and each of the dynamic reconfiguration approaches, i.e., by how much the means of the performance result distributions differ. The change rate augments the A/A tests’ binary decision, by showing *how* different the benchmark results become when applying dynamic reconfiguration.

The second row of Table 2 shows the mean change rate across all benchmarks in percent and its standard deviation. The mean change rate between the three stoppage criteria and the static approach is $\sim 3\%$ or lower for all three. Note that, following a rigorous measurement methodology, $\sim 3\%$ could still be caused by JVM instabilities unrelated to our approach [19]. Again, RCIW is the best criterion with $1.4\% \pm 3.8\%$. Finally, the last three rows show how many benchmarks have a change rate below 1%, 2%, and 3% for all stoppage criteria. We observe that RCIW outperforms the other two significantly, followed by KLD. $\sim 73\%$ of the benchmarks have a change rate below 1%, $\sim 87\%$ below 2%, and $\sim 92\%$ below 3%. This suggests that RCIW is a highly effective technique for stopping benchmark executions.

Figure 4 depicts the change rate distributions per project and stoppage criterion, where every data point corresponds to a benchmark’s mean performance change. Considering the median change rate of a project’s benchmarks, RCIW performs best for all projects except *jenetics*, *jmh-core*, and *SquidLib* where KLD is slightly superior. CV consistently has the largest change rates of the three stoppage criteria; nonetheless, it performs only slightly worse in most cases. Considering the mean change rate, RCIW is the most accurate stoppage criteria for 9/10 projects, with only *jmh-core* being more stable when KLD is employed. Note that for the projects where RCIW is not the best stoppage criterion, both mean and median




Figure 4: Mean change rate per study subject and stoppage criteria. The bar indicates the median, the diamond the mean, the box the IQR, and the whiskers $[Q_1|Q_3] + 1.5 * IQR$.

change rates are below 1%. The projects with the most diverging benchmarks between static configuration and dynamic reconfiguration execution are *byte-buddy*, *JCTools*, *log4j2*, and *SquidLib*. The benchmarks of these projects are less stable compared to the other projects, likely due to executing non-deterministic behaviour such as concurrency and input/output (I/O). Results from benchmarks that are less stable will potentially have statistically different distributions and, therefore, not maintain the same result quality.

Unreachable Stability Criteria. If the stability function *stable* never evaluates the measurements after a warmup iteration or a fork as stable, the maximum number of warmup iterations (wi^{max}) or forks (f^{max}) are executed. This corresponds to the static configuration of JMH. We analyzed how often stability is not achieved according to the three stoppage criteria across all study subjects. CV is the most lenient criterion with only 1.0% of the benchmarks’ forks not considered stable after 50 warmup iterations and 12% of the benchmarks insufficiently accurate after five forks. KLD achieves similar numbers (0.8%) for warmup iterations, however 46.4% of the benchmarks were not considered stable after five forks. RCIW is even more restrictive where 46.7% and 37.9% of the benchmarks do not reach the stability criteria after wi^{max} and f^{max} , respectively. This restrictiveness impacts the A/A test and mean change rate results, leading to benchmark results with higher quality. Not reaching the stability criteria can either happen if the threshold t is too restrictive or the benchmark is inherently variable, which is a common phenomenon [32, 33].

RQ 1 Summary. Applying dynamic reconfiguration *does not* change the result quality of the majority of the benchmarks, when compared to the static configuration. The RCIW stoppage criteria outperforms KLD and CV, with 87.6% of the benchmarks maintaining their result quality and a mean performance change rate of 1.4%.

5.4.2 RQ2: Time Saving. The main goal of dynamic reconfiguration is to save time executing benchmark suites. For this, and to answer RQ 2, we (1) measure the runtime overhead of the three stoppage criteria, (2) estimate the time saving for all projects compared to the

static configuration, and (3) show at which checkpoint (warmup or fork) more time can be saved.

Runtime Overhead. To measure the runtime overhead of the three stoppage criteria, we execute the benchmark suite of *log4j2* once with standard JMH 1.21 (i.e., static configuration) and once for each stoppage criteria with our JMH fork implementing dynamic reconfiguration. To ensure a valid comparison between the four measurements (static configuration + dynamic reconfiguration of three stoppage criteria), we use the same configuration for the static and the dynamic approaches of $C^b = \langle 5, 90, 10, 2, 5, 0, 1s, 1s \rangle$, but do not stop at the stoppage checkpoints. We measure the end-to-end execution time t^b of every benchmark b when executed through JMH’s CLI. This time includes JVM startup, benchmark fixtures, benchmark execution, and stoppage criteria computation, which is negligible compared to the duration of the measurement. Note that the number of data points used for the stoppage criteria calculation is independent of the study subject by construction of JMH and our approach; therefore it is sufficient to measure the overhead based on one project (see a discussion on this in Section 7).

The overheads $o \in O$ of all benchmarks for a stoppage criteria is $O = \bigcup_{b \in B} t_{dyn}^{b'} / t_{sta}^{b'} - 1$, where $t_{dyn}^{b'}$ is the execution time of the dynamic reconfiguration with a specific stoppage criteria, and $t_{sta}^{b'}$ is the execution time of the static configuration. The overheads o are independent of the number of iterations and forks executed, because they are factors of the runtime difference between dynamic reconfiguration with one stoppage criterion and the static configuration (i.e., standard JMH), and all our overhead measurements use the same configuration C^b .

The overheads we measure are $o^{CV} = 0.88\% \pm 0.34\%$ for CV, $o^{RCIW} = 10.92\% \pm 0.63\%$ for RCIW, $o^{KLD} = 4.32\% \pm 0.65\%$ for KLD. Note that changing the iteration time of 1s and executing benchmarks on different hardware might affect the overhead. The considerable difference in overhead is explained by the complexity of the stoppage criteria calculations. Whereas CV is computationally cheap (it only needs to compute standard deviation, mean, and their difference), RCIW is computationally intensive due to the simulations required for bootstrap. Because there is hardly any overhead variability ($< 1\%$) among all benchmarks, we consider the overhead constant and use the mean value for the remainder of the experiments.

Time Saving Estimation. To estimate the overall time that can be saved with dynamic reconfiguration, we adapt the execution time equation t^b (see Section 2) to incorporate the stoppage criteria. The dynamic reconfiguration benchmark execution time is then $t_{dyn}^b = \sum_{f \in forks} [(1 + o) * wi_f * wt + mi * mt]$. *forks* corresponds to the number of executed forks f of a benchmark according to the stoppage criterion, wi_f to the number of warmup iterations in this fork f , and the rest according to C^b from Section 4.1. For simplicity and because of the low variability between benchmark overheads, we disregard benchmark fixture times. The total benchmark suite execution time when using dynamic reconfiguration is then $T_{dyn} = \sum_{b \in B'} t_{dyn}^b$, where B' is the set of benchmark parameter combinations.

Table 3: Time saving per project and stoppage criteria

Project	Time Saving		
	CV	RCIW	KLD
<i>byte-buddy</i>	4.42h (81.7%)	2.62h (48.4%)	4.22h (77.8%)
<i>JCTools</i>	17.42h (84.8%)	11.45h (55.7%)	17.13h (83.3%)
<i>jdk</i>	157.32h (82.0%)	135.57h (70.7%)	154.41h (80.5%)
<i>jenetics</i>	4.78h (86.0%)	3.37h (60.7%)	4.52h (81.4%)
<i>jmh</i>	12.76h (83.5%)	12.69h (83.1%)	12.42h (81.3%)
<i>log4j2</i>	54.56h (77.0%)	39.12h (55.2%)	55.96h (79.0%)
<i>protostuff</i>	3.43h (79.6%)	2.91h (67.7%)	3.44h (79.8%)
<i>RxJava</i>	147.91h (83.1%)	121.55h (68.3%)	138.68h (77.9%)
<i>SquidLib</i>	43.07h (84.5%)	30.70h (60.2%)	41.11h (80.7%)
<i>zipkin</i>	6.17h (72.8%)	4.93h (58.2%)	6.59h (77.8%)
Total	451.84h (82.0%)	364.92h (66.2%)	438.48h (79.5%)

Table 3 shows the time saving per project and stoppage criteria in absolute numbers (hours) and relative to the static configuration. We observe that dynamic reconfiguration with all three stoppage criteria enables drastic time reductions compared to static configuration. In total, CV and KLD save ~80% and RCIW ~66% of the benchmark suite execution times of all projects combined. For individual projects, the time saving ranges between 72.8% and 86.0% for CV, 48.4% and 83.1% for RCIW, and 77.8% and 83.3% for KLD. Even with the computationally most expensive technique, i.e., RCIW, we can save at least 48.4% of time. In total numbers, the savings are between 3.43h and 157.32h for CV, 2.62h and 135.57h for RCIW, and 3.44h and 154.41h for KLD.

Stoppage Criteria Checkpoints. Dynamic reconfiguration defines two points during benchmark execution when to stop: (1) after the warmup phase if measurements are stable within a fork and (2) after a fork if measurements across forks are stable. In our analysis, the range of warmup iterations is from five (w_i^{min}) to 50 (w_i^{max}), and forks are between two (f^{min}) and five (f^{max}) (see C^b in Section 4.1). Although CV and KLD save a similar amount of time, they have different stoppage behavior. Where CV requires more warmup iterations (18.5 ± 9.4) than KLD (14.1 ± 6.9), the opposite is the case for forks with 3.1 ± 1.2 vs. 4.1 ± 1.2 , respectively. RCIW, which saves considerably less time, demands more warmup iterations (34.6 ± 16.6) to consider a fork stable but lies between CV and KLD in terms of forks (3.3 ± 1.4). The reported numbers are arithmetic means (of warmup iterations and forks) with standard deviations across all benchmarks of all study subjects. Generally, warmup iterations are more reduced than forks in our setup, indicating that fork-to-fork variability is more present than within-fork variance, that is variability across multiple JVMs compared to within a JVM, respectively. Dynamic reconfiguration enables finding the sweet spot between shortening warmup iterations and forks in combination with a certain stoppage criteria.

RQ 2 Summary. With runtime overheads between <1% and ~11%, dynamic reconfiguration enables reducing benchmark suite runtimes by 48.4% to 86.0% compared to JMH’s default runtime.

6 DISCUSSION AND RECOMMENDATIONS

Our pre-study (see Section 3) shows that developers often drastically reduce benchmark execution times. We see two potential reasons for this: (1) the benchmark suite runtimes are too long, and, consequently, developers trade shorter runtimes for inaccurate results; or (2) JMH defaults are overly conservative, and benchmarks with shorter runtimes often still produce results that are considered sufficiently accurate. We hypothesize that the former is more likely, but leave the developer perspective for configuration choices for future work. In any case, the proposed dynamic reconfiguration approach enables reducing time while maintaining similar benchmark results, as our empirical evaluation shows.

Recommendations for Developers. Developers are advised to either assess their benchmark accuracies when executed in their environment and adjust configurations accordingly, or employ dynamic reconfiguration which is able to adjust to different execution environments. The choice of stoppage criteria depends on the required result quality and, therefore, the performance change sizes desired to be detected. For slightly less accurate results but more time reduction, we recommend using KLD, otherwise RCIW is preferred. The exact threshold t depends on the stability of the execution environments the benchmarks that are run in it. If a controlled, bare-metal environment is available, we suggest the thresholds of our study. In a virtualized or cloud environment, the thresholds need to be adjusted (see also He et al. [20]). The effectiveness of our technique in non-bare-metal environments, such as in the cloud, is subject to future research. Moreover, whether a combination of different stoppage criteria, e.g., stopping when both KLD and RCIW deem a benchmark run to be stable, improves result accuracy also requires further research. Such a combination would, however, negatively affect the runtime overhead of dynamic reconfiguration.

Microbenchmarks in CI. The long benchmark execution times (see Section 3 and [24, 32, 45]) are a major obstacle for including microbenchmarks in CI [6]. To overcome this hurdle, a combination of our technique with benchmark selection [14], benchmark prioritization [39], and risk analysis on commits [24] would reduce the required time for microbenchmarking and potentially enable CI integration. Continuously assessing software performance would increase confidence that a change does not degrade performance and likely be beneficial for performance bug root cause analysis.

Choosing JMH Configuration Parameters. Choosing JMH configuration parameters that keep execution time low and result accuracy high is non-trivial, and developers decrease configurations drastically. Our results show the importance of setting the warmup phase correctly and utilizing multiple forks for benchmark accuracy. With a large number of benchmarks, expecting developers to pick the “right” values becomes unrealistic. Our dynamic reconfiguration approach helps in this regard by deciding based on data and per benchmark when the results are accurate enough.

Iteration Time and Forks. The warmup and measurement times affect benchmark result accuracy and control the frequency with which stability checkpoints occur. JMH 1.21 changed the iteration time from 1s to 10s, and reduced the number of forks from ten to

five [47, 48]. The OpenJDK team argued that 1s is too short for large workloads [48]. We performed an additional analysis whether result accuracy changes when switching from 10s to 1s but did not observe differences in most cases. Hence, we decided for 1s iterations to give the dynamic reconfiguration approach more checkpoints to assess a benchmark's stability. Whereas 10s is a safe choice for static configurations, we believe that 1s provides more flexibility and works better with dynamic reconfiguration. Our results support reducing to five forks, which indicates that most fork-to-fork variability is captured.

Choosing Stability Criteria Parameters. Choosing optimal meta-parameters for the stability criteria can affect the effectiveness and efficiency of the overall approach. Dynamic reconfiguration supports the sliding-window size s^W , the threshold t when a stability criterion value (CV, RCIW, or KLD) is considered stable, and stability-criterion-dependent parameters (see Section 5.3). We base our parameters on common statistical practice and previous research [18, 20] (see Section 5.3). Only the sliding window size s^W is manually set by us. Our empirical evaluation shows that the employed parameters work well across all study subjects. However, future research should explore the meta-parameter space through experimentation. It is important to emphasize that choosing these meta-parameters is an offline activity, which is done once and before executing the benchmarks; hence, the cost for choosing these parameters is not part of the overhead estimations in RQ 2.

Unreachable Stability Criteria. Although the stability criteria is frequently not met for warmup iterations or forks of individual benchmarks, at least when using KLD and RCIW, the overall runtime of the full benchmark suites is considerably reduced (see Section 5.4). Dynamic reconfiguration uses upper bounds for warmup iterations (wi^{max}) and forks (f^{max}); therefore, it does not exceed the runtime of standard JMH with static configuration. In case of an unreachable stability criteria, our JMH implementation warns the developer, who can then adjust this benchmark's upper bounds to obtain better results. Our approach could also automatically lift the caps if the desired result quality is not reached, which should be explored by future research.

7 THREATS TO VALIDITY

Construct Validity. Our pre-study (see Section 3) relies on information extracted from source code, i.e., configurations based on JMH annotations. We do not consider overwritten configurations through CLI arguments, which might be present in build scripts or documentation in the repositories. Reported runtimes do not consider fixture (setup and teardown) times, JVM startup, and time spent in the benchmark harness of JMH; and they assume iteration times are as configured, while in reality they are minimum times. Therefore, reported times might slightly underestimate the real execution times.

The results and implications from RQ 1 are based on the notion of benchmark result similarity. We assess this through statistical A/A tests (based on bootstrap confidence intervals for the ratio of means) and mean performance change rate, similar to previous work [10, 33]. Other tests for the similarity of benchmark results,

such as non-parametric hypothesis tests and effect sizes [12, 33], might lead to different outcomes.

We base the time savings from RQ 2 on overhead calculations from a single project and assume this overhead is constant for all stoppage points and benchmarks. There is hardly any reason to believe that overheads change between study subjects, benchmarks, and stoppage points, because the number of data points used for stoppage criteria computation are similar. This is due to how JMH and our approach work (see Sections 2 and 4), and how our experiment is designed (see Section 5): (1) the measurement time mt is fixed, irrespective of the benchmark workload; (2) the number of iterations mi and forks f is fixed; (3) benchmark fixtures, i.e., setup and teardown, are constant and of negligible duration compared to the measurement duration; and (4) the stoppage criteria calculation uses a sliding-window approach (s^W) and, therefore, the number of iterations used for the calculation is constant.

Further, we perform post-hoc analysis on a single benchmark execution data set for all stoppage criteria. That is, we execute the benchmark suites with five forks and 100 measurement iterations à 1s and then compute the stoppage points. Computing the stoppage points while executing test suites might lead to slightly different results.

Finally, we use a sliding-window approach for determining the end of the warmup phase with a window size (s^W) of five. Different window sizes might impose a larger runtime overhead and change the stoppage point outcomes.

Internal Validity. Internal validity is mostly concerned with our performance measurement methodology and the employed thresholds. We follow measurement best practice [18] and run experiments on a bare-metal machine [49] to reduce measurement bias [11, 15, 37, 40]. We did not explicitly turn off software and hardware optimizations, which might affect benchmark variability and, therefore, our results.

Regarding the thresholds, we started from previous works [18, 20] and adapted them to fit the context of microbenchmarks. As we used the same thresholds for all benchmarks and projects, we are confident that they are generally applicable for Java microbenchmarks executed on a similar machine to ours.

Further, the times reported in Section 3 rely on the JMH version of a benchmark; we applied simple heuristics to extract the version, which might not be fully accurate in case of, for instance, multi-module projects or dynamic JMH version declarations.

External Validity. Generalizability might be affected with respect to the studied projects. We only focus on OSS projects from Github, and it is unclear whether our findings are equally valid in the context of industrial software or projects hosted on other platforms. Especially, the ten selected projects for our empirical evaluation (see Section 5) might not be a representative sample for all JMH projects. Due to the long benchmark suite execution times, more projects would not have been feasible to study. We aimed for a diverse set of projects, spanning multiple domains (see Table 1), covering ~8% of the benchmarks from the pre-study (see Section 3).

The effectiveness and efficiency results of dynamic reconfiguration depends on the environment used for executing the benchmarks. Our experimental evaluation favors internal validity over external validity by using a controlled, bare-metal environment.

Therefore, we can be more confident that the reported comparisons between study subjects and stoppage criteria are indeed correct and not due to uncontrollable factors present in virtualized and cloud environments. Executing benchmarks with dynamic reconfiguration in such other environments might lead to different results.

Moreover, our focus has been on Java projects that use JMH as their benchmarking framework. Although the concepts from Section 4 also translate to other frameworks and languages, the exact results might be different. We opted for Java/JMH because (1) it is a dynamically compiled language where warmup phases and multiple forks are essential, (2) JMH benchmark suites are long running [32] and can benefit greatly from dynamic reconfiguration, and (3) JMH is a mature framework with many features offering great opportunities for our approach.

Finally, switching to different Java virtual machines, such as Eclipse OpenJ9 or Graal, might change the results due to different performance characteristics.

8 RELATED WORK

Performance testing is a form of measurement-based performance engineering [52], which comes in two main flavors: system-level tests and method/statement-level tests. Historically, research focussed on system-level tests [26, 38, 51], such as load and stress testing, with more recent advances targeting industrial applicability and practice [17, 41]. The other flavor, i.e., software microbenchmarks and performance unit tests, has only recently gained popularity in research. Studies on OSS projects [36, 49] found that adoption lags behind their functional counter-parts, i.e., unit tests. One problem is that handling performance tests is complex and requires in-depth knowledge from developers. To reduce this friction, Ding et al. [16] studied utilizing unit tests for assessing performance properties. Bulej et al. [9] proposed a framework that lets developers specify performance assertions and handles rigorous statistical evaluation. Horký et al. [23] compose performance unit test outcomes into code documentation to raise performance awareness, and Damasceno Costa et al. [12] uncover bad practices in microbenchmark code through static analyses. Generating tests removes the need to write tests by hand: AutoJMH helps avoiding pitfalls rooted in compiler optimization [44], Pradel et al. [42] generate performance regression tests for concurrent classes, and PerfSyn synthesizes inputs through mutation that expose worst-case performance behaviour [50]. Our work is orthogonal to the aforementioned works: it dynamically adapts software microbenchmark configurations to stop their execution once their result is stable.

Long execution times [17, 24, 32] and uncertain results [33, 37] are well-known to complicate the usage of performance tests in general, including software microbenchmarks. There are a few approaches that reduce the time spent in performance testing activities without considering result quality: (1) predicting commits that are likely to impact performance [24, 45], (2) prioritizing [39] and (3) selecting [3, 14] the tests in a suite that are more likely to expose performance changes. Our approach pursues the same goal of reducing benchmarking time, but with a focus on running all benchmarks (similar to prioritization) as long as necessary while maintaining the same result quality.

Result quality is impaired by not running enough measurements as well as measurement bias, which requires careful experiment planning and execution [7, 11, 15, 18, 19, 28, 40]. To mitigate measurement bias, Georges et al. [18] outlined a rigorous methodology how to assess performance of Java programs, which we base our measurement technique on. Using the correct statistical techniques to assess performance is paramount, with estimated confidence intervals using bootstrap being the state-of-the-art [8, 9, 27, 33]. One of our stopping criteria is based on and our result quality evaluation uses confidence intervals with bootstrap. To decide how many measurements are enough, approaches using statistical techniques have been proposed, employing CV [18, 37], confidence intervals [25, 37], and the Kullback-Leibler divergence (KLD) [20]. With these, performance experiments such as benchmark executions run until their results are accurate/stable enough and then abort execution, ideally reducing execution time. Our stoppage criteria use these three techniques and apply them in the context of software microbenchmarks after the warmup phase and after every fork.

Closest to our approach are the ones by Maricq et al. [37] and He et al. [20]. Maricq et al. [37] estimate the number of trials and iterations using a bootstrap technique. While they perform this estimation before executing benchmarks, we evaluate result quality during execution. He et al. [20] stop system-level performance tests executed in cloud environments, once they reach a certain stability criteria. Different from the benchmarks used in their study, microbenchmarks are much shorter, with runtimes in the order of seconds instead of multiple hours. Our work builds on top of their statistics-based approach using KLD for system benchmarks, adapts it for microbenchmarks and extends it to other stoppage criteria.

9 CONCLUSIONS

This paper introduced a dynamic reconfiguration approach for software microbenchmarks, which reduces benchmark execution time and maintains the same result quality.

In a pre-study based on real-world configurations of 13,387 microbenchmarks coming from 753 projects, we find that developers make extensive use of custom configurations to considerably reduce runtimes for 34% of the benchmarks. Still, about 15% of the projects have benchmark suite runtimes of more than 3 hours.

Our dynamic reconfiguration approach implements data-driven decisions to stop microbenchmark executions, assisting developers with the intricate task of correctly configuring microbenchmarks. With overheads between 1% and 11%, it achieves a time reduction of 48.4% to 86.0%, with between 78.8% and 87.6% of the microbenchmarks preserving their result quality.

These results show that dynamic reconfiguration is highly effective and efficient, and we envision it to enable regular performance microbenchmarking activities, such as part of CI.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the Swiss National Science Foundation (SNSF) under project number 165546 and the Swedish Research Council (VR) under grant number 2018-04127.

REFERENCES

- [1] Hammam M. AlGhamdi, Cor-Paul Bezemer, Weiyi Shang, Ahmed E. Hassan, and Parminder Flora. 2020. Towards reducing the time needed for load testing. *Journal of Software: Evolution and Process* (July 2020). <https://doi.org/10.1002/smr.2276>
- [2] Hammam M. AlGhamdi, Mark D. Syer, Weiyi Shang, and Ahmed E. Hassan. 2016. An Automated Approach for Recommending When to Stop Performance Tests. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME 2016)*. 279–289. <https://doi.org/10.1109/ICSME.2016.46>
- [3] Deema Alshoaibi, Kevin Hannigan, Hiten Gupta, and Mohamed Wiem Mkaouer. 2019. PRICE: Detection of Performance Regression Introducing Code Changes Using Static and Dynamic Metrics. In *Proceedings of the 11th International Symposium on Search Based Software Engineering (Tallinn, Estonia) (SSBSE 2019)*. Springer Nature, 75–88. https://doi.org/10.1007/978-3-030-27455-9_6
- [4] Eytan Bakshy and Eitan Frachtenberg. 2015. Design and Analysis of Benchmarking Experiments for Distributed Internet Services. In *Proceedings of the 24th International Conference on World Wide Web (Florence, Italy) (WWW 2015)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 108–118. <https://doi.org/10.1145/2736277.2741082>
- [5] Sebastian Balmes and Paul Ralph. 2020. Sampling in Software Engineering Research: A Critical Review and Guidelines. *CoRR abs/2002.07764* (2020). arXiv:2002.07764 <https://arxiv.org/abs/2002.07764>
- [6] Cor-Paul Bezemer, Simon Eismann, Vincenzo Ferme, Johannes Grohmann, Robert Heinrich, Pooyan Jamshidi, Weiyi Shang, André van Hoorn, Monica Villavicencio, Jürgen Walter, and Felix Willnecker. 2019. How is Performance Addressed in DevOps?. In *Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering (Mumbai, India) (ICPE 2019)*. ACM, New York, NY, USA, 45–50. <https://doi.org/10.1145/3297663.3309672>
- [7] Stephen M. Blackburn, Amer Diwan, Matthias Hauswirth, Peter F. Sweeney, José Nelson Amaral, Tim Brecht, Lubomír Bulej, Cliff Click, Lieven Eeckhout, Sebastian Fischmeister, and et al. 2016. The Truth, The Whole Truth, and Nothing But the Truth: A Pragmatic Guide to Assessing Empirical Evaluations. *ACM Transactions on Programming Languages and Systems* 38, 4, Article 15 (Oct. 2016), 20 pages. <https://doi.org/10.1145/2983574>
- [8] Lubomír Bulej, Vojtěch Horký, Petr Tůma, François Farquet, and Aleksandar Prokopec. 2020. Duet Benchmarking: Improving Measurement Accuracy in the Cloud. In *Proceedings of the 2020 ACM/SPEC International Conference on Performance Engineering (ICPE 2020)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3358960.3379132>
- [9] Lubomír Bulej, Tomáš Bureš, Vojtěch Horký, Jaroslav Kotrč, Lukáš Marek, Tomáš Trojáněk, and Petr Tůma. 2017. Unit testing performance with Stochastic Performance Logic. *Automated Software Engineering* 24, 1 (01 March 2017), 139–187. <https://doi.org/10.1007/s10515-015-0188-0>
- [10] Lubomír Bulej, Vojtěch Horký, and Petr Tůma. 2019. Initial Experiments with Duet Benchmarking: Performance Testing Interference in the Cloud. In *2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. 249–255. <https://doi.org/10.1109/MASCOTS.2019.00035>
- [11] Charlie Curtsinger and Emery D. Berger. 2013. STABILIZER: Statistically Sound Performance Evaluation. In *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (Houston, Texas, USA) (ASPLOS 2013)*. ACM, New York, NY, USA, 219–228. <https://doi.org/10.1145/2451116.2451141>
- [12] Diego Elias Damasceno Costa, Cor-Paul Bezemer, Philipp Leitner, and Artur Andrezejak. 2019. What's Wrong With My Benchmark Results? Studying Bad Practices in JMH Benchmarks. *IEEE Transactions on Software Engineering* (2019), 1–1. <https://doi.org/10.1109/TSE.2019.2925345>
- [13] Anthony C. Davison and D Hinkley. 1997. Bootstrap Methods and Their Application. *J. Amer. Statist. Assoc.* 94 (Jan. 1997).
- [14] Augusto Born de Oliveira, Sebastian Fischmeister, Amer Diwan, Matthias Hauswirth, and Peter F. Sweeney. 2017. Perphhecy: Performance Regression Test Selection Made Simple but Effective. In *2017 IEEE International Conference on Software Testing, Verification and Validation (ICST)*. 103–113. <https://doi.org/10.1109/ICST.2017.17>
- [15] Augusto Born de Oliveira, Jean-Christophe Petkovich, Thomas Reidemeister, and Sebastian Fischmeister. 2013. DataMill: Rigorous Performance Evaluation Made Easy. In *Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering (Prague, Czech Republic) (ICPE 2013)*. ACM, New York, NY, USA, 137–148. <https://doi.org/10.1145/2479871.2479892>
- [16] Zishuo Ding, Jinfu Chen, and Weiyi Shang. 2020. Towards the Use of the Readily Available Tests from the Release Pipeline as Performance Tests. Are We There Yet?. In *Proceedings of the 42nd International Conference on Software Engineering (Seoul, South Korea) (ICSE 2020)*. ACM, New York, NY, USA, 12.
- [17] King Chun Foo, Zhen Ming (Jack) Jiang, Bram Adams, Ahmed E. Hassan, Ying Zou, and Parminder Flora. 2015. An Industrial Case Study on the Automated Detection of Performance Regressions in Heterogeneous Environments. In *Proceedings of the 37th International Conference on Software Engineering - Volume 2 (Florence, Italy) (ICSE 2015)*. IEEE Press, Piscataway, NJ, USA, 159–168. <https://doi.org/10.1109/icse.2015.144>
- [18] Andy Georges, Dries Buytaert, and Lieven Eeckhout. 2007. Statistically Rigorous Java Performance Evaluation. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications (Montreal, Quebec, Canada) (OOPSLA 2007)*. ACM, New York, NY, USA, 57–76. <https://doi.org/10.1145/1297027.1297033>
- [19] Joseph Yossi Gil, Keren Lenz, and Yuval Shimron. 2011. A Microbenchmark Case Study and Lessons Learned. In *Proceedings of the Compilation of the Co-located Workshops on DSM'11, TMC'11, AGERE! 2011, AOOPEs'11, NEAT'11, & VMIL'11 (Portland, Oregon, USA) (SPLASH 2011 Workshops)*. ACM, New York, NY, USA, 297–308. <https://doi.org/10.1145/2095050.2095100>
- [20] Sen He, Glenna Manns, John Saunders, Wei Wang, Lori Pollock, and Mary Lou Soffa. 2019. A Statistics-based Performance Testing Methodology for Cloud Applications. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Tallinn, Estonia) (ESEC/FSE 2019)*. ACM, New York, NY, USA, 188–199. <https://doi.org/10.1145/3338906.3338912>
- [21] Tim C. Hesterberg. 2015. What Teachers Should Know About the Bootstrap: Resampling in the Undergraduate Statistics Curriculum. *The American Statistician* 69, 4 (2015), 371–386. <https://doi.org/10.1080/00031305.2015.1089789>
- [22] Michael Hilton, Timothy Tunnell, Kai Huang, Darko Marinov, and Danny Dig. 2016. Usage, Costs, and Benefits of Continuous Integration in Open-Source Projects. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering (Singapore, Singapore) (ASE 2016)*. ACM, New York, NY, USA, 426–437. <https://doi.org/10.1145/2970276.2970358>
- [23] Vojtěch Horký, Peter Libić, Lukáš Marek, Antonín Steinhäuser, and Petr Tůma. 2015. Utilizing Performance Unit Tests To Increase Performance Awareness. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering (Austin, Texas, USA) (ICPE 2015)*. ACM, New York, NY, USA, 289–300. <https://doi.org/10.1145/2668930.2688051>
- [24] Peng Huang, Xiao Ma, Dongcai Shen, and Yuanyuan Zhou. 2014. Performance Regression Testing Target Prioritization via Performance Risk Analysis. In *Proceedings of the 36th International Conference on Software Engineering (Hyderabad, India) (ICSE 2014)*. ACM, New York, NY, USA, 60–71. <https://doi.org/10.1145/2568225.2568232>
- [25] Raj Jain. 1991. *The Art of Computer Systems Performance Analysis*. Wiley.
- [26] Zhen Ming Jiang and Ahmed E. Hassan. 2015. A Survey on Load Testing of Large-Scale Software Systems. *IEEE Transactions on Software Engineering* 41, 11 (Nov. 2015), 1091–1118. <https://doi.org/10.1109/TSE.2015.2445340>
- [27] Tomas Kalibera and Richard Jones. 2012. *Quantifying Performance Changes with Effect Size Confidence Intervals*. Technical Report 4–12. University of Kent. 55 pages. <https://www.cs.kent.ac.uk/pubs/2012/3233>
- [28] Tomas Kalibera and Richard Jones. 2013. Rigorous Benchmarking in Reasonable Time. In *Proceedings of the 2013 International Symposium on Memory Management (Seattle, Washington, USA) (ISMM 2013)*. ACM, New York, NY, USA, 63–74. <https://doi.org/10.1145/2464157.2464160>
- [29] Solomon Kullback and Richard A. Leibler. 1951. On Information and Sufficiency. *Annals of Mathematical Statistics* 22, 1 (March 1951), 79–86. <https://doi.org/10.1214/aoms/1177729694>
- [30] Christoph Laaber. 2020. *bencher - JMH Benchmark Analysis and Prioritization*. <https://github.com/chrstphlbr/bencher>
- [31] Christoph Laaber. 2020. *pa - Performance (Change) Analysis using Bootstrap*. <https://github.com/chrstphlbr/pa>
- [32] Christoph Laaber and Philipp Leitner. 2018. An Evaluation of Open-Source Software Microbenchmark Suites for Continuous Performance Assessment. In *Proceedings of the 15th International Conference on Mining Software Repositories (Gothenburg, Sweden) (MSR 2018)*. ACM, New York, NY, USA, 119–130. <https://doi.org/10.1145/3196398.3196407>
- [33] Christoph Laaber, Joel Scheuner, and Philipp Leitner. 2019. Software Microbenchmarking in the Cloud. How Bad is it Really? *Empirical Software Engineering* (17 April 2019), 40. <https://doi.org/10.1007/s10664-019-09681-1>
- [34] Christoph Laaber, Stefan Würsten, Harald C. Gall, and Philipp Leitner. 2020. JMH with Dynamic Reconfiguration. <https://github.com/sealuzh/jmh>
- [35] Christoph Laaber, Stefan Würsten, Harald C. Gall, and Philipp Leitner. 2020. Replication Package "Dynamically Reconfiguring Software Microbenchmarks: Reducing Execution Time Without Sacrificing Result Quality". <https://doi.org/10.6084/m9.figshare.11944875>
- [36] Philipp Leitner and Cor-Paul Bezemer. 2017. An Exploratory Study of the State of Practice of Performance Testing in Java-Based Open Source Projects. In *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering (L'Aquila, Italy) (ICPE 2017)*. ACM, New York, NY, USA, 373–384. <https://doi.org/10.1145/3030207.3030213>
- [37] Aleksander Maricq, Dmitry Duplyakin, Ivo Jimenez, Carlos Maltzahn, Ryan Stutsman, and Robert Ricci. 2018. Taming Performance Variability. In *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation (Carlsbad, CA, USA) (OSDI 2018)*. USENIX Association, USA, 409–425. <https://www.usenix.org/conference/osdi18/presentation/maricq>

- [38] Daniel A. Menascé. 2002. Load testing of Web sites. *IEEE Internet Computing* 6, 4 (July 2002), 70–74. <https://doi.org/10.1109/MIC.2002.1020328>
- [39] Shaikh Mostafa, Xiaoyin Wang, and Tao Xie. 2017. PerfRanker: Prioritization of Performance Regression Tests for Collection-Intensive Software. In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis* (Santa Barbara, CA, USA) (*ISSTA 2017*). ACM, New York, NY, USA, 23–34. <https://doi.org/10.1145/3092703.3092725>
- [40] Todd Mytkowicz, Amer Diwan, Matthias Hauswirth, and Peter F. Sweeney. 2009. Producing Wrong Data Without Doing Anything Obviously Wrong!. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems* (Washington, DC, USA) (*ASPLOS XIV*). ACM, New York, NY, USA, 265–276. <https://doi.org/10.1145/1508244.1508275>
- [41] Thanh H. D. Nguyen, Meiyappan Nagappan, Ahmed E. Hassan, Mohamed Nasser, and Parminder Flora. 2014. An Industrial Case Study of Automatically Identifying Performance Regression-Causes. In *Proceedings of the 11th Working Conference on Mining Software Repositories* (Hyderabad, India) (*MSR 2014*). ACM, New York, NY, USA, 232–241. <https://doi.org/10.1145/2597073.2597092>
- [42] Michael Pradel, Markus Huggler, and Thomas R. Gross. 2014. Performance Regression Testing of Concurrent Classes. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis* (San Jose, CA, USA) (*ISSTA 2014*). ACM, New York, NY, USA, 13–25. <https://doi.org/10.1145/2610384.2610393>
- [43] Shiquan Ren, Hong Lai, Wenjing Tong, Mostafa Aminzadeh, Xuezhong Hou, and Shenghan Lai. 2010. Nonparametric bootstrapping for hierarchical data. *Journal of Applied Statistics* 37, 9 (2010), 1487–1498. <https://doi.org/10.1080/02664760903046102>
- [44] Marcelino Rodriguez-Cancio, Benoit Combemale, and Benoit Baudry. 2016. Automatic Microbenchmark Generation to Prevent Dead Code Elimination and Constant Folding. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering* (Singapore, Singapore) (*ASE 2016*). Association for Computing Machinery, New York, NY, USA, 132–143. <https://doi.org/10.1145/2970276.2970346>
- [45] Juan Pablo Sandoval Alcocer, Alexandre Bergel, and Marco Tulio Valente. 2016. Learning from Source Code History to Identify Performance Failures. In *Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering* (Delft, The Netherlands) (*ICPE 2016*). ACM, New York, NY, USA, 37–48. <https://doi.org/10.1145/2851553.2851571>
- [46] Marija Selakovic and Michael Pradel. 2016. Performance Issues and Optimizations in JavaScript: An Empirical Study. In *Proceedings of the 38th International Conference on Software Engineering* (Austin, Texas) (*ICSE 2016*). ACM, New York, NY, USA, 61–72. <https://doi.org/10.1145/2884781.2884829>
- [47] Aleksey Shipilev. 2018. Reconsider defaults for fork count. <https://bugs.openjdk.java.net/browse/CODETOOLS-7902170>
- [48] Aleksey Shipilev. 2018. Reconsider defaults for warmup and measurement iteration counts, durations. <https://bugs.openjdk.java.net/browse/CODETOOLS-7902165>
- [49] Petr Stefan, Vojtěch Horký, Lubomir Bulej, and Petr Tůma. 2017. Unit Testing Performance in Java Projects: Are We There Yet?. In *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering* (L'Aquila, Italy) (*ICPE 2017*). ACM, New York, NY, USA, 401–412. <https://doi.org/10.1145/3030207.3030226>
- [50] Luca Della Toffola, Michael Pradel, and Thomas R. Gross. 2018. Synthesizing Programs That Expose Performance Bottlenecks. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization* (Vienna, Austria) (*CGO 2018*). Association for Computing Machinery, New York, NY, USA, 314–326. <https://doi.org/10.1145/3168830>
- [51] Elaine J. Weyuker and Filippos I. Vokolos. 2000. Experience with Performance Testing of Software Systems: Issues, an Approach, and Case Study. *IEEE Transactions on Software Engineering* 26, 12 (Dec. 2000), 1147–1156. <https://doi.org/10.1109/32.888628>
- [52] Murray Woodside, Greg Franks, and Dorina C. Petriu. 2007. The Future of Software Performance Engineering. In *2007 Future of Software Engineering (FOSE 2007)*. IEEE Computer Society, Washington, DC, USA, 171–187. <https://doi.org/10.1109/FOSE.2007.32>