# Renaissance: Benchmarking Suite for Parallel Applications on the JVM

Aleksandar Prokopec Oracle Labs Switzerland aleksandar.prokopec@oracle.com

Gilles Duboscq Oracle Labs Switzerland gilles.m.duboscq@oracle.com

> Lubomír Bulej Charles University Czech Republic bulej@d3s.mff.cuni.cz

Doug Simon Oracle Labs Switzerland doug.simon@oracle.com Andrea Rosà Università della Svizzera italiana Switzerland andrea.rosa@usi.ch

Petr Tůma Charles University Czech Republic petr.tuma@d3s.mff.cuni.cz

Yudi Zheng Oracle Labs Switzerland yudi.zheng@oracle.com

Thomas Würthinger Oracle Labs Switzerland thomas.wuerthinger@oracle.com David Leopoldseder Johannes Kepler Universität Linz Austria david.leopoldseder@jku.at

Martin Studener Johannes Kepler Universität Linz Austria martinstudener@gmail.com

Alex Villazón Universidad Privada Boliviana Bolivia avillazon@upb.edu

Walter Binder Università della Svizzera italiana Switzerland walter.binder@usi.ch

## Abstract

Established benchmark suites for the Java Virtual Machine (JVM), such as DaCapo, ScalaBench, and SPECjvm2008, lack workloads that take advantage of the parallel programming abstractions and concurrency primitives offered by the JVM and the Java Class Library. However, such workloads are fundamental for understanding the way in which modern applications and data-processing frameworks use the JVM's concurrency features, and for validating new just-in-time (JIT) compiler optimizations that enable more efficient execution of such workloads. We present *Renaissance*, a new benchmark suite composed of modern, real-world, concurrency primitives of the JVM. We show that the use of concurrency primitives in these workloads reveals optimization opportunities that were not visible with the existing workloads.

ACM ISBN 978-1-4503-6712-7/19/06...\$15.00 https://doi.org/10.1145/3314221.3314637 We use Renaissance to compare performance of two stateof-the-art, production-quality JIT compilers (HotSpot C2 and Graal), and show that the performance differences are more significant than on existing suites such as DaCapo and SPECjvm2008. We also use Renaissance to expose four new compiler optimizations, and we analyze the behavior of several existing ones. Evaluating these optimizations using four benchmark suites shows a more prominent impact on the Renaissance workloads than on those of other suites.

**CCS Concepts** • General and reference  $\rightarrow$  Empirical studies; Evaluation; *Metrics*; • Software and its engineering  $\rightarrow$  Just-in-time compilers; Dynamic compilers; Runtime environments; *Object oriented languages; Functional languages*; • Computing methodologies  $\rightarrow$  Parallel programming languages; Concurrent programming languages;

*Keywords* benchmarks, JIT compilation, parallelism, concurrency, JVM, object-oriented programming benchmarks, functional programming benchmarks, Big Data benchmarks

#### **ACM Reference Format:**

Aleksandar Prokopec, Andrea Rosà, David Leopoldseder, Gilles Duboscq, Petr Tůma, Martin Studener, Lubomír Bulej, Yudi Zheng, Alex Villazón, Doug Simon, Thomas Würthinger, and Walter Binder. 2019. Renaissance: Benchmarking Suite for Parallel Applications on the JVM. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '19), June* 22–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 26 pages. https://doi.org/10.1145/3314221.3314637

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *PLDI '19, June 22–26, 2019, Phoenix, AZ, USA* 

<sup>© 2019</sup> Copyright held by the owner/author(s). Publication rights licensed to ACM.

## 1 Introduction

To drive innovation in managed runtimes and just-in-time (JIT) compilers, researchers and practitioners rely on benchmarks suites that capture representative patterns of application behavior. This allows developing and validating new optimizations and memory management techniques. Even though an improvement demonstrated on a benchmark is not always a sufficient condition, it is usually a necessary condition for a new optimization or technique to be considered useful. The Java Virtual Machine (JVM) is at the forefront of managed runtime research and development, and due to its popularity and widespread use, many benchmarking suites targeting the JVM have been developed. Of those most established, the DaCapo benchmark suite introduced a set of workloads to help understand memory behavior of complex Java applications [21]. The ScalaBench suite [100] focused on Scala programs, which exhibit a significantly different behavior compared to Java programs [99], despite Scala being compiled to Java bytecode and executed on the JVM. The SPECjvm2008 focused on computationally intensive workloads [1], while the more recent SPECjbb2015 benchmark simulates the IT infrastructure of an online supermarket [4].

Still, as we argue in this paper, neither of the existing benchmark suites for the JVM was specifically focused on concurrency and parallelism. Moreover, the JVM has evolved considerably over the years, and gained support for atomic operations [54], lock-free concurrent data structures [51], and non-blocking I/O [19], as well as new language and JVM features such as Java Lambdas [36], invokedynamic and method handles [97], to name a few. At the same time, the arrival of multicore processors prompted a plethora of new programming models and frameworks, such as software transactional memory [24, 41, 101], fork-join [50], dataparallel collections [33, 81, 111], and actors [5, 39, 105]. These workloads potentially present new optimization opportunities for compilers and virtual machines.

This paper proposes a new representative set of benchmarks that cover modern JVM concurrency and parallelism paradigms. While we do not argue that the proposed benchmark suite is exhaustive in that it represents all relevant workloads or optimization opportunities, we have found these benchmarks useful, since they helped us identify, implement and assess new compiler optimizations. As we show in the paper, three other established benchmark suites did not exhibit similar characteristics.

The contributions in this work are as follows:

- We propose a new benchmark suite, which consists of 21 parallelism- and concurrency-oriented benchmarks for the JVM, and relies on multiple existing state-of-the-art Java and Scala frameworks (Section 2).
- We identify a set of metrics that reflect the usage of the basic concurrency primitives on the JVM, along with a normalization technique (Section 3).

- We show that the proposed benchmarks are diverse in the sense that they span the space of concurrency metrics better than the existing suites (Section 4).
- We demonstrate that the proposed benchmarks reveal new opportunities for JIT compilers by implementing four new optimizations in the Graal JIT compiler [32] (Section 5).
- We evaluate the impact of the four new optimizations, plus three pre-existing optimizations, on Renaissance, DaCapo, ScalaBench and SPECjvm2008 suites, showing that all 7 optimizations exhibit over 5% impact on Renaissance, compared to only 2 on ScalaBench, 1 on DaCapo, and 3 on SPECjvm2008 (Section 6).
- We show that the proposed benchmark suite is as complex as DaCapo and ScalaBench, and much more complex than SPECjvm2008, by evaluating six Chidamber and Kemerer metrics, and by inspecting the compiled code size and the hot method count of all the benchmarks (Section 7),

We present related work in Section 8, and we summarize our observations in Section 9.

## 2 Motivation and Benchmark Selection

Our work has several motivating factors. First, we realized that the state-of-the-art JIT compilers, such as C2 and Graal, have very comparable performance, e.g., on the DaCapo suite, and that it has become extremely difficult to demonstrate significant performance gains on the existing work-loads when exploring new compiler optimizations. Second, we determined that many modern programming abstractions, such as Java Lambdas [36], Streams [33], or Futures [53], were not represented in the existing benchmark suites. Moreover, we found that the optimizations in existing JIT compilers rarely focused on concurrency-related primitives.

### 2.1 Selection Methodology

To identify a core set of benchmarks which would reflect the use of modern concurrency-related abstractions and programming paradigms, we manually collected a large number of candidate workloads using popular Java and Scala frameworks. In addition to manual filtering, we also developed a tool, inspired by the AutoBench toolchain [113], to scan the online corpus of open-source projects on GitHub<sup>1</sup> for candidate workloads. When selecting benchmarks for the Renaissance suite, our main goals were as follows.

**Modern Concurrency.** Benchmarks in the suite should rely on different styles of concurrent and parallel programming. These include data-parallelism, task-parallelism, streaming

<sup>&</sup>lt;sup>1</sup>We included GitHub projects with updates in the last 3 years (2015–2017) and higher number of recent commits, pull requests, and number of stars.

and pipelined parallelism, message-based concurrency, software transactional memory, lock-based and lock-free concurrent data structures and in-memory databases, as well as asynchronous programming and network communication. **Realistic Workloads.** Benchmarks should reflect real-world applications. To this end, we focused on workloads using popular Java and Scala frameworks, such as Java Stream API [33], Apache Spark [111], the Java Reactive Extensions [15, 16], the Java Fork/Join framework [50], ScalaSTM [22], and Twitter Finagle [12]. We reused existing workloads where possible, and we adopted several standalone benchmarks that target specific paradigms, such as STMBench7 [38], J. Paumard's Scrabble benchmark [69], and the CloudSuite MovieLens benchmark for Apache Spark [35]. Overall, we collected roughly 100 benchmarks since starting the effort in 2016.

Workload Diversity. Benchmarks should be sufficiently diverse so that they exercise different concurrency-related features of the JVM, such as atomic instructions, Java synchronized blocks, thread-park operations, or guarded blocks (i.e., the wait-notify pattern). At the same time, the benchmarks should be object-oriented to exercise the parts of the JVM responsible for efficient execution of code patterns commonly associated with abstraction in object-oriented programs, i.e., frequent object allocation and virtual dispatch. To this end, we analyzed candidate projects by running their testing code and recording various concurrency-related metrics (details in Section 3) to ensure adequate coverage of the metric space. This served as the basis for the diversity analysis in Section 4. Deterministic Execution. Benchmarks should run deterministically, even though it is not possible to achieve full determinism in concurrent benchmarks due to non-determinism inherent to thread scheduling. However, apart from that, the control flow of a benchmark should not include nondeterministic choices on data produced, e.g., by a random generator seeded by the current time instead of a constant. Open-Source Availability. Benchmarks should be opensource, and rely only on open-source frameworks, whenever possible. This goal is important for several reasons. First of all, it enables inspection of the workloads by the community and allows multiple parties to collaborate on improving and maintaining the suite so that it can evolve along with the target platform ecosystem. Moreover, it enables source-code level analysis of the benchmarks, and allows evaluating the actionability of profiler results. And last, but not least, opensource software is less likely to cause licensing issues.

Avoiding Known Pitfalls. The benchmarks should avoid pitfalls that were identified in other benchmark suites over time and which hinder adoption in a broader community. The lack of benchmark source code is one such pitfall. Another pitfall is the use of timeouts, which make it difficult to analyze the benchmark execution using dynamic analysis tools that require heavy-weight instrumentation, because the timeouts keep triggering due to high overhead. Yet another pitfall is benchmark correctness, avoiding benchmarks with resource leaks and, especially in the case of concurrent benchmarks, data races and deadlocks.

#### 2.2 Renaissance Benchmarks and Harness

The set of benchmarks that we ultimately decided to include in the Renaissance suite is shown in Table 1, along with short descriptions. In total, 14 out of 21 benchmarks were adapted from existing standalone benchmarks (*akka-uct*, *als, chi-square, dec-tree, fj-kmeans, log-regression, movie-lens, naive-bays, philosophers, reactors-savina, rx-scrabble, scrabble, stm-bench7* and *streams-mnemonics*), and the rest were gathered from production usages of different companies, and usages in existing applications.

All benchmarks run within a single JVM process, and only rely on the JDK as an external dependency. Some of the benchmarks use network communication, and they are encoded as multiple threads that exercise the network stack within a single process (using the loopback interface). We realize that this does not reflect realistic network conditions, but it does exercise code responsible for spanning multiple address spaces, which is common for data-parallel computational frameworks. The default execution time of each benchmark is tuned to take several seconds.

We provide a harness that allows to run the benchmarks and collect the results, and also allows to easily add new benchmarks. The harness also provides an interface for custom measurement plugins, which can latch onto benchmark execution events to perform additional operations. For example, most of the metrics in Section 3 were collected using custom plugins. In addition, the harness allows running the benchmarks with JMH [102] (a standard microbenchmarking tool for the JVM) as a frontend to avoid common measurement pitfalls associated with benchmarking.

#### 3 Characterizing Metrics

To meet the goals outlined in Section 2.1, we define a set of metrics to characterize the usage of concurrency primitives, basic primitives of object-oriented programming, and modern programming primitives introduced in JDK 7 or later. Notably, our goal is not to define an exhaustive set of metrics to cover all of these features. Instead, we aim at profiling metrics that 1) can reasonably characterize the use of these features, 2) can be collected with simple instrumentation, and 3) can be easily understood.

In this section, we present the metrics used in benchmark selection, and the details of how we collect them. Subsequent sections compare these metrics across different benchmarks and relate them with compiler optimizations.

#### 3.1 Description of the Metrics

Table 2 presents the metrics used during benchmark selection. The initial several metrics in the list reflect the usage of some basic concurrency primitives on the JVM [94]:

| Benchmark         | Description  | Focus                                      |
|-------------------|--|--|
| akka-uct          | Unbalanced Cobwebbed Tree computation using Akka [5].                      | actors, message-passing                    |
| als               | Alternating Least Squares algorithm using Spark.                           | data-parallel, compute-bound               |
| chi-square        | Computes a Chi-Square Test in parallel using Spark ML [62].                | data-parallel, machine learning            |
| db-shootout       | Parallel shootout test on Java in-memory databases.                        | query-processing, data structures          |
| dec-tree          | Classification decision tree algorithm using Spark ML [62].                | data-parallel, machine learning            |
| dotty             | Compiles a Scala codebase using the Dotty compiler for Scala.              | data-structures, synchronization           |
| finagle-chirper   | Simulates a microblogging service using Twitter Finagle [12].              | network stack, futures, atomics            |
| finagle-http      | Simulates a high server load with Twitter Finagle [12] and Netty [11].     | network stack, message-passing             |
| fj-kmeans         | K-means algorithm using the Fork/Join framework [50].                      | task-parallel, concurrent data structures  |
| future-genetic    | Genetic algorithm function optimization using Jenetics [9].                | task-parallel, contention                  |
| log-regression    | Performs the logistic regression algorithm on a large dataset.             | data-parallel, machine learning            |
| movie-lens        | Recommender for the MovieLens dataset using Spark ML [62].                 | data-parallel, compute-bound               |
| naive-bayes       | Multinomial Naive Bayes algorithm using Spark ML [62].                     | data-parallel, machine learning            |
| neo4j-analytics   | Analytical queries and transactions on the Neo4J database [10].            | query processing, transactions             |
| page-rank         | PageRank using the Apache Spark framework [111].                           | data-parallel, atomics                     |
| philosophers      | Dining philosophers using the ScalaSTM framework [22].                     | STM, atomics, guarded blocks               |
| reactors          | A set of message-passing workloads encoded in the Reactors framework [90]. | actors, message-passing, critical sections |
| rx-scrabble       | Solves the Scrabble puzzle [69] using the RxJava framework.                | streaming                                  |
| scrabble          | Solves the Scrabble puzzle [69] using Java 8 Streams.                      | data-parallel, memory-bound                |
| stm-bench7        | STMBench7 workload [38] using the ScalaSTM framework [22].                 | STM, atomics                               |
| streams-mnemonics | Computes phone mnemonics [64] using Java 8 Streams.                        | data-parallel, memory-bound                |

Table 1. Summary of benchmarks included in Renaissance.

Table 2. Metrics considered during benchmark selection.

| Name      | Description   |
|-----------|---|
| synch     | synchronized methods and blocks executed.                         |
| wait      | Invocations of Object.wait().                                     |
| notify    | <pre>Invocations of Object.notify() and Object.notifyAll().</pre> |
| atomic    | Atomic operations executed.                                       |
| park      | Park operations.  |
| сри       | Average CPU utilization (user and kernel).                        |
| cachemiss | Cache misses, including L1 cache (instruction and data), last-    |
|           | layer cache (LLC), and translation lookaside buffer (TLB; in-     |
|           | struction and data).  |
| object    | Objects allocated.  |
| array     | Arrays allocated.   |
| method    | Methods invoked with invokevirtual, invokeinterface or            |
|           | invokedynamic bytecodes.  |
| idynamic  | invokedynamic bytecodes executed.                                 |

(1) synchronized blocks and methods (i.e. critical sections), (2) wait/notify calls (i.e. guarded blocks), (3) atomic memoryaccess operations (e.g., compare-and-swap), (4) thread parking and unparking. Since most high-level concurrency abstractions are implemented from these basic primitives, their usage rates estimate the use of concurrency abstractions during execution. We therefore collect the dynamic invocation counts of these primitives<sup>2</sup>. Next, we define a metric for the average CPU utilization, and a metric for the number of cache misses during execution. Both these metrics are indirectly related to concurrency – for example, unless several CPUs or cores are utilized by the benchmark, it is unlikely that the benchmark will exhibit interesting synchronization behavior. Similarly, a high cachemiss rate may indicate contention between threads.

Then, we track the object-allocation rate, the array-allocation rate, and the dynamic-dispatch rate (i.e., the execution counts of the invokedynamic, invokeinterface, or invokevirtual bytecodes). These metrics estimate the usage of the basic primitives in object-oriented programming, and usually correlate with the complexity of the code<sup>3</sup>.

Finally, we track the execution rate of the invokedynamic bytecode (introduced in JDK 7 [97]), which supports dynamic languages on the JVM [66], and is also used to implement Lambdas [36]. Java Lambdas are used in various data-processing libraries – for example, Java Streams [33] or Reactive Extensions [15] expose operations such as map and reduce that typically take Lambda values as arguments. Therefore, counting the invokedynamic occurrences allows estimating the usage rate of high-level operations.

<sup>&</sup>lt;sup>2</sup>We omitted unparks, since they correlated with parks in all benchmarks.

<sup>&</sup>lt;sup>3</sup>Some data-parallel and streaming frameworks allocate intermediate arrays.

#### 3.2 Normalization

Absolute values of the metrics from Table 2 are not very indicative, since a high value may reflect a long execution rather than a frequent usage of a given primitive, which is our focus. To this end, we use rates instead of absolute values, and we normalize each metric (except cpu) by the number of *references cycles* executed by the benchmark, defined as machine cycles measured at a constant nominal frequency, even if the real CPU frequency scales up or down [46].

We use reference cycles as a measure of the amount of computations executed by an application. Using reference cycles allows comparing the metrics between different benchmarks, even if they are executed under different or fluctuating CPU frequencies. Moreover, reference cycles represent runtime conditions more accurately, since they account for instruction complexity (as complex instructions take more cycles) and for latencies due to e.g. cache misses.

#### 3.3 Metric Collection

To collect the metrics, we rely on several tools. To profile the usage of concurrency primitives (synch, wait, notify, atomic, park), allocations (object, array), invocations (method), and invokedynamic executions (idynamic), we developed a profiler based on bytecode instrumentation, built on top of the open-source dynamic program-analysis framework DiSL<sup>4</sup> [61]. The profiler exploits the full bytecode coverage offered by DiSL (i.e., it can instrument every method that has a bytecode representation), so the metrics can be collected in every loaded class, including the JDK classes, as well as the classes that are dynamically generated or fetched from a remote location. The profiler uses a deployment setting of DiSL called Shadow VM [60] to enforce isolation between analysis and application code, preventing certain well-known problems that may be introduced by less isolated approaches [49]. We profile atomic and park by intercepting the respective method calls in sun.misc.Unsafe. CPU utilization is sampled every ~150ms using top [57]. Finally, we use perf [70] to collect cache misses and reference cycles (used for normalization). While instrumentation may influence thread interleaving and bias the collection, we made sure to keep the perturbations low with minimal instrumentation that avoids heap allocations, and by keeping the profiling data structures in a separate process, on a separate NUMA node. Additional details of the experimental setup are in Section B of the supplemental material.

## 4 Diversity

Having collected the metrics described in Section 3, we used these metrics to narrow down the list of benchmarks, and to compare Renaissance to the established benchmark suites. In this section, we show that Renaissance represents the selected concurrency primitives better than the existing suites,

<sup>4</sup>https://disl.ow2.org/

that it is comparable to suites such as DaCapo and ScalaBench in terms of object-allocation rates and dynamic dispatch, and that it exercises invokedynamic more often.

#### 4.1 Benchmark Suites

We compare Renaissance with benchmarks included in several established and widely-used benchmark suites for the JVM: DaCapo [21], ScalaBench [100], and SPECjvm2008 [1]<sup>5</sup>. These suites can be executed with different input sizes. In DaCapo and ScalaBench, we use the largest input defined for each benchmark, while in SPECjvm2008 we use the "lagom" workload (a fixed-size workload, intended for research purposes [2]). Table 6 in Section A of the supplemental material lists all benchmarks that we considered.

All benchmarks have a *warm-up* phase, which takes either a number of iterations (Renaissance, DaCapo, ScalaBench) or some execution time (SPECjvm2008). Execution after the warmup is classified as *steady-state*, and we always measure steady-state execution in this paper.

#### 4.2 Methodology

To study the differences between Renaissance and other suites, we first collect the metrics defined in Section 3 for all workloads;<sup>6</sup> then, we visually compare them by means of scatter plots. Rather than analyzing benchmarks on a coordinate system composed of 11 dimensions (one dimension for each collected metric) we resort to principal component analysis (PCA) [48] to reduce data dimensionality, comparing benchmarks on the resulting coordinate system. PCA is a statistical technique that transforms a K-dimensional space into another space of linearly uncorrelated variables, called principal components (PCs). PCA applies linear transformations such that the first PC has the largest possible variance, while each subsequent component has the highest possible variance, under the constraint of being orthogonal to the preceeding components. Considering variance as a measure of information, PCA allows one to reduce the dimensionality of the original dataset by considering only the first components (i.e., those retaining most of the variance) while preserving as much information as possible from the original dataset.

We apply PCA as follows. Matrix X contains the set of collected metrics. Each component  $x_{ij} \in X$  ( $i \in [1, N]$ ;  $j \in [1, K]$ ; where N is the number of benchmarks analyzed and K the number of collected metrics) is the normalized value of metric *j* obtained on benchmark *i*. The metrics are profiled during a single steady-state benchmark execution, as discussed in Section **B** of the supplemental material.

<sup>&</sup>lt;sup>5</sup>We use the latest released versions of the suites at the time of writing: DaCapo v.9.12-MR1 (dated Jan. 2018), ScalaBench v.0.1.0 (dated Feb. 2012) and SPECjvm2008 v.1.01 (dated May 2009). According to the guidelines on the DaCapo website, we use the *lusearch-fix* benchmark instead of *lusearch*. <sup>6</sup>Table 7 in Section D of the supplemental material reports the collected unnormalized metrics for all analyzed benchmarks.

**Table 3.** Loadings of metrics (defined in Table 2) on the first four principal components (PCs), sorted by absolute value (descending order).

| PC1       |       | PC2       |       | PC3       |       | PC4       |       |
|-----------|-------|-----------|-------|-----------|-------|-----------|-------|
| Metric    | load. | Metric    | load. | Metric    | load. | Metric    | load. |
| object    | +0.50 | atomic    | +0.67 | cachemiss | +0.58 | idynamic  | +0.56 |
| сри       | -0.49 | park      | +0.65 | notify    | +0.50 | array     | +0.42 |
| method    | +0.44 | method    | +0.20 | wait      | +0.41 | notify    | +0.42 |
| array     | +0.40 | notify    | +0.18 | cpu       | -0.28 | method    | -0.35 |
| idynamic  | +0.27 | idynamic  | -0.17 | synch     | -0.25 | cachemiss | +0.28 |
| synch     | -0.17 | cpu       | -0.16 | park      | -0.20 | cpu       | +0.22 |
| notify    | -0.13 | cachemiss | -0.08 | idynamic  | -0.18 | atomic    | +0.18 |
| atomic    | -0.13 | object    | +0.05 | array     | -0.13 | wait      | -0.15 |
| cachemiss | -0.07 | array     | -0.03 | method    | +0.10 | object    | +0.13 |
| park      | -0.06 | synch     | -0.02 | object    | -0.04 | synch     | +0.11 |
| wait      | -0 02 | wait      | -0 00 | atomic    | -0 03 | nark      | +0 05 |



**Figure 1.** Scatter plots of benchmark scores over the first four principal components (PCs).

Each metric vector  $X_j$  is standardized to zero mean and unit variance, yielding a new vector  $Y_j = \frac{X_j - \overline{x_j}}{s_j} \in Y$  (where  $\overline{x_j}$  and  $s_j$  are the mean and variance of  $X_j$ , respectively). According to established practices, we apply PCA on the standardized matrix Y rather than on X. PCA produces a new matrix S = YL, where each row vector in S is the projection of the corresponding row vector in Y on the new coordinate system formed by the PCs (called *score*), while  $l_{ij} \in L$  ( $l_{ij} \in$  $[-1, 1]; i, j \in [1, K]$ ) is the *loading* of metric *i* on the *j*-th PC (henceforth called PC*j*). The absolute value of a loading quantifies the correlation of a metric to a given PC, while its sign determines if the correlation is positive or negative.

#### 4.3 Analysis

Table 3 reports the loading of each metric on the first four PCs, ranked by their absolute values in descending order, while in Figure 1, we plot the scores of the considered benchmarks against the four PCs. <sup>7</sup> The first four components account for ~60% of the variance present in the data. <sup>8</sup>

As shown in Table 3, object allocation, dynamic dispatch, and array allocation correlate positively with PC1, and CPU utilization correlates negatively. Figure 1(a) shows that the SPECjvm2008 benchmarks are clustered in the bottom-left portion of the figure, while the benchmarks of the other suites are well distributed along PC1. This is a sign that Renaissance, DaCapo and ScalaBench contain applications that nicely exercise object allocation and dynamic dispatch, which is often a sign of using abstractions in those workloads.

In contrast, PC2 exhibits a strong positive correlation with two concurrency primitives, i.e., atomic and park. As demonstrated by the wide distribution of Renaissance benchmarks along PC2, as shown in Figure 1(a), Renaissance benchmarks extensively use these primitives. On the other hand, benchmarks from the other suites span a limited space along PC2. Similarly to PC2, PC3 is well correlated with metrics estimating concurrency, particularly cache misses and the usage of wait/notify. The benchmark distribution along PC3 is similar to the one along PC2 (as shown in Figure 1(b)), which suggests that generally, the use of concurrency primitives in DaCapo, ScalaBench, and SPECjvm2008 is limited.

Finally, idynamic contributes most to PC4, with a strong positive correlation. The presence of numerous Renaissance benchmarks in the top part of Figure 1(b) demonstrates the frequent execution of invokedynamic bytecodes in the workloads (indeed, 10 out of 21 benchmarks execute this bytecode at runtime), which in turn may represent the use of functional-style operations commonly found in parallelism frameworks. This is expected, as other suites were released before the introduction of invokedynamic in JDK 7 [97] and of Lambas and Streams in JDK 8 [36] (while a recent DaCapo version was released in Jan. 2018, there was no major change to the constituent benchmarks, originally released in 2009).

## 5 Optimization Opportunities

Once we identified a sufficiently diverse set of benchmarks, as explained previously in Sections 2, 3 and 4, we wanted to check if these benchmarks help us in identifying new optimization opportunities. We therefore investigated if we can develop new JIT-compiler optimizations that are useful for the code patterns that appear in Renaissance benchmarks.

In this section, we analyze several compiler optimizations for which we noticed a significant impact on Renaissance benchmarks, and we explain the reason for the impact. For the purposes of the analysis, we used the Graal JIT compiler to implement 4 new optimizations<sup>9</sup> (escape analysis with atomic operations, loop-wide lock coarsening, atomicoperation coalescing, and method-handle simplification), and to study 3 existing optimizations (speculative guard movement, loop vectorization, and dominance-based duplication simulation). Notably, we do not claim that these are the

 $<sup>^7\</sup>mathrm{A}$  larger version of Figure 1 can be found in Section E of the supplemental material.

<sup>&</sup>lt;sup>8</sup>As discussed in Section B of the supplemental material, we excluded benchmarks *tradebeans*, *actors* and *scimark.monte\_carlo* from the PCA.

<sup>&</sup>lt;sup>9</sup>Graal's compliance test suite is available at GitHub [6], and we ran it to check that our implementations are correct.

only optimizations that are important for the Renaissance benchmarks – we only discuss those parallelism-related optimizations that we were able to identify. Also, we focused on high-level optimizations for the existing backends of the Graal compiler. In particular, we did not study fence optimizations although they are interesting for platforms such as PowerPC or AArch64<sup>10</sup>.

In the rest of the section, we briefly explain each optimization and discuss its impact on the benchmarks. We also establish a relationship to the metrics from Section 3 and illustrate performance with data cited from Section 6. We present a minimal example for each optimization – while such examples are rare in the source code, they do appear in the compiler after transformations such as inlining.

In sections where optimization soundness is not apparent, we reason as follows. Each program P has some set R of possible results, where a result is a sequence of external sideeffects that the program's threads make. An external sideeffect is an output of the program that the user can observe (for example, the exit status or the I/O). Program's executions can produce different results because the program-threads' execution schedule is non-deterministic. When writing a concurrent program, users consider the program correct if its executions produce some result from the set of allowed results R. However, users are not allowed to assume any probability distribution of the program's results across executions. Thus, for an optimization to be valid, it must transform an original program P with a set of possible results R into a program *P'* with a set of possible results *R'*, such that  $R' \subseteq R$ . For each optimization, we informally reason why the set of results of the transformed program is either equal to that of the original program or its strict subset.

#### 5.1 Escape Analysis with Atomic Operations

Partial Escape Analysis (PEA) [106] is an optimization that reduces the amount of object allocations by postponing heap allocation of an object until it escapes and can be seen by other threads. PEA does this by analyzing the reads and writes to the allocated objects. So far, the PEA in Graal was limited to regular reads and writes, but it did not consider atomic operations such as Compare-And-Swap (CAS).

Consider the example on the right: if the CASes are not accounted for during PEA, then the A and B objects must be allocated before the first CAS.

1 o = new A(v); 2 CAS(o.x, v, new B(v2)); 3 CAS(o.x.y, v2, v3) 4 return o.x;

A CAS operation involves up to three objects: the object containing the field, the expected value and the new value. If the object containing the field has not yet escaped, the CAS can simply update the state of the scalar-replaced object. This might replace the two other objects with scalars. This



**Figure 2.** Number of atomic operations executed, normalized by reference cycles.

also helps when the objects escape later, because the CAS operation can be removed and the objects can be directly initialized in their mutated state before being published. In particular, this initialization can be performed with—potentially cheaper—regular writes rather than with atomic instructions.

For example, in the previous code, assuming the constructors return new B(v3); don't let the objects escape, the allocation of the A object can be completely eliminated and the value of the field in the B object can directly be set to its updated value (v3).

We observed this pattern in usages of java.util.Random, com.twitter.util.Promise, and java.util.concurrent. atomic.AtomicReference. These classes change their internal state using a CAS operation. An application can create an instance of these classes and change its internal state a few times before discarding it (full escape analysis), or publishing it to the rest of the program (partial escape analysis).

This optimization has a 24% impact on *finagle-chirper*. As seen in Figure 2, this benchmark exhibits a higher value for the atomic metric than any of the benchmarks from the existing suites. This is coherent, since the optimization targets an operation counted by the atomic metric.

**Soundness.** Objects that have not yet escaped cannot be observed or mutated by other threads by definition. Consequently, any schedule of the original program P will yield the same outcome for the CAS operation: the transformed program P' emulates this behavior on the thread which executed the CAS in *P*. If the object containing the memory location subject to CAS can be entirely escape-analyzed away in P', no other thread can observe the effects of this CAS in P. In this case, P' does not contain any external side-effect for this CAS and any result of P' is a possible result of P. If the object escapes at a later point in the program, its fully initialized state containing the emulated effects of the CAS is safely published during the side-effect that causes the object to escape. In this case any schedule of P' can be mapped to a schedule of P where all side-effects due to the initialization of this object happened before the escape point.

<sup>&</sup>lt;sup>10</sup>Graal has no PowerPC backend that we know of and its AArch64 backend has not yet been optimized for fine-grained use of fences.



Figure 3. Number of synchronized methods or blocks executed, normalized by reference cycles.

#### 5.2 Loop-Wide Lock Coarsening

Consider the Java loop shown in 1 for (; condition ;) the snippet on the right enclosing synchronized(x) { 2 a synchronized statement, which 3 ... region ... 4 } protects a critical region. The time

spent entering and exiting the critical region can critically affect performance if the loop is frequently executed.

| -                              | ,                                   |
|--------------------------------|-------------------------------------|
| 1 <b>int</b> b = 0;            | This synchronization pattern        |
| 2 for (; condition ;)          | occurs when accessing syn-          |
| <pre>3 synchronized(x) {</pre> | chronized collections such as       |
| 4 <b>for</b> (; condition;) {  | the IDV is we set il Vester and     |
| 5 region                       | the JDK Java. util. vector and      |
| 6 b + + ;                      | SynchronizedList (whose op-         |
| 7 if (b % T == 0)              | erations use locks [13, 14]) inside |
| <pre>8 break; } }</pre>        | a loop                              |
|                                | a 1000                              |

To reduce the cost of locking and unlocking within the synchronized block, we extended Graal with an optimization that coarsens the lock [30] by tiling the iteration space into C-sized chunks, as seen in the second snippet. The lock is held in the entire nested loop, so the total number of monitor operations decreases. The necessary condition is that the code motion on condition and region is legal (e.g., condition must not obtain another lock).

Existing lock optimizations in HotSpot's C2 [68] compiler can optimize loops that have a statically known number of iterations, by completely unrolling and coarsening the lock across the unrolled iterations. By contrast, our optimization tiles the iteration space of any loop to coarsen the lock, without pulling the lock out of the loop completely. Lock coarsening has an effect on fairness, but since synchronized regions in Java are inherently not fair, this optimization does not change the program semantics - applications relying on fairness must anyway use more elaborate locking mechanisms, such as the ReentrantLock class from the JDK.

Loop-wide lock coarsening improves *fj-kmeans* by about 71%. We found that a chunk size of C = 32 works well for this benchmark. Generally, the impact depends on the size of the critical region, as well as the number of loop iterations. Figure 3 shows the synchronization counts across the suites, normalized by the CPU reference cycles. Note that *fj-kmeans* 

uses the synchronized primitive considerably more often, which made it possible to identify this optimization.

Soundness. Any schedule s' of the transformed program P' can be mapped to a schedule s of the original program *P*, such that the respective external side-effects r' and r are equal. Consider a schedule s' of P' in which some thread T is about to execute a loop that contains a synchronized region, and which has external side-effects  $r_0$  and the memory state  $m_0$  before executing that loop. There exists some execution schedule s of P in which all other threads pause during the execution of the C iterations of the loop. That execution schedule has the same external side-effects  $r_0$  and the memory state  $m_0$  before executing that loop. Since condition does not acquire any locks, it is easy to show that both s and s' have the same set of external side-effects and the same memory state after executing C iterations of the loop. By induction, the two execution schedules have the same external side-effects r' = r, so the transformation is correct.

#### 5.3 Atomic-Operation Coalescing

A typical retry-loop in lock-free algorithms reads the value at a memory address, and uses a CAS instruction to atomically replace the old value with the newly computed value. If the CAS fails due to a concurrent write by another

```
1 do {
2 v = READ(x)
3
   nv = f1(v)
4 } while(!CAS(x,v,nv))
5 do {
6 \quad v = READ(x)
7 nv = f_2(v)
8 } while(!CAS(x,v,nv))
```

thread, the loop repeats, or otherwise terminates. The code snippet on the right shows two such consecutive retry-loops. This pattern rarely appears directly in user programs, but it does occur after inlining high-level operations, such as random-number generators or concurrent counters.

Consider an execution schedule in which no other thread makes progress between the READ in line 2 up to the CAS in line 8, In this execution schedule, the intermediate CAS in line 4 is not observed by any other thread. According to the Java Memory Model [94], threads are not guaranteed to observe other execution schedules of this snippet, so user programs should not assume that they will ever see f1(v).

| 1 | do {                             |
|---|----------------------------------|
| 2 | v = READ(x)                      |
| 3 | nv = f2(f1(v))                   |
| 4 | <pre>} while(!CAS(x.v.nv))</pre> |

Consequently, if the functions f1 and f2 are referentially transparent, the new value nv of the first CAS can be directly replaced

with the new value nv of the second CAS, and the second CAS can be removed; the resulting snippet is on the left.

We observed this pattern in java.util.Random usages, where the method nextDouble consecutively executes a CAS twice. This optimization mostly affects future-genetic, improving it by  $\approx 24\%$ . The high atomic operations rate in *future*genetic, shown previously in Figure 2, is due to the shared use of a pseudo-random generator, and this optimization reduces the overall contention by eliminating some CASes. **Soundness.** Call the original program *P*, and the transformed program P'. Let S be the subset of all the execution schedules

of *P* in which some thread *T* executes the CAS instructions in lines 4-8, during which all other threads are paused. We will show that any schedule of *P'* can be mapped to  $s \in S$ .

First, consider some schedule s' of P' in which the CAS in line 4 succeeds. There must exist an execution schedule  $s \in S$ of P that is exactly the same as s' until the lines 4-8. Before T executes the CAS in line 4 of P', the set of external sideeffects must have been  $r_0$  in both s and s', and the state of the memory must have been  $m_0 \cup \{x \rightarrow v\}$ . After T executes the CAS in line 4 of P', external side-effects remain  $r_0$  because f1 and f2 are referentially transparent, and memory state becomes  $m_0 \cup \{x \rightarrow f2(f1(v))\}$ . Since the CAS in line 4 of P' succeeds in s', then CAS in line 4 of P succeeds in s. Since  $s \in S$ , the CAS in line 8 must also succeed. Consequently, after T executes the lines 4-8 of P, the set of external side-effects is  $r_0$ , and  $m_0 \cup \{x \rightarrow f2(f1(v))\}$  is the state of the memory.

Therefore, in the schedule  $s \in S$ , after the other threads resume, *P* has the same memory state as *P'* in *s'*, so the remaining external side-effects  $r_1$  must be the same in both *s* and *s'*, and the execution has the same result  $r' = r = r_0 \cdot r_1$ . The proof is similar when the CAS in line 4 of *P'* fails.

In conclusion, any schedule s' of P' can be mapped to a schedule s in P that has the same result r = r'. By the previous definition,  $R' \subseteq R$  and the transformation is correct.

#### 5.4 Method-Handle Simplification

Java 8 Lambdas are used to express declarative operations in multiple frameworks such as Java Streams [33], futures [12, 51, 53, 54], and Rx [15]. In the bytecode, the lambda-value creation is encoded with an invokedynamic instruction [97]. The first time this instruction is executed, a special *bootstrap method* generates an anonymous class for the lambda, and returns a *method handle*. The method handle's invoke method is then used to call the lambda.

... invokedynamic #14, 0 aload 11 invokevirtual #15, 2 ...

Consider the sequence of bytecodes shown on the left. By the time this code gets compiled, the invokedynamic gets replaced with

the address of the generated method-handle object. The figure on the right shows the cor-

responding program segment in the intermediate representation used by Graal [32]. The Invoke node represents the second

C(0x2FB7B44C) Arguments Invoke

invoke, while Arguments encapsulates the argument list. Importantly, the first argument C is a constant that represents the address of the method-handle in memory.

In frameworks that parameterize operations with lambda values, inlining the body of the lambda typically triggers other optimizations. Unfortunately, the method handle's invoke method is polymorphic, which prevents the inlining. Even though the method-handle object has a constant address, the compiler does not know *the actual code of the JVM method that the method-handle object refers to*. We therefore used the existing JVM compiler interface [96] to resolve the

PLDI '19, June 22-26, 2019, Phoenix, AZ, USA



**Figure 4.** Number of invokedynamic bytecodes executed, normalized by reference cycles.

method-handle-object address to the JVM method that is encapsulated within the method handle. This allowed us to replace method-handle calls with direct calls, which can be inlined in the subsequent phases in the compiler.

We found that this generally improves performance across multiple methods. For example, using Oracle Developer Studio 12.6, we recorded the hottest methods in *scrabble*, along with the execution times with and without this optimization.

| with<br>(ms) | w/o<br>(ms) | Compilation unit   |
|--------------|-------------|--|
| 303.4        | 350.0       | <total></total>  |
| 25.9         | 29.6        | JavaScrabble.lambda\$run\$2                                |
| 23.7         | 23.1        | java.util.Spliterator\$OfInt.forEachRemaining              |
| 22.4         | 28.8        | <pre>java.util.stream.ReferencePipeline\$3\$1.accept</pre> |
| 18.1         | 20.4        | JavaScrabble.lambda\$run\$5                                |

In the preceding table, method-handle simplification reduces the execution time for most methods, but not for the forEachRemainining method of the Spliterator interface. The reason is as follows: parallel Stream operations split collections into subsets, pass some subsets to worker threads, and some to the calling thread. The invokedynamic instruction that creates the method-handle is executed on the calling thread. Thus, the method-handle type can be statically determined on the calling thread, but not on the worker threads, which invoke the forEachRemaining method. This indicates additional opportunities for optimizations in the future.

Method lambda\$run\$2 corresponds to the Java code on the left: a lambda that produces a histogram of characters for a

given string. Using Graal, we inspected the intermediate representation of this method, both with and without the optimization. The additional inlining induced by methodhandle simplification triggers other optimizations that reduce the number of callsites from 19 to 1, remove 37 out of 61 dynamic type- and null-checks, and reduce the total number of IR nodes in that method from 696 to 490.

Overall, method-handle simplification impacts the running time of *finagle-chirper* by 4%, *scrabble* by 22%, *streamsmnemonics* by 7%, and *rx-scrabble* by 1%. These results correlate with the invokedynamic counts normalized by CPU reference cycles, shown in Figure 4. On DaCapo, ScalaBench, and SPECjvm2008 (which precede invokedynamic), the calls to invokedynamic usually occur only indirectly through the Java class library.

#### 5.5 Speculative Guard Motion

JIT compilers for the JVM often use speculative optimizations. Some of these optimizations require the insertion of *guards*, which ensure the speculation is correct at runtime. If the specu-

lation is not correct, these guards divert the execution to an environment that does not use speculations (i.e., an interpreter, or code compiled by a baseline compiler). New profiling then takes place and a re-compilation that does not use the respective speculation is triggered. In order to avoid useless re-compilations, guards are not typically hoisted out of the branch where they are needed.

However, for some loops, it is beneficial to speculatively hoist guards out of loops even if the control-flow in the loop does not always lead to that guard.

This is beneficial because the hoisted guards are executed less often. To capture typical cases such as bounds checks, comparisons of induction variables can be rewritten to loop-invariant versions. Speculative Guard Motion is described in more details by Duboscq [31]. This optimization has a 15% impact on the *log-regression* benchmark and an 8% impact on the *dec-tree* benchmark. To better understand this difference, we measured the number of guards executed with and without speculative guard motion for the *log-regression* benchmark.

| Execution         | ns   | Guard type                       |  |  |  |
|-------------------|------|----------------------------------|--|--|--|
| Without           |      |                                  |  |  |  |
| 57 487 131        | 1%   | Others                           |  |  |  |
| 529 958 424       | 10%  | UnreachedCode                    |  |  |  |
| 1 639 903 998     | 32%  | NullCheckException               |  |  |  |
| 2 917 610 059 57% |      | BoundsCheckException             |  |  |  |
| 5144959612 100%   |      | Total                            |  |  |  |
| With              |      |                                  |  |  |  |
| 55 660 483        | 7%   | Others                           |  |  |  |
| 18 314 637        | 2%   | Speculative NullCheckException   |  |  |  |
| 22 056 283        | 3%   | Speculative UnreachedCode        |  |  |  |
| 28 570 247        | 3%   | Speculative BoundsCheckException |  |  |  |
| 62 046 637        | 7%   | BoundsCheckException             |  |  |  |
| 166 341 946       | 20%  | NullCheckException               |  |  |  |
| 499 933 160       | 59%  | UnreachedCode                    |  |  |  |
| 852 923 393       | 100% | Total                            |  |  |  |

First, note that the total number of executed guards is reduced by 83%. Moreover, note the entry for *speculative* guard variants, which accounts for the hoisted guards. In particular, we can see the large effect on BoundsCheckException and NullCheckException guards where most occurrences have been replaced by the lower-frequency speculative versions.

In some cases, removing the guards enables additional optimizations. Loop vectorization, discussed next, is one such example – we found that by disabling speculative guard motion, loop vectorization almost never triggers.

**Soundness.** Hoisting guards of loop-invariant conditions is sound since executing unnecessary guards is always sound: guards have no external side-effects, they only run the risk of reducing performance. In this case performance is conserved by not doing this transformation again if a deoptimization already happened for this loop. The rewrite of loop-variant inequalities ( $\langle, \rangle, \leq, \rangle$ , signed and unsigned) involving loop induction variables requires the induction variable to change monotonically which can be statically or dynamically checked above the loop. Once this is established, if the inequality holds at both bounds, it holds over the whole range. As a result, the hoisted guard implies the original one, ensuring the transformed program will deoptimize in at least as many cases as the original.

#### 5.6 Loop Vectorization

Modern CPUs include vector instruction sets that ensure better performance by operating on multiple memory addresses at once [3]. Most languages do not have dedicated abstractions for these instructions, so the compiler transforms parts of the program to use vector instructions, when possible.

Consider a simple Java loop shown on the right. When separate loop iterations have no data dependencies, a compiler

can aggregate arithmetic opera-

tions from several consecutive iterations into a single vector operation. However, on the JVM, the code must also perform a bounds check on each array access, which introduces additional guards into the loop, and normally prevents vectorization. Speculative guard motion makes loop vectorization possible, since it moves the bounds-check guards outside of the loop.

We found that combining loop vectorization with speculative guard motion results in a  $\approx 10\%$  impact on *als*, and a  $\approx 3\%$  impact on *dec-tree*.

#### 5.7 Dominance-Based Duplication Simulation

| Dominance-based duplication  |
|------------------------------|
| simulation (DBDS) [56] is an |
| optimization that simplifies |

1 if (x instanceof C) a()
2 else b()
3 if (x instanceof C) c()

control flow by duplicating the code after control-flow merges. Consider the pattern on the right, which consists of two subsequent instanceof-checks on the same class C.

After duplicating the second instanceof-check within both preceding branches, this second check becomes dominated by the first check. It can therefore be eliminated, resulting in the code shown on the left.

| 1 <b>if</b> (x : | instanceof | C) | { |
|------------------|------------|----|---|
| 2 a();           | c ()       |    |   |
| 3 } else         | b()        |    |   |

Since tail-duplication optimizations typically enable other optimizations by devirtualizing

the callsites, we decided to include DBDS in the impact analysis. We found that this optimization mostly affects *streamsmnemonics*, with a performance impact of around 22%.

## 6 Performance Evaluation

To investigate how the new optimization opportunities in Renaissance compare to the other suites, we provide a comprehensive performance evaluation that identifies the impact of each optimization discussed in Section 5 on each individual benchmark. As a consistent measure of impact across different optimizations and benchmarks, we define the impact of an optimization as the change in the benchmark execution time observed when the optimization is selectively disabled. The measure thus accounts for the cumulative impact of an optimization, including its possible (enabling or disabling) effects on other optimizations performed by the compiler.

The results are summarized in Figure 5. Evaluation results shown in this section have been obtained using the G1 collector. More details on the experimental setup used for collecting the measurements are described in Section C of the supplemental material. The figure supports our claim that the optimizations outlined in Section 5 benefit Renaissance more than other benchmark suites – at significance level  $\alpha = 0.01$ , all (7 of 7) evaluated optimizations have an impact of at least 5% on some Renaissance benchmark, compared with 2 of 7 for ScalaBench, 1 of 7 for DaCapo, and 3 of 7 for SPECjvm2008. At the same significance level, the median impact is 6.4% for Renaissance, compared with 2.8% for ScalaBench, 1.8% for DaCapo, and 3.9% for SPECjvm2008. In Section G of the supplemental material, we also provide some information on how each optimization impacts the overall compilation time.

The architecture of Graal enables rapid development and deployment of new optimizations, making it an obvious choice for our experiments with optimization opportunities in modern workloads. We found it relatively straightforward to implement new optimizations in the open-source Graal compiler. However, if Graal were a poorly performing compiler, the impact of each optimization relative to baseline might appear magnified. To prove that Graal can serve as a reasonable baseline, we also compared performance of the four benchmark suites on two variants of the HotSpot JVM, one equipped with the standard C2 compiler, other with Graal, both in the role of the final-tier compiler. The results, relative to the baseline HotSpot JVM using the C2 compiler, are shown in Figure 6. Graal provides better performance than C2 in 51 out of 68 benchmarks, the opposite is true in 10 benchmarks, for the remaining 7 benchmarks the difference is not statistically significant as the 99% confidence interval straddles 1.0. On benchmarks where Graal is better than C2, the median speedup is 20%. On benchmarks where C2 is better than Graal, the median slowdown is 4%. This justifies our choice of Graal as the experimental evaluation platform.

## 7 Code Complexity

In this section, we compare the software complexity and the compiled code size of Renaissance against the one of DaCapo, ScalaBench and SPECjvm2008.

#### 7.1 Software Complexity Metrics

Here, we compare Renaissance against the other benchmark suites by computing the Chidamber and Kemerer (CK) metrics [27], which were previously used to evaluate software complexity of benchmarks from the DaCapo and  $CD_x$  [47] suites. We used the *ckjm* [104] tool to calculate the metrics. To analyze only the classes used in each benchmark (and not all the classes in the classpath), we also developed a JVMTI native agent [65] for the HotSpot JVM that receives class-loading events from the VM, and forwards the loaded classes to *ckjm*, which then calculates the metrics.

We use the following CK metrics:

1. *Weighted methods per class* (**WMC**), which is calculated as the number of declared methods per loaded class.

2. *Depth of the inheritance tree* (**DIT**), defined as the depth of the inheritance tree along all loaded classes.

3. *Number of children* (**NOC**) is computed as the number of immediate subclasses of a class.

4. *Coupling between object classes* (CBO) counts the number of classes coupled to a given class, either by method calls, field accesses, inheritance, method signatures or exceptions.
5. *Response for a class* (RFC) counts the number of different methods that are potentially (transitively) executed when a single method is invoked.

6. *Lack of cohesion* (**LCOM**) counts how many methods of a class are not related by accessing the same field of that class.

The summarized results of the analysis are shown in Table 4. For each benchmark suite, we report a geometric mean, the minimum, and the maximum for the sum and for the average (arithmetic mean) of each metric across all benchmarks of the suite. When looking at the sums, we overall observed the highest values on five out of six metrics on Renaissance, while the LCOM metric was the highest on ScalaBench. With respect to the average complexity across all benchmarks of the suite, we can see that ScalaBench has higher values than Renaissance in 5 cases. In general we cannot say whether Renaissance is more complex than ScalaBench. However,



**Figure 5.** Optimization impact on individual benchmarks. Results with black outline significant at  $\alpha = 0.01$ .



Figure 6. Graal compiler performance relative to HotSpot second tier JIT compiler.

we see that Renaissance is in the same ball park as ScalaBench and the other suites. The reason why Renaissance is more complex on some of the metrics is highly related to the number of loaded classes of the benchmark suites. We took a look at the sum of all loaded classes across each benchmark suite and the number of unique loaded classes for each benchmark suite in Table 5. We observed that Renaissance benchmarks on average load many classes, compared to the other benchmark suites.

#### 7.2 Compiled Code Size

This section presents several metrics about the size of the compiled programs. Figure 7 shows the code size and compiled-method count distribution for the DaCapo, ScalaBench, SPEC-jvm2008, and Renaissance suites. Both the code size and the

method count refer to the hot code, compiled with the secondtier optimizing compiler, and were measured using Oracle's Graal compiler [7] by letting the benchmark run for 120 seconds.

The JVM uses the method's invocation count and its loop iteration counts to decide if a method is hot (i.e. frequently executed). Hot methods are compiled using the second-tier optimizing compiler, so the total size and the method count of the hot code indicate how much of the program was deemed important for the overall performance.

The geomean of the compiled code size (across all benchmarks) is  $\approx 6.87$ MB for Renaissance, which is slightly lower than  $\approx 7.98$ MB for DaCapo, and  $\approx 10.03$ MB for ScalaBench. The geomean of the total number of hot methods is  $\approx 1.636$ for Renaissance, which is slightly higher than  $\approx 1.599$  for DaCapo, and somewhat lower than  $\approx 1.853$  for Scalabench.

|                 | WMC              | DIT   | СВО       | NOC   | RFC    | LCOM    |
|-----------------|------------------|-------|-----------|-------|--------|---------|
| Renaissance     |                  |       |           |       |        |         |
| min-sum         | 21830            | 3066  | 21757     | 1571  | 41799  | 455958  |
| max-sum         | 206933           | 23936 | 184901    | 14029 | 369131 | 7360650 |
| geomean-sum     | 55533            | 7842  | 55147     | 4212  | 105105 | 1358042 |
| min-avg         | 11.07            | 1.79  | 12.57     | 0.97  | 20.78  | 141.02  |
| max-avg         | 18.48            | 2.29  | 17.54     | 1.16  | 33.84  | 706.54  |
| geomean-avg     | 13.58            | 1.92  | 13.49     | 1.03  | 25.71  | 332.19  |
|                 |                  | DaC   | Capo [21] |       |        |         |
| min-sum         | 12466            | 1687  | 10640     | 857   | 22378  | 161333  |
| max-sum         | 124191           | 5753  | 45597     | 6421  | 97757  | 616524  |
| geomean-sum     | 32470            | 3377  | 23275     | 2160  | 48461  | 336192  |
| min-avg         | 11.74            | 1.24  | 11.90     | 0.97  | 22.07  | 130.36  |
| max-avg         | 42.43            | 2.13  | 14.45     | 2.19  | 32.27  | 310.86  |
| geomean-avg     | 17.97            | 1.87  | 12.88     | 1.20  | 26.82  | 186.05  |
|                 | ScalaBench [100] |       |           |       |        |         |
| min-sum         | 24693            | 2657  | 19713     | 1481  | 45364  | 1080228 |
| max-sum         | 150895           | 7789  | 66036     | 7890  | 124240 | 2298594 |
| geomean-sum     | 44505            | 4291  | 32810     | 2735  | 71840  | 1489515 |
| min-avg         | 14.04            | 1.70  | 12.67     | 1.00  | 27.82  | 440.66  |
| max-avg         | 42.60            | 1.97  | 16.17     | 2.49  | 34.26  | 836.09  |
| geomean-avg     | 18.85            | 1.82  | 13.90     | 1.16  | 30.43  | 631.02  |
| SPECjvm2008 [1] |                  |       |           |       |        |         |
| min-sum         | 30560            | 3832  | 29966     | 2190  | 65741  | 546396  |
| max-sum         | 55044            | 5744  | 45745     | 3480  | 103373 | 1131251 |
| geomean-sum     | 33195            | 4142  | 32187     | 2383  | 70280  | 578408  |
| min-avg         | 13.55            | 1.72  | 13.48     | 0.99  | 28.69  | 206.19  |
| max-avg         | 16.90            | 1.86  | 14.04     | 1.07  | 31.73  | 347.22  |
| geomean-avg     | 14.00            | 1.75  | 13.58     | 1.01  | 29.65  | 244.03  |

**Table 4.** CK Metrics Summary: Minimum, maximum and geometric mean across the sum and arithmetic mean of all loaded classes of a benchmark suite. Bold values indicate the highest value for a metric across all benchmark suites.

| Benchmark Suite | Sum All<br>Loaded Classes | Sum Unique<br>Loaded Classes |
|-----------------|---------------------------|------------------------------|
| Renaissance     | 109 004                   | 29 157                       |
| DaCapo          | 27 923                    | 12073                        |
| ScalaBench      | 27 911                    | 12 207                       |
| SPECjvm2008     | 50 099                    | 5 274                        |

Table 5. Loaded classes per benchmark suite.

While DaCapo, ScalaBench and Renaissance are roughly in the same range, the SPECjvm2008 has a geomean of  $\approx$ 486 hot methods, and a geomean code size of  $\approx$ 1.17MB (for readability, the names of individual SPECjvm2008 benchmarks are not shown in Figure 7). This indicates that the SPECjvm2008 workloads are considerably smaller.

## 8 Related Work

By differentiating between the optimizations that are worth pursuing, benchmark suites were at the core of innovation throughout the JVM's history. Since its introduction in 2006, the DaCapo suite [21] has been a de facto standard for JVM benchmarking. While much of the original motivation for the DaCapo suite was to understand object and memory behavior in complex Java applications, this suite is still actively used to evaluate not only JVM components such as JIT compilers [34, 56, 83, 85, 106] and garbage collectors [17, 63],



Figure 7. Code Size Metrics

but also tools such as profilers [25, 98], data-race detectors [20, 110], memory monitors and contention analyzers [42, 109], static analyzers [37, 107], and debuggers [58].

The subsequently proposed ScalaBench suite [99, 100] identified a range of typical Scala programs, and argued that Scala and Java programs have considerably different distributions of instructions, polymorphic calls, object allocations, and method sizes. This observation that benchmark suites tend to over-represent certain programming styles was also noticed in other languages, (e.g., JavaScript [95]). On the other hand, the SPECjvm2008 benchmark suite [1] focused more on the core Java functionality. Most of the SPECjvm2008 benchmarks are considerably smaller than the DaCapo and ScalaBench benchmarks, and do not use a lot of object-oriented abstractions – SPECjvm2008 exercises classic JIT compiler optimizations, such as instruction scheduling and loop optimizations [31].

We did not include SPECjbb2015 [4] in our analysis, because it consists of a single benchmark designed to run over 2 hours, typically executed on multiple JVM instances. Using SPECjbb2015 to detect small statistically significant changes is therefore relatively costly (needs long execution time).

A brief evaluation suggests that the combined effect of all optimizations considered here accounts for about 1.2% improvement in the reported peak throughput (max jOPS) in single JVM configuration.

Another older JVM benchmark suite is Java Grande [103], but it consists mostly of microbenchmarks.

The tuning of compilers such as C2 [68] and Graal [7, 32] was heavily influenced by the DaCapo, ScalaBench, and SPECjvm2008 suites. Given that these existing benchmark suites do not exercise many frameworks and language extensions that gained popularity in the recent years, we looked for workloads exercising frameworks such as Java Streams [33] and Parallel Collections [81, 92, 93], Reactive Extensions [15], Akka [5], Scala actors [39] and Reactors [72, 75, 84, 90], coroutines [8, 86, 87], Apache Spark [111], futures and promises [40], Netty [11], Twitter Finagle [12], and Neo4J [10]. Most of these frameworks either assist in structuring concurrent programs, or enable programmers to declaratively specify data-parallel processing tasks. In both cases, they achieve these goals by providing a higher level of abstraction - for example, Finagle supports functional-style composition of future values, while Apache Spark exposes data-processing combinators for distributed datasets. By inspecting the IR of the open-source Graal compiler (c.f. Section 5), we found that many of the benchmarks exercise the interaction between different types of JIT compiler optimizations: optimizations, such as inlining, duplication [56], and partial escape analysis [106], typically start by reducing the level of abstraction in these frameworks, and then trigger more low-level optimizations such as guard motion [31], vectorization, or atomic-operation coalescing. Aside from a challenge in dealing with highlevel abstractions, the new concurrency primitives in modern benchmarks pose new optimization opportunities, such as contention elimination [59], application-specific workstealing [89], NUMA-aware node replication [26], speculative spinning [73], access path caching [74, 76–78], or other traditional compiler optimizations applied to concurrent programs [108]. Many of these newer optimizations may be applicable to domains such as concurrent data structures, which have been extensively studied on the JVM [18, 23, 28, 52, 67, 71, 79, 80, 82, 88, 91].

Unlike some other suites whose goal was to simulate deployment in clusters and Cloud environments, such as Cloud-Suite [35], our design decision was to follow the philosophy of DaCapo and ScalaBench, in which benchmarks are executed within a single JVM instance, whose execution characteristics can be understood more easily. Still, we found some alternative suites useful: for example, we took the *movielens* benchmark for Apache Spark from CloudSuite, and we adapted it to use Spark's single-process mode. The Unbalanced Cobwebbed Tree (UCT) benchmark [112] was designed for better actor-scheduler comparison on nonuniform workloads, and our *akka-uct* benchmark is the implementation of that benchmark in the Akka actor framework for the JVM [5]. A suite that specialized exclusively on benchmarks for the actor model is the Savina suite [43].

Several other benchmarks were either inspired by or adapted from existing workloads. The naive-bayes, log-regression, als, dec-tree and chi-square benchmarks directly work with several machine-learning algorithms from Apache Spark ML-Lib, and some of these benchmarks were inspired by the SparkPerf suite [29]. The Shakespeare plays Scrabble benchmark [69] was presented by José Paumard at the Virtual Technology Summit 2015 to demonstrate an advanced usage of Java Streams, and we directly adopted it as our scrabble benchmark. The rx-scrabble is a version of the scrabble benchmark that uses the Reactive Extensions framework instead of Java Streams. The streams-mnemonics benchmark is rewritten from the Phone Mnemonics benchmark that was originally used to demonstrate the usage of Scala collections [64]. The stm-bench7 benchmark is STMBench7 [38] applied to ScalaSTM [22, 24], a software transactional memory implementation for Scala, while the philosophers benchmark is ScalaSTM's Reality-Show Philosophers usage example.

#### 9 Conclusion

We presented Renaissance - a benchmark suite consisting of 21 benchmarks that represent modern concurrency and parallelism workloads, written in popular Java and Scala frameworks. To obtain these benchmarks, we gathered a large list of candidate workloads, both manually and by scanning an online corpus of GitHub projects. We then defined a set of basic metrics to filter potentially interesting workloads, and to ensure that the selection is sufficiently diverse. Our PCA analysis revealed that the set of benchmarks selected this way covers the metric space differently than the existing benchmark suites. To confirm that the thus-selected benchmarks are useful, we then analyzed them for performance-critical patterns, which lead us to implement four new optimizations in the Graal JIT compiler. These optimizations have a considerably smaller impact on existing suites, such as DaCapo, ScalaBench and SPECjvm2008, indicating that Renaissance helped in identifying new compiler optimizations. We also identified three existing optimizations whose performance impact is prominent. Furthermore, by comparing two production-quality JIT compilers, Graal and HotSpot C2, we determined that performance varies much more on Renaissance than on DaCapo and SPECjvm2008.

While we showed that Renaissance aids in identifying new compiler optimizations, we believe that the suite might be beneficial for other domains as well, such as garbage collectors, profilers, data-race detectors, and debuggers. We leave these and other investigations to future work.

#### PLDI '19, June 22-26, 2019, Phoenix, AZ, USA

## References

- [1] 2008. SPECjvm2008. https://www.spec.org/jvm2008/.
- [2] 2008. SPECjvm2008 User's Guide. https://www.spec.org/jvm2008/ docs/UserGuide.html
- [3] 2013. AVX 512 Instructions. https://software.intel.com/en-us/blogs/ 2013/avx-512-instructions.
- [4] 2015. SPECjbb2015. https://www.spec.org/jbb2015/.
- [5] 2018. Akka Documentation. http://akka.io/docs/.
- [6] 2018. Graal JTTTest Source Code. https://github.com/oracle/graal/ blob/master/compiler/src/org.graalvm.compiler.jtt/src/org/graalvm/ compiler/jtt/JTTTest.java.
- [7] 2018. GraalVM Website. https://www.graalvm.org/downloads/
- [8] 2018. Kotlin Coroutines. https://github.com/Kotlin/kotlinx. coroutines/blob/master/coroutines-guide.md. Accessed: 2018-11-15.
- [9] 2018. Open-Source Jenetics Repository at GitHub. https://github. com/jenetics/jenetics.
- [10] 2018. Open-Source Neo4J Repository at GitHub. https://github.com/ neo4j/neo4j.
- [11] 2018. Open-Source Netty Repository at GitHub. https://github.com/ netty/netty.
- [12] 2018. Open-Source Twitter Finagle Repository at GitHub. https: //github.com/twitter/finagle.
- [13] 2018. OpenJDK SynchronizedList Source Code. http: //hg.openjdk.java.net/jdk8/jdk8/jdk/file/687fd7c7986d/src/share/ classes/java/util/Collections.java.
- [14] 2018. OpenJDK Vector Source Code. http://hg.openjdk.java.net/jdk8/ jdk8/jdk/file/687fd7c7986d/src/share/classes/java/util/Vector.java.
- [15] 2018. ReactiveX project. http://reactivex.io/languages.html.
- [16] 2018. RxJava repository. https://github.com/ReactiveX/RxJava.
- [17] Shoaib Akram, Jennifer B. Sartor, Kathryn S. McKinley, and Lieven Eeckhout. 2018. Write-rationing Garbage Collection for Hybrid Memories. In PLDI. 62–77.
- [18] Dmitry Basin, Edward Bortnikov, Anastasia Braginsky, Guy Golan-Gueta, Eshcar Hillel, Idit Keidar, and Moshe Sulamy. 2017. KiWi: A Key-Value Map for Scalable Real-Time Analytics. In Proceedings of the 22Nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '17). ACM, New York, NY, USA, 357–369. https://doi.org/10.1145/3018743.3018761
- [19] Alan Bateman and Doug Lea. 2011. Java Specification Request 203: More New I/O APIs for the JavaTM Platform ("NIO.2"). https: //jcp.org/en/jsr/detail?id=203.
- [20] Swarnendu Biswas, Man Cao, Minjia Zhang, Michael D. Bond, and Benjamin P. Wood. 2017. Lightweight Data Race Detection for Production Runs. In CC. 11–21.
- [21] Stephen M. Blackburn, Robin Garner, Chris Hoffmann, Asjad M. Khang, Kathryn S. McKinley, Rotem Bentzur, Amer Diwan, Daniel Feinberg, Daniel Frampton, Samuel Z. Guyer, Martin Hirzel, Antony Hosking, Maria Jump, Han Lee, J. Eliot B. Moss, Aashish Phansalkar, Darko Stefanović, Thomas VanDrunen, Daniel von Dincklage, and Ben Wiedermann. 2006. The DaCapo Benchmarks: Java Benchmarking Development and Analysis. *SIGPLAN Not.* 41, 10 (Oct. 2006), 169–190.
- [22] Nathan Bronson, Jonas Boner, Guy Korland, Aleksandar Prokopec, Krishna Sankar, Daniel Spiewak, and Peter Veentjer. 2011. ScalaSTM Expert Group. https://nbronson.github.io/scala-stm/expert\_group. html.
- [23] Nathan G. Bronson, Jared Casper, Hassan Chafi, and Kunle Olukotun. 2010. A Practical Concurrent Binary Search Tree. *SIGPLAN Not.* 45, 5 (Jan. 2010), 257–268. https://doi.org/10.1145/1837853.1693488
- [24] Nathan G. Bronson, Hassan Chafi, and Kunle Olukotun. 2010. CCSTM: A library-based STM for Scala. In *The First Annual Scala Workshop at Scala Days*.
- [25] Rodrigo Bruno and Paulo Ferreira. 2017. POLM2: Automatic Profiling for Object Lifetime-aware Memory Management for Hotspot Big Data Applications. In *Middleware*. 147–160.

- [26] Irina Calciu, Siddhartha Sen, Mahesh Balakrishnan, and Marcos K. Aguilera. 2017. Black-box Concurrent Data Structures for NUMA Architectures. In Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '17). ACM, New York, NY, USA, 207–221. https://doi.org/10.1145/3037697.3037721
- [27] Shyam R Chidamber and Chris F Kemerer. 1994. A Metrics Suite for Object Oriented Design. *IEEE Transactions on Software Engineering* 20, 6 (1994), 476–493.
- [28] Cliff Click. 2007. Towards a Scalable Non-Blocking Coding Style. http://www.azulsystems.com/events/javaone\_2007/2007\_ LockFreeHash.pdf
- [29] Databricks. 2018. Spark Performance Tests. https://github.com/ databricks/spark-perf.
- [30] Pedro C. Diniz and Martin C. Rinard. 1998. Lock Coarsening: Eliminating Lock Overhead in Automatically Parallelized Object-based Programs. J. Parallel and Distrib. Comput. 49, 2 (1998), 218–244.
- [31] Gilles Duboscq. 2016. Combining Speculative Optimizations with Flexible Scheduling of Side-effects. Ph.D. Dissertation. Johannes Kepler University, Linz.
- [32] Gilles Duboscq, Thomas Würthinger, Lukas Stadler, Christian Wimmer, Doug Simon, and Hanspeter Mössenböck. 2013. An Intermediate Representation for Speculative Optimizations in a Dynamic Compiler. In VMIL. 1–10.
- [33] Michael Duigou. 2011. Java Enhancement Proposal 107: Bulk Data Operations for Collections. http://openjdk.java.net/jeps/107.
- [34] Josef Eisl, Matthias Grimmer, Doug Simon, Thomas Würthinger, and Hanspeter Mössenböck. 2016. Trace-based Register Allocation in a JIT Compiler. In *PPPJ*. 14:1–14:11.
- [35] Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafaee, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, and Babak Falsafi. 2012. Clearing the Clouds: A Study of Emerging Scale-out Workloads on Modern Hardware. SIGPLAN Not. 47, 4 (March 2012), 37–48.
- [36] Rémi Forax, Vladimir Zakharov, Kevin Bourrillion, Dan Heidinga, Srikanth Sankaran, Andrey Breslav, Doug Lea, Bob Lee, Brian Goetz, Daniel Smith, Samuel Pullara, and David Lloyd. 2014. Java Specification Request 335: Lambda Expressions for the JavaTM Programming Language. https://jcp.org/en/jsr/detail?id=335.
- [37] Neville Grech, George Fourtounis, Adrian Francalanza, and Yannis Smaragdakis. 2018. Shooting from the Heap: Ultra-scalable Static Analysis with Heap Snapshots. In ISSTA. 198–208.
- [38] Rachid Guerraoui, Michal Kapalka, and Jan Vitek. 2007. STMBench7: A Benchmark for Software Transactional Memory. In *EuroSys.* 315– 324.
- [39] Philipp Haller and Martin Odersky. 2007. Actors That Unify Threads and Events. In COORDINATION. 171–190.
- [40] Philipp Haller, Aleksandar Prokopec, Heather Miller, Viktor Klang, Roland Kuhn, and Vojin Jovanovic. 2012. Scala Improvement Proposal: Futures and Promises (SIP-14). http://docs.scala-lang.org/sips/ pending/futures-promises.html
- [41] Tim Harris, Simon Marlow, Simon Peyton-Jones, and Maurice Herlihy. 2005. Composable Memory Transactions. In PPoPP. 48–60.
- [42] Peter Hofer, David Gnedt, Andreas Schörgenhumer, and Hanspeter Mössenböck. 2016. Efficient Tracing and Versatile Analysis of Lock Contention in Java Applications on the Virtual Machine Level. In *ICPE*. 263–274.
- [43] Shams M. Imam and Vivek Sarkar. 2014. Savina An Actor Benchmark Suite: Enabling Empirical Evaluation of Actor Libraries. In AGERE! 67–80.
- [44] Intel. 2018. Hyper-Threading Technology. https://www. intel.com/content/www/us/en/architecture-and-technology/ hyper-threading/hyper-threading-technology.html.

- [45] Intel. 2018. Turbo Boost Technology 2.0. https://www.intel.com/ content/www/us/en/architecture-and-technology/turbo-boost/ turbo-boost-technology.html.
- [46] Intel. 2019. Intel 64 and IA-32 Architectures Developer's Manual, Section 18.18. https://www.intel. com/content/www/us/en/architecture-and-technology/ 64-ia-32-architectures-software-developer-manual-325462.html.
- [47] Tomas Kalibera, Jeff Hagelberg, Petr Maj, Filip Pizlo, Ben Titzer, and Jan Vitek. 2010. A Family of Real-time Java Benchmarks. *Concurrency* and Computation: Practice and Experience 23, 14 (2010), 1679–1700.
- [48] Karl Pearson. 1901. On lines and planes of closest fit to systems of points in space. *Philos. Mag.* 2, 11 (1901), 559–572.
- [49] Stephen Kell, Danilo Ansaloni, Walter Binder, and Lukáš Marek. 2012. The JVM is Not Observable Enough (and What to Do About It). In VMIL. 33–38.
- [50] Doug Lea. 2000. A Java Fork/Join Framework. In JAVA. 36-43.
- [51] Doug Lea. 2012. Java Enhancement Proposal 155: Concurrency Updates. http://openjdk.java.net/jeps/155.
- [52] Doug Lea. 2014. Doug Lea's Workstation. http://g.oswego.edu/
- [53] Doug Lea. 2015. Java Enhancement Proposal 266: More Concurrency Updates. http://openjdk.java.net/jeps/266.
- [54] Doug Lea, Joshua Bloch, Sam Midkiff, David Holmes, Joseph Bowbeer, and Tim Peierls. 2004. Java Specification Request 166: Concurrency Utilities. https://jcp.org/ja/jsr/detail?id=166.
- [55] Philipp Lengauer, Verena Bitto, Hanspeter Mössenböck, and Markus Weninger. 2017. A Comprehensive Java Benchmark Study on Memory and Garbage Collection Behavior of DaCapo, DaCapo Scala, and SPECjvm2008. In *ICPE*. 3–14.
- [56] David Leopoldseder, Lukas Stadler, Thomas Würthinger, Josef Eisl, Doug Simon, and Hanspeter Mössenböck. 2018. Dominance-based Duplication Simulation (DBDS): Code Duplication to Enable Compiler Optimizations. In CGO. 126–137.
- [57] Linux man. 2013. top(1). https://linux.die.net/man/1/top.
- [58] Bozhen Liu and Jeff Huang. 2018. D4: Fast Concurrency Debugging with Parallel Differential Analysis. In PLDI. 359–373.
- [59] Honghui Lu, Alan L. Cox, and Willy Zwaenepoel. 2001. Contention Elimination by Replication of Sequential Sections in Distributed Shared Memory Programs. In Proceedings of the Eighth ACM SIGPLAN Symposium on Principles and Practices of Parallel Programming (PPoPP '01). ACM, New York, NY, USA, 53–61. https: //doi.org/10.1145/379539.379568
- [60] Lukáš Marek, Stephen Kell, Yudi Zheng, Lubomír Bulej, Walter Binder, Petr Tůma, Danilo Ansaloni, Aibek Sarimbekov, and Andreas Sewe. 2013. ShadowVM: Robust and Comprehensive Dynamic Program Analysis for the Java Platform. In *GPCE*. 105–114.
- [61] Lukáš Marek, Alex Villazón, Yudi Zheng, Danilo Ansaloni, Walter Binder, and Zhengwei Qi. 2012. DiSL: A Domain-specific Language for Bytecode Instrumentation. In AOSD. 239–250.
- [62] Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, Doris Xin, Reynold Xin, Michael J. Franklin, Reza Zadeh, Matei Zaharia, and Ameet Talwalkar. 2016. MLlib: Machine Learning in Apache Spark. *Journal of Machine Learning Research* 17, 34 (2016), 1–7.
- [63] Khanh Nguyen, Lu Fang, Guoqing Xu, Brian Demsky, Shan Lu, Sanazsadat Alamian, and Onur Mutlu. 2016. Yak: A High-performance Big-data-friendly Garbage Collector. In OSDI. 349–365.
- [64] Martin Odersky. 2011. State of Scala. http://days2011.scala-lang.org/ sites/days2011/files/01.%20Martin%20Odersky.pdf.
- [65] Oracle. 2018. JVM Tool Interface. https://docs.oracle.com/javase/8/ docs/platform/jvmti/jvmti.html.
- [66] F. Ortin, P. Conde, D. Fernandez-Lanvin, and R. Izquierdo. 2014. The Runtime Performance of invokedynamic: An Evaluation with a Java Library. *IEEE Software* 31, 4 (July 2014), 82–90.

- [67] Rotem Oshman and Nir Shavit. 2013. The SkipTrie: Low-depth Concurrent Search Without Rebalancing. In Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC '13). ACM, New York, NY, USA, 23–32. https://doi.org/10.1145/2484239.2484270
- [68] Michael Paleczny, Christopher Vick, and Cliff Click. 2001. The Java HotspotTM Server Compiler. In *JVM*.
- [69] José Paumard. 2018. JDK8 Stream/Rx Comparison. https://github. com/JosePaumard/jdk8-stream-rx-comparison.
- [70] perf. 2015. Linux profiling with performance counters. https://perf. wiki.kernel.org.
- [71] Aleksandar Prokopec. 2015. SnapQueue: Lock-free Queue with Constant Time Snapshots. In Proceedings of the 6th ACM SIGPLAN Symposium on Scala (SCALA 2015). ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/2774975.2774976
- [72] Aleksandar Prokopec. 2016. Pluggable Scheduling for the Reactor Programming Model. In AGERE! 41–50.
- [73] Aleksandar Prokopec. 2017. Accelerating by Idling: How Speculative Delays Improve Performance of Message-Oriented Systems. Springer International Publishing, Cham, 177–191. https://doi.org/10.1007/ 978-3-319-64203-1\_13
- [74] Aleksandar Prokopec. 2017. Analysis of Concurrent Lock-Free Hash Tries with Constant-Time Operations. ArXiv e-prints (Dec. 2017). arXiv:cs.DS/1712.09636
- [75] Aleksandar Prokopec. 2017. Encoding the Building Blocks of Communication. In Proceedings of the 2017 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2017). ACM, New York, NY, USA, 104–118. https://doi.org/10.1145/3133850.3133865
- [76] Aleksandar Prokopec. 2018. Cache-tries: Concurrent Lock-free Hash Tries with Constant-time Operations. In Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '18). ACM, New York, NY, USA, 137–151. https://doi.org/10.1145/3178487.3178498
- [77] Aleksandar Prokopec. 2018. Efficient Lock-Free Removing and Compaction for the Cache-Trie Data Structure. Springer International Publishing, Cham.
- [78] Aleksandar Prokopec. 2018. Efficient Lock-Free Removing and Compaction for the Cache-Trie Data Structure. https://doi.org/10.6084/m9.figshare.6369134.
- [79] Aleksandar Prokopec, Phil Bagwell, and Martin Odersky. 2011. Cache-Aware Lock-Free Concurrent Hash Tries. Technical Report.
- [80] Aleksandar Prokopec, Phil Bagwell, and Martin Odersky. 2011. Lock-Free Resizeable Concurrent Tries. Springer Berlin Heidelberg, Berlin, Heidelberg, 156–170. https://doi.org/10.1007/978-3-642-36036-7\_11
- [81] Aleksandar Prokopec, Phil Bagwell, Tiark Rompf, and Martin Odersky. 2011. A Generic Parallel Collection Framework. In *Euro-Par.* 136–147.
- [82] Aleksandar Prokopec, Nathan Grasso Bronson, Phil Bagwell, and Martin Odersky. 2012. Concurrent Tries with Efficient Non-blocking Snapshots. In Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '12). ACM, New York, NY, USA, 151–160. https://doi.org/10.1145/2145816.2145836
- [83] Aleksandar Prokopec, Gilles Duboscq, David Leopoldseder, and Thomas Würthinger. 2019. An Optimization-driven Incremental Inline Substitution Algorithm for Just-in-time Compilers. In Proceedings of the 2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO 2019). IEEE Press, Piscataway, NJ, USA, 164–179. http://dl.acm.org/citation.cfm?id=3314872.3314893
- [84] Aleksandar Prokopec, Philipp Haller, and Martin Odersky. 2014. Containers and Aggregates, Mutators and Isolates for Reactive Programming. In *Proceedings of the Fifth Annual Scala Workshop (SCALA '14)*. ACM, New York, NY, USA, 51–61. https://doi.org/10.1145/2637647. 2637656
- [85] Aleksandar Prokopec, David Leopoldseder, Gilles Duboscq, and Thomas Würthinger. 2017. Making Collection Operations Optimal with Aggressive JIT Compilation. In SCALA. 29–40.

- [86] Aleksandar Prokopec and Fengyun Liu. 2018. On the Soundness of Coroutines with Snapshots. CoRR abs/1806.01405 (2018). arXiv:1806.01405 https://arxiv.org/abs/1806.01405
- [87] Aleksandar Prokopec and Fengyun Liu. 2018. Theory and Practice of Coroutines with Snapshots. In 32nd European Conference on Object-Oriented Programming, ECOOP 2018, July 16-21, 2018, Amsterdam, The Netherlands. 3:1–3:32. https://doi.org/10.4230/LIPIcs.ECOOP.2018.3
- [88] Aleksandar Prokopec, Heather Miller, Tobias Schlatter, Philipp Haller, and Martin Odersky. 2012. FlowPools: A Lock-Free Deterministic Concurrent Dataflow Abstraction. In *LCPC*. 158–173.
- [89] Aleksandar Prokopec and Martin Odersky. 2014. Near Optimal Work-Stealing Tree Scheduler for Highly Irregular Data-Parallel Workloads. In *Languages and Compilers for Parallel Computing*, Călin Cascaval and Pablo Montesinos (Eds.). Springer International Publishing, Cham, 55–86.
- [90] Aleksandar Prokopec and Martin Odersky. 2015. Isolates, Channels, and Event Streams for Composable Distributed Programming. In Onward! 171–182.
- [91] Aleksandar Prokopec and Martin Odersky. 2016. Conc-Trees for Functional and Parallel Programming. Springer International Publishing, Cham, 254–268. https://doi.org/10.1007/978-3-319-29778-1\_16
- [92] Aleksandar Prokopec, Dmitry Petrashko, and Martin Odersky. 2014. On Lock-Free Work-stealing Iterators for Parallel Data Structures. (2014), 10.
- [93] A. Prokopec, D. Petrashko, and M. Odersky. 2015. Efficient Lock-Free Work-Stealing Iterators for Data-Parallel Collections. In 2015 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing. 248–252. https://doi.org/10.1109/PDP. 2015.65
- [94] William Pugh, Sarita Adve, and Doug Lea. 2004. Java Specification Request 133: JavaTM Memory Model and Thread Specification Revision. https://jcp.org/ja/jsr/detail?id=133.
- [95] Paruj Ratanaworabhan, Benjamin Livshits, and Benjamin G. Zorn. 2010. JSMeter: Comparing the Behavior of JavaScript Benchmarks with Real Web Applications. In *WebApps*. 3–3.
- [96] John Rose. 2014. Java Enhancement Proposal 243: Java-Level JVM Compiler Interface. http://openjdk.java.net/jeps/243.
- [97] John Rose, Bini Ola, William Cook, Rémi Forax, Samuele Pedroni, and Jochen Theodorou. 2011. Java Specification Request 292: Supporting Dynamically Typed Languages on the JavaTM Platform. https: //jcp.org/en/jsr/detail?id=292.
- [98] Andreas Schörgenhumer, Peter Hofer, David Gnedt, and Hanspeter Mössenböck. 2017. Efficient Sampling-based Lock Contention Profiling for Java. In *ICPE*. 331–334.

- [99] Andreas Sewe, Mira Mezini, Aibek Sarimbekov, Danilo Ansaloni, Walter Binder, Nathan Ricci, and Samuel Z. Guyer. 2012. new Scala() instanceof Java: A Comparison of the Memory Behaviour of Java and Scala Programs. In *ISMM*. 97–108.
- [100] Andreas Sewe, Mira Mezini, Aibek Sarimbekov, and Walter Binder. 2011. Da Capo Con Scala: Design and Analysis of a Scala Benchmark Suite for the Java Virtual Machine. In OOPSLA. 657–676.
- [101] Nir Shavit and Dan Touitou. 1995. Software Transactional Memory. In PODC. 204–213.
- [102] Aleksei Shipilev. 2018. Code Tools: jmh. http://openjdk.java.net/ projects/code-tools/jmh/.
- [103] L. A. Smith, J. M. Bull, and J. Obdrizalek. 2001. A Parallel Java Grande Benchmark Suite. In SC.
- [104] Diomidis Spinellis. 2005. Tool Writing: A Forgotten Art? IEEE Software 4 (2005), 9–11.
- [105] Sriram Srinivasan and Alan Mycroft. 2008. Kilim: Isolation-Typed Actors for Java. In ECOOP. 104–128.
- [106] Lukas Stadler, Thomas Würthinger, and Hanspeter Mössenböck. 2014. Partial Escape Analysis and Scalar Replacement for Java. In CGO. 165:165–165:174.
- [107] Rei Thiessen and Ondřej Lhoták. 2017. Context Transformations for Pointer Analysis. In PLDI. 263–277.
- [108] Jaroslav Ševčík. 2011. Safe Optimisations for Shared-memory Concurrent Programs. In Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '11). ACM, New York, NY, USA, 306–316. https://doi.org/10.1145/1993498. 1993534
- [109] Markus Weninger and Hanspeter Mössenböck. 2018. User-defined Classification and Multi-level Grouping of Objects in Memory Monitoring. In *ICPE*. 115–126.
- [110] Benjamin P. Wood, Man Cao, Michael D. Bond, and Dan Grossman. 2017. Instrumentation Bias for Dynamic Data Race Detection. *Proc.* ACM Program. Lang. 1, OOPSLA (Oct. 2017), 69:1–69:31.
- [111] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster Computing with Working Sets. In *HotCloud*. 10–10.
- [112] Xinghui Zhao and Nadeem Jamali. 2013. Load Balancing Non-uniform Parallel Computations. In AGERE! 97–108.
- [113] Yudi Zheng, Andrea Rosà, Luca Salucci, Yao Li, Haiyang Sun, Omar Javed, Lubomír Bulej, Lydia Y. Chen, Zhengwei Qi, and Walter Binder. 2016. AutoBench: Finding Workloads That You Need Using Pluggable Hybrid Analyses.. In SANER. 639–643.

**Table 6.** Benchmarks considered in the paper, including input size and number of operations executed (# ops).

| Benchmark                    | Input   | Benchmark            | Input      |  |  |
|------------------------------|---------|----------------------|------------|--|--|
| DaCapo [21]                  | size    | ScalaBench [100]     | size       |  |  |
| avrora                       | large   | actors               | huge       |  |  |
| batik                        | large   | apparat              | gargantuan |  |  |
| eclipse                      | large   | factorie             | gargantuan |  |  |
| fop                          | default | kiama                | default    |  |  |
| h2                           | huge    | scalac               | large      |  |  |
| jython                       | large   | scaladoc             | large      |  |  |
| luindex                      | default | scalap               | large      |  |  |
| lusearch-fix                 | large   | scalariform          | huge       |  |  |
| pmd                          | large   | scalatest            | default    |  |  |
| sunflow                      | large   | scalaxb              | huge       |  |  |
| tomcat                       | huge    | specs                | large      |  |  |
| tradebeans                   | huge    | tmt                  | huge       |  |  |
| tradesoap                    | huge    |                      |            |  |  |
| xalan                        | large   |                      |            |  |  |
| Benchmark<br>SPECjvm2008 [1] | # ops   | Benchmark            | # ops      |  |  |
| compiler.compiler            | 20      | scimark.lu.small     | 125        |  |  |
| compiler.sunflow             | 20      | scimark.monte_carlo  | 900        |  |  |
| compress                     | 50      | scimark.sor.large    | 15         |  |  |
| crypto.aes                   | 20      | scimark.sor.small    | 75         |  |  |
| crypto.rsa                   | 150     | scimark.sparse.large | 10         |  |  |
| crypto.signverify            | 125     | scimark.sparse.small | 25         |  |  |
| derby                        | 30      | serial               | 25         |  |  |
| mpegaudio                    | 50      | sunflow              | 30         |  |  |
| scimark.fft.large            | 10      | xml.transform        | 7          |  |  |
| scimark.fft.small            | 100     | xml.validation       | 40         |  |  |
| scimark.lu.large             | 4       |                      |            |  |  |

## A Analyzed Benchmarks

Table 6 lists the benchmarks from the DaCapo, ScalaBench and SPECjvm2008 suites that were considered in the main paper, along with the used input size (expressed as number of operations executed in SPECjvm2008).

## **B** Experimental Setup for Metric Profiling and Principal Component Analysis

Here, we detail the experimental setup for the collection of metrics described in Table 2 and analyzed in Section 4 and 5 of the main paper.

The metrics are profiled during a single steady-state benchmark execution. Before collecting the metrics, we let the benchmarks warp-up until dynamic compilation and GC ergonomics are stabilized, following the methodology of Lengauer et al. [55]. We could not collect metrics for benchmarks *tradebeans*, *actors* and *scimark.monte\_carlo* either because bytecode instrumentation causes a premature workload termination with a TimeoutException (*tradebeans*, *actors*) or because profiling takes an excessive amount of time, exceeding 7 days (*scimark.monte\_carlo*). Therefore we excluded such benchmarks from the PCA analysis (Section 4 of the main paper). We collect the metrics on a machine with two NUMA nodes, each containing an Intel Xeon E5-2680 (2.7 GHz) processor with 8 physical cores and 64 GB of RAM, running under Ubuntu 16.04.03 LTS (kernel GNU/Linux 4.4.0-112-generic x86\_64). We configure top to sample CPU utilization only for the NUMA node where the benchmark is executing, to increase the accuracy of the collected measurements (as the computational resources used by perf and top are not accounted). We disable Turbo Boost [45] and Hyper-Threading [44]. We use Java OpenJDK 1.8.0\_161-b12.

We collect the metrics in two runs, profiling OS- and hardware-layer metrics (cpu and cachemiss) in the first run on the original program, and the other metrics in the second run (using DiSL instrumentation). This way, we obtain more precise metrics at the OS- and hardware-layer, which do not account for the execution of instrumentation code. During metric collection, no other CPU-, memory-, or IO-intensive application is executing on the system to reduce measurement perturbations. In addition, we pin the execution to an exclusive NUMA node, to reduce performance interference caused by other running processes.

## C Experimental Setup for Performance Evaluation

Here, we describe the experimental setup for the performance evaluation described in Section  $\frac{6}{6}$  of the main paper.

The performance measurement experiments are conducted on 8-core Intel servers, equipped with an Intel Xeon E5-2620v4 CPU (2.1 GHz, 8 cores, 20 MB cache, Hyper Threading disabled), 64 GB RAM, running Fedora Linux 27 (kernel 4.15.6). For stable measurement, power management features are disabled and the processor is run at the nominal frequency. Prior to each benchmark execution, the physical memory pool is randomized. We use Oracle JDK 8u172 with Graal 1.0.0-rc9 as virtual machine. The heap size is fixed at 12 GB with the G1 collector, and except for the selection of individual compiler optimizations used to produce Figure 5, no other option is used.

For each benchmark and each optimization configuration, we execute the measurements in a new JVM 15 times. Each execution consists of a warm-up period of 5 minutes, followed by 60 seconds of steady-state execution, rounded up to the next complete benchmark iteration. The duration of the warm-up period is chosen so that major performance fluctuations due to compilation happen before actual measurement (verified manually). To provide for meaningful comparison across benchmarks, we always collect the execution times of the main benchmark operation (we have modified the SPECjvm2008 benchmark harness to achieve this, the benchmark would normally report aggregated throughput). Winsorized filtering is used to remove outliers from Figure 5.

## **D** Collected Metrics

Table 7 reports the metrics (listed in Table 2) collected on all analyzed benchmarks, before being normalized by reference cycles. The experimental setup used for metric collection is detailed in Section **B** of the main paper.

#### **E** Principal Component Analysis

In Figure 8, we report a larger version of the scatter plots shown in Figure 1 and discussed in Section 4 of the main paper.

## F Additional Data for the Software Complexity Metrics

In Tables 8 to 11, we present additional data for the Chidamber & Kemerer metrics (Section 7.1 in the main paper), across all four suites. Tables 8 and 9 contain the sum for each metric across all benchmarks of a suite, while in Tables 10 and 11 we present the arithmetic mean for each metric across all benchmarks of a suite.

## G Additional Data for the Optimization Impact Measurements

In Tables 12 to 15, we provide numerical data for the optimization impact overview from Figure 5. The seven columns - AC, DS, EAWA, GM, LV, LLC and MHS – stand for the seven optimizations considered, namely Atomic-Operation Coalescing, Dominance-Based Duplication Simulation, Escape Analysis with Atomic Operations, Speculative Guard Motion, Loop Vectorization, Loop-Wide Lock Coarsening, and Method-Handle Simplification. In each column, the first number gives the change in benchmark execution times observed when the relevant optimization is turned off, relative to a baseline with all optimizations turned on (positive numbers mean optimization speeds up execution, negative numbers mean optimization slows down execution). The second number gives the p-value as computed by the Welch's t-test.

Table 16 provides estimate on the compilation overhead associated with each of the seven optimizations considered. In each row, the value gives the relative reduction in compiler thread execution time when the particular optimization is disabled, measured over the entire warm up period. The values are aggregated across all benchmarks.

| Benchmark            | synch                | wait                 | notify               | atomic               | park                 | cpu            | cachemiss            | object               | array                | method               | idynamic             |
|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------|----------------------|----------------------|----------------------|----------------------|----------------------|
|                      |                      |                      |                      |                      | Renaissanc           | e              |                      |                      |                      |                      |                      |
| akka-uct             | 4.27E+05             | 1.00E+00             | 2.00E+00             | 1.18E+07             | 1.75E+05             | 94.45          | 6.24E+08             | 1.16E+08             | 6.00E+05             | 1.96E+09             | 0.00E+00             |
| als                  | 3.01E+06             | 1.15E+02             | 1.31E+04             | 1.81E+06             | 2.39E+03             | 58.90          | 9.63E+08             | 1.10E+08             | 2.38E+07             | 2.91E+09             | 0.00E+00             |
| db-shootout          | 1.52E+06             | 8.70E+01             | 4.50E+01             | 1.58E+05<br>2.72E+07 | 7.20E+01             | 20.19<br>45.53 | 4.97E+08             | 1.75E+08             | 2.40E+07             | 2.39E+09             | 0.00E+00<br>1.00E+06 |
| dec-tree             | 5.83E±05             | 5.20L+01<br>8.80F±01 | 1.37E±03             | 5.45E±05             | 4.01L+03             | 45.55          | 2.91E+09             | 2.10L+08             | 2.84E±07             | 2.96F±09             | 0.00E+00             |
| dotty                | 5.63E+06             | 4 00E+01             | 2.56E+04             | 4 33E+04             | 0.00E+00             | 15.68          | 7.59E+08             | 4 92E+07             | 1 42E+07             | 1.26E+09             | 7.22E+06             |
| finagle-chirper      | 1.29E+07             | 1.72E+03             | 1.74E+03             | 1.02E+08             | 1.72E+04             | 69.82          | 2.52E+09             | 1.43E+08             | 1.01E+07             | 4.44E+09             | 2.36E+03             |
| finagle-http         | 2.72E+04             | 2.00E+01             | 0.00E+00             | 5.20E+04             | 6.66E+02             | 14.72          | 4.14E+08             | 2.81E+08             | 6.40E+04             | 3.09E+09             | 5.80E+02             |
| fj-kmeans            | 1.01E+08             | 6.57E+02             | 6.62E+02             | 1.89E+04             | 1.19E+03             | 69.59          | 4.23E+08             | 1.35E+08             | 2.45E+03             | 7.08E+08             | 0.00E+00             |
| future-genetic       | 6.72E+05             | 2.37E+04             | 2.40E+04             | 5.00E+07             | 1.59E+05             | 55.85          | 7.04E+08             | 2.11E+08             | 2.64E+05             | 1.58E+09             | 2.34E+06             |
| log-regression       | 2.09E+05             | 1.09E+02             | 9.81E+02             | 4.77E+05             | 6.62E+02             | 24.83          | 6.58E+08             | 5.39E+07             | 1.68E+07             | 1.86E+09             | 0.00E+00             |
| movie-lens           | 1.24E+07             | 5.78E+02             | 2.22E+05             | 3.11E+07             | 3.98E+04             | 44.17          | 3.37E+09             | 2.00E+08             | 2.58E+07             | 7.32E+09             | 2.16E+02             |
| naive-bayes          | 2.33E+05             | 3.50E+01             | 1.09E+02             | 1.81E+04             | 1.32E+02             | 76.90          | 1.04E+09             | 3.60E+08             | 8.21E+07             | 3.65E+09             | 0.00E+00             |
| neo4j-analytics      | 1.37E+07             | 4.23E+02             | 1.54E+05             | 2.05E+06             | 2.02E+02             | 59.74          | 1.05E+18             | 1.42E+09             | 3.29E+07             | 2.22E+10             | 2.49E+07             |
| page-rank            | 2.03E+06             | 9.10E+01             | 1.26E+02             | 9.25E+06             | 1.38E+02             | 56.14          | 1.23E+09             | 2.01E+08             | 2.94E+06             | 5.15E+09             | 0.00E+00             |
| reactors             | 2.21E+00<br>2.50E+08 | 1.32E+04             | 0.00E+04             | 1.16E+08             | 2.32E+04             | 56.62          | 1.10E+09             | 2.71E+08             | 4.01L+07             | 0.20E+09             | 0.00E+00             |
| rx-scrabble          | 2.59E+06             | 0.00E+00             | 0.00E+00             | 7.49E+05             | 3.32E+00<br>8.90E+01 | 25.10          | 4.22E+09<br>8.11E+07 | 1.07E+07             | 0.00E+00             | 1.24L+10<br>1.02E+08 | 1.71E+06             |
| scrabble             | 2.00E+00             | 1.00E+00             | 1.00E+00             | 3.00E+01             | 1.00E+00             | 66.70          | 2.81E+08             | 5.65E+07             | 3.44E+06             | 4.99E+08             | 2.73E+07             |
| stm-bench7           | 3.56E+03             | 1.00E+01             | 3.00E+00             | 2.92E+06             | 0.00E+00             | 49.44          | 3.48E+08             | 3.03E+07             | 2.96E+06             | 8.15E+08             | 0.00E+00             |
| streams-mnemonics    | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 19.24          | 7.99E+08             | 2.04E+08             | 2.09E+08             | 1.15E+09             | 2.15E+07             |
|                      |                      |                      |                      |                      | DaCapo               |                |                      |                      |                      |                      |                      |
| avrora               | 7.22E+06             | 1.75E+06             | 1.68E+05             | 0.00E+00             | 0.00E+00             | 20.28          | 7.11E+17             | 7.22E+06             | 2.05E+06             | 2.78E+09             | 0.00E+00             |
| batik                | 1.67E+06             | 6.00E+00             | 0.00E+00             | 7.00E+00             | 0.00E+00             | 24.54          | 1.96E+08             | 8.77E+05             | 2.03E+05             | 3.46E+07             | 0.00E+00             |
| eclipse              | 6.80E+07             | 1.88E+04             | 3.60E+05             | 1.27E+05             | 0.00E+00             | 13.91          | 4.66E+09             | 9.06E+07             | 9.89E+07             | 2.41E+09             | 0.00E+00             |
| fop                  | 2.45E+06             | 0.00E+00             | 0.00E+00             | 1.60E+01             | 0.00E+00             | 6.25           | 5.10E+07             | 1.63E+06             | 7.12E+05             | 3.50E+07             | 0.00E+00             |
| h2                   | 7.76E+08             | 4.65E+03             | 0.00E+00             | 2.82E+07             | 0.00E+00             | 17.78          | 2.14E+10             | 2.91E+08             | 1.23E+08             | 2.62E+10             | 0.00E+00             |
| Jython               | 1.06E+08             | 0.00E+00             | 0.00E+00             | 1.72E+07             | 0.00E+00             | 11.41          | 8.86E+08             | 1.38E+08             | 2.80E+07             | 4.14E+09             | 0.00E+00             |
| lundex               | 2.77E+05             | 1.00E+00<br>1.38E+02 | 1.25E+05             | 1.00E+01<br>5.12E+02 | 0.00E+00             | 5.57<br>85.00  | 5.74E+07             | 1.85E+05<br>1.04E+07 | 0.40E+04             | 7.81E+07             | 0.00E+00             |
| nmd                  | 3.05E+06             | 0.00E+02             | 3.34F+03             | 4.62E+02             | 3.00E+00             | 22.26          | 4.06E+08             | 1.04E+07             | 2.86E+06             | 1.73E+08             | 0.00E+00             |
| sunflow              | 1 53E+03             | 5.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 79 55          | 1.13E+09             | 1.01E+07             | 4 34E+06             | 4 19E+09             | 0.00E+00             |
| tomcat               | 2.28E+08             | 6.04E+02             | 2.18E+05             | 7.84E+06             | 1.93E+05             | 27.51          | 1.61E+18             | 1.07E+08             | 7.61E+07             | 4.44E+09             | 0.00E+00             |
| tradesoap            | 7.31E+08             | 2.12E+02             | 1.29E+06             | 2.39E+06             | 1.30E+05             | 64.92          | 2.96E+10             | 6.64E+08             | 2.44E+08             | 1.50E+10             | 1.40E+02             |
| xalan                | 2.12E+08             | 3.48E+02             | 1.01E+05             | 0.00E+00             | 0.00E+00             | 97.89          | 5.11E+09             | 6.12E+07             | 4.00E+07             | 3.84E+09             | 0.00E+00             |
|                      |                      |                      |                      |                      | ScalaBencl           | ı              |                      |                      |                      |                      |                      |
| apparat              | 1.35E+07             | 5.64E+03             | 5.16E+05             | 1.19E+06             | 4.54E+04             | 15.80          | 2.69E+10             | 3.22E+08             | 2.55E+07             | 1.00E+11             | 0.00E+00             |
| factorie             | 3.10E+07             | 3.00E+00             | 0.00E+00             | 9.81E+07             | 0.00E+00             | 12.04          | 1.43E+10             | 7.43E+09             | 1.16E+08             | 6.00E+10             | 0.00E+00             |
| kiama                | 6.47E+04             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 3.12           | 6.28E+07             | 9.67E+06             | 2.10E+06             | 9.10E+07             | 0.00E+00             |
| scalac               | 2.52E+06             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 15.45          | 6.36E+08             | 4.69E+07             | 6.45E+06             | 1.27E+09             | 0.00E+00             |
| scaladoc             | 1.90E+06             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.11           | 5.74E+08<br>2.05E+07 | 3.92E+07             | 7.02E+00<br>3.40E+05 | 9.76E+08             | 0.00E+00             |
| scalariform          | 1.05E+04             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 15.65          | 3.12E+08             | 5.70E+07             | 4.17E+06             | 5 78E+08             | 0.00E+00             |
| scalatest            | 7.83E+05             | 6.45E+02             | 1.93E+04             | 6.53E+04             | 3.30E+01             | 20.00          | 2.56E+08             | 2.61E+06             | 7.83E+05             | 3.51E+07             | 0.00E+00             |
| scalaxb              | 1.76E+05             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 0.00E+00             | 12.42          | 9.55E+09             | 1.22E+08             | 4.08E+06             | 1.14E+10             | 0.00E+00             |
| specs                | 1.14E+06             | 4.10E+01             | 1.38E+04             | 9.48E+04             | 5.10E+01             | 10.53          | 3.08E+08             | 1.33E+07             | 1.93E+06             | 1.12E+08             | 0.00E+00             |
| tmt                  | 1.35E+08             | 5.56E+03             | 8.70E+01             | 5.13E+04             | 5.00E+03             | 36.54          | 1.26E+19             | 3.47E+09             | 1.75E+07             | 7.19E+10             | 0.00E+00             |
|                      |                      |                      |                      | 5                    | SPECjvm20            | 08             |                      |                      |                      |                      |                      |
| compiler.compiler    | 4.50E+06             | 5.00E+00             | 1.00E+00             | 2.60E+01             | 0.00E+00             | 98.30          | 1.31E+10             | 4.17E+08             | 4.78E+07             | 1.01E+10             | 0.00E+00             |
| compiler.sunflow     | 3.31E+07             | 1.52E+02             | 1.00E+00             | 3.00E+01             | 0.00E+00             | 97.85          | 2.38E+10             | 1.02E+09             | 1.72E+08             | 2.98E+10             | 0.00E+00             |
| compress             | 6.18E+05             | 4.00E+00             | 1.00E+00             | 5.50E+01             | 0.00E+00             | 98.56          | 7.06E+10             | 2.15E+05             | 1.43E+05             | 1.56E+11             | 0.00E+00             |
| crypto.aes           | 2.94E+07             | 5.00E+00             | 1.00E+00             | 7.70E+01             | 0.00E+00             | 97.63          | 8.29E+09             | 2.80E+05             | 3.83E+05             | 2.91E+09             | 0.00E+00             |
| crypto.isa           | 2.68E±00             | 4.00E+00             | 1.00E+00             | 9.74E+02<br>7.30E±01 | 0.00E+00             | 97.55          | 4.17E+09             | 1.4/E+08             | 2.48E±07             | 2.55E±10             | 1.00E+00             |
| derby                | 4 39F+08             | 1.50E+04             | 2.60E+07             | 1.97E+06             | 1.50E+01             | 97.92          | 2.05E+10             | 1.51E+07             | 4 25E+08             | 1.43E+10             | 0.00E+00             |
| mpegaudio            | 9.38E+06             | 8.50E+01             | 3.00E+00             | 3.19E+03             | 0.00E+00             | 98.29          | 2.86E+10             | 1.50E+05             | 4.81E+06             | 1.73E+10             | 1.00E+00             |
| scimark.fft.large    | 3.36E+08             | 6.00E+00             | 1.00E+00             | 3.70E+01             | 0.00E+00             | 95.31          | 5.94E+10             | 4.01E+03             | 2.84E+03             | 3.36E+08             | 0.00E+00             |
| scimark.fft.small    | 3.36E+09             | 4.00E+00             | 1.00E+00             | 3.60E+01             | 0.00E+00             | 98.14          | 3.27E+11             | 4.90E+05             | 5.29E+05             | 3.36E+09             | 0.00E+00             |
| scimark.lu.large     | 1.34E+08             | 4.00E+00             | 1.00E+00             | 3.20E+01             | 0.00E+00             | 92.28          | 1.01E+11             | 1.80E+03             | 1.39E+03             | 1.34E+08             | 0.00E+00             |
| scimark.lu.small     | 4.02E+09             | 4.00E+00             | 1.00E+00             | 3.30E+01             | 0.00E+00             | 97.83          | 2.95E+11             | 6.12E+05             | 9.16E+05             | 4.02E+09             | 0.00E+00             |
| scimark.sor.large    | 5.03E+08             | 5.00E+00             | 1.00E+00             | 3.90E+01             | 0.00E+00             | 92.38          | 1.78E+11             | 5.34E+03             | 3.35E+03             | 5.03E+08             | 0.00E+00             |
| scimark.sor.small    | 6.00E+08             | 5.00E+00             | 1.00E+00             | 3.90E+01             | 0.00E+00             | 98.22          | 7.24E+10             | 7.00E+04             | 5.13E+04             | 6.01E+08             | 0.00E+00             |
| scimark.sparse.large | 3.36E+08             | 4.00E+00             | 1.00E+00             | 3.20E+01             | 0.00E+00             | 87.15          | 1.96E+11             | 3.64E+03             | 2.71E+03             | 3.36E+08             | 0.00E+00             |
| serial               | 2.40E+08<br>2.25E±00 | 5.00E+00<br>2.20E±02 | 1.00E+00<br>7.50E±01 | 3.20E+01<br>8.35E±02 | 0.00E+00<br>0.00E±00 | 90.42<br>08.00 | 3.10E+10<br>3.80E+10 | 2.30£+04<br>1.78F±00 | 3.33£+04<br>1.05E±00 | 2.40E+08<br>3.70E±10 | 0.00E+00<br>1.00E+00 |
| sunflow              | 1.03E+05             | 4.49E+02             | 1.00E+01             | 5.14E+02             | 0.00E+00             | 96.95          | 1.34E+10             | 2.54E+09             | 6.26E+07             | 6.23E+10             | 0.00E+00             |
| xml.transform        | 4.76E+08             | 7.00E+00             | 1.00E+00             | 2.40E+01             | 0.00E+00             | 97.82          | 5.80E+09             | 1.75E+08             | 7.74E+07             | 7.73E+09             | 0.00E+00             |
| xml.validation       | 8.99E+08             | 5.00E+00             | 1.00E+00             | 1.22E+02             | 0.00E+00             | 98.80          | 2.05E+10             | 5.49E+08             | 2.11E+08             | 2.41E+10             | 0.00E+00             |

 Table 7. Unnormalized metrics collected on all analyzed benchmarks.



Figure 8. Scatter plots of benchmark scores over the first four principal components (PCs) - Larger version.

| Benchmark         | WMC      | DIT     | СВО      | NOC     | RFC       | LCOM       |
|-------------------|----------|---------|----------|---------|-----------|------------|
|                   |          | Rena    | aissance |         |           |            |
| akka-uct          | 34607    | 4938    | 35384    | 2718    | 66665     | 674747     |
| als               | 96524    | 13465   | 95480    | 7488    | 184730    | 5049619    |
| chi-square        | 116483   | 13963   | 108544   | 8110    | 232191    | 3465588    |
| db-shootout       | 57652    | 7393    | 51499    | 4285    | 99878     | 1874929    |
| dec-tree          | 206933   | 23936   | 184901   | 14029   | 369131    | 7360650    |
| dotty             | 65887    | 7185    | 62533    | 4121    | 120656    | 1824595    |
| finagle-chirper   | 71465    | 13894   | 78437    | 6322    | 137705    | 1429201    |
| finagle-http      | 65465    | 13122   | 72146    | 5835    | 126281    | 1169093    |
| fj-kmeans         | 22425    | 3092    | 22061    | 1592    | 42584     | 461842     |
| future-genetic    | 26198    | 3499    | 25430    | 1883    | 49263     | 508615     |
| log-regression    | 163424   | 21841   | 161667   | 12057   | 307276    | 5569868    |
| movie-lens        | 101483   | 14335   | 100517   | 8050    | 192950    | 5118756    |
| naive-bayes       | 88885    | 12871   | 91563    | 7130    | 174908    | 1850846    |
| neo4j-analytics   | 119743   | 22172   | 141185   | 11666   | 224669    | 1524820    |
| page-rank         | 93537    | 13939   | 97078    | 7732    | 183346    | 1541349    |
| philosophers      | 24617    | 3432    | 24161    | 1821    | 46714     | 494658     |
| reactors          | 32644    | 4097    | 29610    | 2251    | 60899     | 1066392    |
| rx-scrabble       | 25981    | 3752    | 25829    | 1958    | 49387     | 576353     |
| scrabble          | 24333    | 3380    | 24176    | 1759    | 46212     | 484610     |
| stm-bench7        | 28074    | 3829    | 27159    | 2083    | 52889     | 635890     |
| streams-mnemonics | 21830    | 3066    | 21757    | 1571    | 41799     | 455958     |
| min               | 21830    | 3066    | 21757    | 1571    | 41799     | 455958     |
| max               | 206933   | 23936   | 184901   | 14029   | 369131    | 7360650    |
| geomean           | 55533.33 | 7842.22 | 55146.79 | 4212.47 | 105104.97 | 1358042.08 |

Table 8. CK metrics for Renaissance: Sum across all loaded classes of a benchmark.

| Benchmark            | WMC      | DIT     | СВО      | NOC     | RFC      | LCOM       |
|----------------------|----------|---------|----------|---------|----------|------------|
|                      |          | Da      | Саро     |         |          |            |
| avrora               | 13488    | 2328    | 13719    | 1145    | 25357    | 169973     |
| batik                | 31205    | 4958    | 30436    | 2655    | 60275    | 333328     |
| eclinse              | 66318    | 5753    | 45597    | 4346    | 97757    | 600483     |
| fon                  | 28162    | 4272    | 28427    | 2278    | 56725    | 324197     |
| h2                   | 20230    | 2358    | 16492    | 1200    | 37467    | 297210     |
| ivthon               | 66079    | 3595    | 34583    | 2978    | 71881    | 616258     |
| Juindex              | 14361    | 1927    | 12737    | 1004    | 26549    | 181134     |
| lusearch-fix         | 12466    | 1687    | 10640    | 857     | 22378    | 161333     |
| nmd                  | 24238    | 3406    | 21589    | 1746    | 46074    | 526595     |
| sunflow              | 18088    | 2359    | 17493    | 1266    | 37288    | 245761     |
| tomcat               | 63591    | 5347    | 37448    | 4239    | 85921    | 616524     |
| tradebeans           | 122123   | 5108    | 36856    | 6347    | 74677    | 463044     |
| tradesoan            | 124191   | 5120    | 36971    | 6421    | 75043    | 466569     |
| xalan                | 18203    | 2852    | 16987    | 1364    | 35126    | 231540     |
| min                  | 12466    | 1687    | 10640    | 857     | 22378    | 161333     |
| max                  | 124191   | 5753    | 45597    | 6421    | 97757    | 616524     |
| deomean              | 32470    | 3376.65 | 23275 3  | 2160.25 | 18/61 37 | 336101 51  |
| geomean              | 52470    | Scala   | Bench    | 2100.25 | 10101.57 | 550171.51  |
| actors               | 41398    | 4900    | 34981    | 2848    | 76794    | 1216221    |
| apparat              | 34994    | 4023    | 29838    | 2121    | 66309    | 1344756    |
| factorie             | 24693    | 2657    | 19713    | 1481    | 45364    | 1080228    |
| kiama                | 31925    | 4054    | 27664    | 2051    | 60955    | 1259018    |
| scalac               | 57337    | 7789    | 66036    | 4111    | 124240   | 2298594    |
| scaladoc             | 50655    | 6343    | 50616    | 3379    | 103649   | 2178195    |
| scalap               | 29661    | 3137    | 23638    | 1743    | 54480    | 1312560    |
| scalariform          | 32871    | 3626    | 28467    | 2041    | 62891    | 1352823    |
| scalatest            | 114544   | 4589    | 37128    | 6687    | 78209    | 1473601    |
| scalaxb              | 30112    | 3402    | 25528    | 1884    | 57511    | 1180527    |
| specs                | 150895   | 6548    | 50104    | 7890    | 105427   | 2215152    |
| tmt                  | 35875    | 3188    | 25590    | 1941    | 64142    | 1565153    |
| min                  | 24693    | 2657    | 19713    | 1481    | 45364    | 1080228    |
| max                  | 150895   | 7789    | 66036    | 7890    | 124240   | 2298594    |
| geomean              | 44505.06 | 4290.27 | 32809.23 | 2734.69 | 71839.71 | 1489515.03 |
|                      | 1        | SPECj   | vm2008   |         |          |            |
| compiler.compiler    | 36385    | 4744    | 36421    | 2728    | 77428    | 598961     |
| compiler.sunflow     | 36456    | 4745    | 36424    | 2728    | 77538    | 600724     |
| compress             | 30586    | 3843    | 30009    | 2200    | 65819    | 546396     |
| crypto.aes           | 33789    | 4134    | 31810    | 2413    | 69028    | 551637     |
| crypto.rsa           | 32724    | 4089    | 31553    | 2381    | 68457    | 549741     |
| crypto.signverify    | 30999    | 3930    | 30639    | 2250    | 66899    | 547493     |
| derby                | 55044    | 5744    | 45745    | 3480    | 103373   | 1131251    |
| mpegaudio            | 31370    | 3925    | 30483    | 2251    | 66884    | 552206     |
| scimark.fft.large    | 30572    | 3835    | 29977    | 2192    | 65767    | 546563     |
| scimark.fft.small    | 30572    | 3835    | 29977    | 2192    | 65767    | 546563     |
| scimark.lu.large     | 30569    | 3833    | 29968    | 2191    | 65761    | 546558     |
| scimark.lu.small     | 30569    | 3832    | 29967    | 2190    | 65761    | 546558     |
| scimark.monte_carlo  | 30560    | 3833    | 29966    | 2191    | 65741    | 546482     |
| scimark.sor.large    | 30565    | 3835    | 29975    | 2192    | 65755    | 546493     |
| scimark.sor.small    | 30565    | 3835    | 29975    | 2192    | 65755    | 546493     |
| scimark.sparse.large | 30561    | 3833    | 29967    | 2191    | 65746    | 546486     |
| scimark.sparse.small | 30561    | 3833    | 29967    | 2191    | 65746    | 546486     |
| serial               | 32690    | 3987    | 31200    | 2310    | 68584    | 570658     |
| sunflow              | 31946    | 4003    | 31463    | 2324    | 69076    | 552092     |
| xml.transform        | 43374    | 5654    | 40968    | 3154    | 86981    | 625179     |
| xml.validation       | 34578    | 4378    | 34225    | 2545    | 74250    | 581626     |
| min                  | 30560    | 3832    | 29966    | 2190    | 65741    | 546396     |
| max                  | 55044    | 5744    | 45745    | 3480    | 103373   | 1131251    |
| geomean              | 55194.67 | 4142.17 | 52187.12 | 2383.19 | /02/9.58 | 5/8408.18  |

Table 9. CK metrics for DaCapo, ScalaBench and SPECjvm2008: Sum across all loaded classes of a benchmark.

| Benchmark        | WMC   | DIT    | СВО   | NOC  | RFC   | LCOM   |
|------------------|-------|--------|-------|------|-------|--------|
|                  | R     | enaiss | ance  |      |       |        |
| akka-uct         | 12.83 | 1.83   | 13.11 | 1.01 | 24.71 | 250.09 |
| als              | 13.51 | 1.88   | 13.36 | 1.05 | 25.85 | 706.54 |
| chi-square       | 14.91 | 1.79   | 13.9  | 1.04 | 29.73 | 443.74 |
| db-shootout      | 14.61 | 1.87   | 13.05 | 1.09 | 25.3  | 475.03 |
| dec-tree         | 16.17 | 1.87   | 14.45 | 1.1  | 28.84 | 575.1  |
| dotty            | 18.48 | 2.02   | 17.54 | 1.16 | 33.84 | 511.81 |
| finagle-chirper  | 11.5  | 2.24   | 12.62 | 1.02 | 22.16 | 230    |
| finagle-http     | 11.41 | 2.29   | 12.57 | 1.02 | 22    | 203.71 |
| fj-kmeans        | 13.76 | 1.9    | 13.53 | 0.98 | 26.13 | 283.34 |
| future-genetic   | 13.94 | 1.86   | 13.53 | 1    | 26.22 | 270.68 |
| log-regression   | 14.39 | 1.92   | 14.23 | 1.06 | 27.05 | 490.31 |
| movie-lens       | 13.31 | 1.88   | 13.18 | 1.06 | 25.31 | 671.4  |
| naive-bayes      | 13.01 | 1.88   | 13.4  | 1.04 | 25.6  | 270.91 |
| neo4j-analytics  | 11.07 | 2.05   | 13.06 | 1.08 | 20.78 | 141.02 |
| page-rank        | 12.68 | 1.89   | 13.16 | 1.05 | 24.86 | 209    |
| philosophers     | 13.29 | 1.85   | 13.05 | 0.98 | 25.22 | 267.09 |
| reactors         | 14.54 | 1.82   | 13.19 | 1    | 27.13 | 475.01 |
| rx-scrabble      | 13.2  | 1.91   | 13.12 | 0.99 | 25.08 | 292.71 |
| scrabble         | 13.57 | 1.89   | 13.48 | 0.98 | 25.77 | 270.28 |
| stm-bench7       | 13.36 | 1.82   | 12.93 | 0.99 | 25.17 | 302.66 |
| stream-mnemonics | 13.54 | 1.9    | 13.5  | 0.97 | 25.93 | 282.85 |
| min              | 11.07 | 1.79   | 12.57 | 0.97 | 20.78 | 141.02 |
| max              | 18.48 | 2.29   | 17.54 | 1.16 | 33.84 | 706.54 |
| geomean          | 13.58 | 1.92   | 13.49 | 1.03 | 25.71 | 332.19 |

Table 10. CK metrics for Renaissance: Average across all loaded classes of a benchmark.

| Benchmark            | WMC   | DIT     | СВО   | NOC  | RFC   | LCOM   |
|----------------------|-------|---------|-------|------|-------|--------|
|                      |       | DaCap   | 0     |      |       |        |
| avrora               | 11.74 | 2.03    | 11.94 | 1    | 22.07 | 147.93 |
| batik                | 12.2  | 1.94    | 11.9  | 1.04 | 23.57 | 130.36 |
| eclipse              | 21.02 | 1.82    | 14.45 | 1.38 | 30.98 | 190.33 |
| fop                  | 12.67 | 1.92    | 12.79 | 1.02 | 25.52 | 145.84 |
| h2                   | 17.42 | 2.03    | 14.2  | 1.03 | 32.27 | 255.99 |
| jython               | 22.81 | 1.24    | 11.94 | 1.03 | 24.81 | 212.72 |
| luindex              | 14.68 | 1.97    | 13.02 | 1.03 | 27.15 | 185.21 |
| lusearch-fix         | 15.13 | 2.05    | 12.91 | 1.04 | 27.16 | 195.79 |
| pmd                  | 14.31 | 2.01    | 12.74 | 1.03 | 27.2  | 310.86 |
| sunflow              | 13.83 | 1.8     | 13.37 | 0.97 | 28.51 | 187.89 |
| tomcat               | 22.76 | 1.91    | 13.4  | 1.52 | 30.75 | 220.66 |
| tradebeans           | 41.89 | 1.75    | 12.64 | 2.18 | 25.62 | 158.85 |
| tradesoap            | 42.43 | 1.75    | 12.63 | 2.19 | 25.64 | 159.4  |
| xalan                | 13.57 | 2.13    | 12.67 | 1.02 | 26.19 | 172.66 |
| min                  | 11.74 | 1.24    | 11.9  | 0.97 | 22.07 | 130.36 |
| max                  | 42.43 | 2.13    | 14.45 | 2.19 | 32.27 | 310.86 |
| geomean              | 17.97 | 1.87    | 12.88 | 1.2  | 26.82 | 186.05 |
|                      | S     | calaBeı | ıch   |      |       |        |
| actors               | 15    | 1.78    | 12.67 | 1.03 | 27.82 | 440.66 |
| apparat              | 16.78 | 1.93    | 14.31 | 1.02 | 31.8  | 644.97 |
| factorie             | 16.67 | 1.79    | 13.31 | 1    | 30.63 | 729.39 |
| kiama                | 15.55 | 1.97    | 13.47 | 1    | 29.69 | 613.26 |
| scalac               | 14.04 | 1.91    | 16.17 | 1.01 | 30.43 | 562.97 |
| scaladoc             | 15.12 | 1.89    | 15.1  | 1.01 | 30.93 | 650.01 |
| scalap               | 17.14 | 1.81    | 13.66 | 1.01 | 31.47 | 758.27 |
| scalariform          | 16.18 | 1.79    | 14.02 | 1    | 30.97 | 666.09 |
| scalatest            | 42.6  | 1.71    | 13.81 | 2.49 | 29.08 | 548.01 |
| scalaxb              | 16.04 | 1.81    | 13.6  | 1    | 30.64 | 628.94 |
| specs                | 40.1  | 1.74    | 13.31 | 2.1  | 28.02 | 588.67 |
| tmt                  | 19.16 | 1.7     | 13.67 | 1.04 | 34.26 | 836.09 |
| min                  | 14.04 | 1.7     | 12.67 | 1    | 27.82 | 440.66 |
| max                  | 42.6  | 1.97    | 16.17 | 2.49 | 34.26 | 836.09 |
| geomean              | 18.85 | 1.82    | 13.9  | 1.16 | 30.43 | 631.02 |
|                      | SP    | ECjvm   | 2008  |      |       |        |
| compiler.compiler    | 13.55 | 1.77    | 13.56 | 1.02 | 28.83 | 222.99 |
| compiler.sunflow     | 13.57 | 1.77    | 13.56 | 1.02 | 28.87 | 223.65 |
| compress             | 13.83 | 1.74    | 13.57 | 1    | 29.77 | 247.13 |
| crypto.aes           | 14.32 | 1.75    | 13.48 | 1.02 | 29.26 | 233.84 |
| crypto.rsa           | 14    | 1.75    | 13.5  | 1.02 | 29.29 | 235.23 |
| crypto.signverify    | 13.73 | 1.74    | 13.57 | 1    | 29.63 | 242.47 |
| derby                | 16.9  | 1.76    | 14.04 | 1.07 | 31.73 | 347.22 |
| mpegaudio            | 13.89 | 1.74    | 13.49 | 1    | 29.61 | 244.45 |
| scimark.fft.large    | 13.87 | 1.74    | 13.6  | 0.99 | 29.84 | 247.99 |
| scimark.fft.small    | 13.87 | 1.74    | 13.6  | 0.99 | 29.84 | 247.99 |
| scimark.lu.large     | 13.88 | 1.74    | 13.6  | 0.99 | 29.85 | 248.1  |
| scimark.lu.small     | 13.88 | 1.74    | 13.61 | 0.99 | 29.86 | 248.21 |
| scimark.monte_carlo  | 13.87 | 1.74    | 13.6  | 0.99 | 29.84 | 248.06 |
| scimark.sor.large    | 13.87 | 1.74    | 13.6  | 0.99 | 29.83 | 247.96 |
| scimark.sor.small    | 13.87 | 1.74    | 13.6  | 0.99 | 29.83 | 247.96 |
| scimark.sparse.large | 13.87 | 1.74    | 13.6  | 0.99 | 29.84 | 248.06 |
| scimark.sparse.small | 13.87 | 1.74    | 13.6  | 0.99 | 29.84 | 248.06 |
| serial               | 14.15 | 1.73    | 13.5  | 1    | 29.68 | 246.93 |
| sunflow              | 13.69 | 1.72    | 13.49 | 1    | 29.61 | 236.64 |
| xml.transform        | 14.31 | 1.86    | 13.51 | 1.04 | 28.69 | 206.19 |
| xml.validation       | 13.62 | 1.72    | 13.48 | 1    | 29.24 | 229.08 |
| min                  | 13.55 | 1.72    | 13.48 | 0.99 | 28.69 | 206.19 |
| max                  | 16.9  | 1.86    | 14.04 | 1.07 | 31.73 | 347.22 |
| geomean              | 14    | 1.75    | 13.58 | 1.01 | 29.65 | 244.03 |

 Table 11. CK metrics for DaCapo, ScalaBench and SPECjvm2008: Average across all loaded classes of a benchmark.

| workload          | A    | С   | D    | S   | EAV  | EAWA |      | A   | LV   | 7   | LLC  |      | MHS  |     |
|-------------------|------|-----|------|-----|------|------|------|-----|------|-----|------|------|------|-----|
| akka-uct          | +1%  | 38% | +2%  | 22% | +5%  | 0%   | +1%  | 44% | +4%  | 1%  | +1%  | 39%  | +3%  | 4%  |
| als               | +0%  | 81% | +1%  | 33% | -1%  | 14%  | +11% | 0%  | +10% | 0%  | +1%  | 29%  | +0%  | 82% |
| chi-square        | +4%  | 3%  | +4%  | 4%  | +5%  | 0%   | +5%  | 0%  | +3%  | 12% | +2%  | 33%  | +4%  | 2%  |
| db-shootout       | -0%  | 35% | -0%  | 63% | +0%  | 14%  | +5%  | 0%  | +0%  | 24% | -0%  | 29%  | +0%  | 48% |
| dec-tree          | +0%  | 60% | +1%  | 0%  | -0%  | 90%  | +8%  | 0%  | +3%  | 0%  | -0%  | 35%  | -0%  | 45% |
| dotty             | +0%  | 1%  | +2%  | 0%  | +0%  | 85%  | +3%  | 0%  | +1%  | 0%  | +0%  | 0%   | +8%  | 0%  |
| finagle-chirper   | -1%  | 90% | -0%  | 96% | +24% | 0%   | -1%  | 88% | +0%  | 91% | +3%  | 23%  | +4%  | 18% |
| finagle-http      | -1%  | 5%  | +4%  | 0%  | -1%  | 12%  | +0%  | 60% | -0%  | 25% | -0%  | 29%  | -0%  | 95% |
| fj-kmeans         | -0%  | 16% | -1%  | 0%  | +0%  | 90%  | +2%  | 0%  | -0%  | 6%  | +71% | 0%   | -0%  | 62% |
| future-genetic    | +24% | 0%  | +0%  | 59% | +2%  | 1%   | +2%  | 0%  | +1%  | 0%  | +1%  | 0%   | +25% | 0%  |
| log-regression    | -0%  | 89% | +1%  | 58% | +0%  | 73%  | +15% | 0%  | +2%  | 6%  | +2%  | 1%   | +1%  | 28% |
| movie-lens        | +1%  | 85% | +0%  | 99% | +1%  | 81%  | +1%  | 84% | -1%  | 80% | -3%  | 18%  | +1%  | 76% |
| naive-bayes       | +1%  | 14% | -3%  | 0%  | +1%  | 25%  | +13% | 2%  | +1%  | 17% | +1%  | 27%  | -0%  | 55% |
| neo4j-analytics   | +0%  | 91% | -4%  | 37% | -7%  | 10%  | +5%  | 24% | -3%  | 49% | -0%  | 100% | -4%  | 27% |
| page-rank         | -1%  | 2%  | -0%  | 51% | -1%  | 2%   | +2%  | 0%  | -0%  | 38% | -1%  | 0%   | -1%  | 0%  |
| philosophers      | -5%  | 5%  | -2%  | 32% | -1%  | 43%  | +2%  | 9%  | +2%  | 22% | -1%  | 64%  | -1%  | 62% |
| reactors          | -0%  | 42% | -2%  | 0%  | -0%  | 11%  | -1%  | 3%  | -1%  | 1%  | -1%  | 4%   | -1%  | 16% |
| rx-scrabble       | -0%  | 93% | +1%  | 6%  | -0%  | 69%  | -1%  | 0%  | -1%  | 8%  | -0%  | 38%  | +1%  | 0%  |
| scrabble          | +1%  | 65% | +1%  | 32% | -2%  | 11%  | +3%  | 6%  | -1%  | 47% | -1%  | 31%  | +22% | 0%  |
| stm-bench7        | +1%  | 21% | +3%  | 0%  | +1%  | 26%  | +1%  | 6%  | +0%  | 12% | +1%  | 1%   | -0%  | 96% |
| streams-mnemonics | +0%  | 35% | +22% | 0%  | +1%  | 2%   | +1%  | 3%  | +2%  | 0%  | +0%  | 59%  | +7%  | 0%  |

 Table 12. Optimization impact – Renaissance benchmarks.

| workload   | A   | С   | D   | S   | EA  | EAWA |     | М   | LV  |     | LLC |     | MHS |     |
|------------|-----|-----|-----|-----|-----|------|-----|-----|-----|-----|-----|-----|-----|-----|
| avrora     | +0% | 3%  | +0% | 4%  | +0% | 90%  | +0% | 17% | +0% | 19% | +0% | 0%  | +0% | 7%  |
| batik      | -0% | 35% | -0% | 0%  | +0% | 91%  | +1% | 0%  | +0% | 81% | -0% | 68% | -0% | 3%  |
| eclipse    | +0% | 26% | +5% | 0%  | -0% | 39%  | +1% | 0%  | +1% | 0%  | +0% | 10% | +0% | 11% |
| fop        | +0% | 90% | +1% | 0%  | +0% | 63%  | +0% | 83% | +1% | 0%  | -0% | 84% | +0% | 72% |
| h2         | +0% | 29% | +2% | 0%  | -0% | 65%  | +1% | 4%  | +0% | 20% | +0% | 8%  | +1% | 1%  |
| jython     | -1% | 26% | +5% | 0%  | +1% | 65%  | +2% | 9%  | -0% | 96% | +1% | 42% | +0% | 96% |
| luindex    | -0% | 39% | +3% | 0%  | -0% | 5%   | +2% | 0%  | +0% | 0%  | -0% | 70% | -1% | 0%  |
| lusearch   | -0% | 17% | +1% | 0%  | -0% | 9%   | -0% | 76% | -0% | 1%  | -0% | 80% | -0% | 62% |
| pmd        | -0% | 10% | -0% | 58% | +0% | 28%  | -1% | 0%  | -0% | 63% | +0% | 76% | +0% | 81% |
| sunflow    | +1% | 1%  | +4% | 0%  | +0% | 78%  | +0% | 24% | +2% | 2%  | +2% | 1%  | +2% | 1%  |
| tomcat     | +0% | 6%  | -0% | 40% | +0% | 76%  | -0% | 54% | +0% | 40% | -0% | 85% | -0% | 91% |
| tradebeans | +0% | 19% | +7% | 0%  | +0% | 46%  | -0% | 33% | +1% | 0%  | +0% | 78% | +0% | 85% |
| tradesoap  | +3% | 1%  | -2% | 0%  | -2% | 4%   | -0% | 80% | +1% | 35% | +0% | 70% | -3% | 0%  |
| xalan      | +1% | 4%  | +1% | 0%  | +0% | 52%  | +0% | 1%  | +0% | 6%  | +0% | 42% | +0% | 2%  |

Table 13. Optimization impact – DaCapo benchmarks.

| workload    | A   | С   | D   | S   | EA  | EAWA |      | Λ   | LV  |     | LLC |     | MHS |      |
|-------------|-----|-----|-----|-----|-----|------|------|-----|-----|-----|-----|-----|-----|------|
| actors      | +0% | 51% | +1% | 0%  | +1% | 0%   | +0%  | 35% | +0% | 5%  | -0% | 86% | +0% | 4%   |
| apparat     | +1% | 2%  | -1% | 14% | -1% | 19%  | +0%  | 83% | +1% | 10% | -0% | 62% | -0% | 74%  |
| factorie    | +2% | 0%  | +7% | 0%  | +1% | 1%   | -2%  | 0%  | +1% | 30% | +1% | 7%  | +1% | 27%  |
| kiama       | -0% | 37% | +4% | 0%  | -0% | 24%  | +1%  | 0%  | +1% | 0%  | +0% | 24% | +0% | 60%  |
| scalac      | -0% | 77% | +1% | 0%  | +0% | 38%  | -0%  | 96% | +0% | 20% | -0% | 32% | -0% | 10%  |
| scaladoc    | -2% | 0%  | +0% | 65% | -3% | 0%   | -2%  | 0%  | -1% | 23% | -1% | 10% | -1% | 40%  |
| scalap      | -0% | 1%  | +1% | 0%  | -0% | 0%   | +9%  | 0%  | +2% | 0%  | -0% | 7%  | -0% | 0%   |
| scalariform | +0% | 5%  | +1% | 0%  | -0% | 49%  | +0%  | 1%  | +0% | 2%  | +0% | 64% | -0% | 22%  |
| scalatest   | +0% | 90% | -1% | 19% | -1% | 34%  | +0%  | 83% | +1% | 2%  | +1% | 41% | +0% | 100% |
| scalaxb     | +1% | 88% | +8% | 5%  | +1% | 87%  | +6%  | 15% | +6% | 18% | +7% | 8%  | +2% | 72%  |
| specs       | -0% | 16% | +0% | 18% | -0% | 11%  | +0%  | 2%  | +0% | 78% | -0% | 20% | -0% | 45%  |
| tmt         | +0% | 10% | +1% | 0%  | +0% | 0%   | +13% | 0%  | +1% | 0%  | +0% | 6%  | +0% | 42%  |

 Table 14. Optimization impact – ScalaBench benchmarks.

| workload             | A   | С   | D   | S   | EA  | EAWA GM |       | L 1 | LV   |     | LLC |      | MHS |     |
|----------------------|-----|-----|-----|-----|-----|---------|-------|-----|------|-----|-----|------|-----|-----|
| compiler.compiler    | +0% | 8%  | +1% | 0%  | -0% | 10%     | +3%   | 0%  | +1%  | 0%  | -0% | 35%  | -0% | 73% |
| compiler.sunflow     | -0% | 42% | +1% | 0%  | +0% | 39%     | +2%   | 0%  | +1%  | 0%  | -0% | 39%  | +0% | 12% |
| compress             | -0% | 33% | -2% | 0%  | +0% | 77%     | +2%   | 0%  | +4%  | 0%  | -0% | 82%  | -0% | 35% |
| crypto.aes           | -0% | 5%  | -0% | 67% | -0% | 37%     | +1%   | 0%  | +1%  | 0%  | -0% | 8%   | -0% | 4%  |
| crypto.rsa           | -0% | 20% | +0% | 2%  | -0% | 36%     | +0%   | 77% | -0%  | 34% | -0% | 29%  | -0% | 15% |
| crypto.signverify    | -0% | 92% | +0% | 69% | -0% | 64%     | +9%   | 0%  | -0%  | 74% | -0% | 100% | +0% | 87% |
| derby                | +0% | 44% | +0% | 59% | -0% | 48%     | -1%   | 1%  | -1%  | 18% | +0% | 58%  | +0% | 72% |
| mpegaudio            | -0% | 84% | -3% | 0%  | +0% | 26%     | +5%   | 0%  | +0%  | 69% | +0% | 50%  | +0% | 31% |
| scimark.fft.large    | -3% | 1%  | -2% | 3%  | -3% | 2%      | -1%   | 19% | -3%  | 0%  | -2% | 7%   | -1% | 49% |
| scimark.fft.small    | -1% | 44% | +2% | 33% | -3% | 9%      | -1%   | 65% | -2%  | 22% | -3% | 4%   | -1% | 68% |
| scimark.lu.large     | -0% | 11% | -0% | 57% | -0% | 8%      | +69%  | 0%  | +29% | 0%  | -0% | 6%   | +0% | 81% |
| scimark.lu.small     | +0% | 40% | +1% | 0%  | +0% | 16%     | +137% | 0%  | +58% | 0%  | +0% | 92%  | +0% | 1%  |
| scimark.monte_carlo  | +2% | 30% | +7% | 0%  | -0% | 83%     | -0%   | 83% | +0%  | 89% | +1% | 61%  | +1% | 62% |
| scimark.sor.large    | +0% | 4%  | -0% | 21% | +0% | 0%      | +34%  | 0%  | -0%  | 25% | +0% | 13%  | -0% | 44% |
| scimark.sor.small    | -0% | 64% | -0% | 44% | +0% | 65%     | +36%  | 0%  | +0%  | 20% | -0% | 32%  | +0% | 38% |
| scimark.sparse.large | +0% | 4%  | +1% | 0%  | +0% | 4%      | +16%  | 0%  | +0%  | 2%  | +0% | 46%  | +0% | 16% |
| scimark.sparse.small | -0% | 2%  | -0% | 0%  | -0% | 6%      | -10%  | 0%  | -0%  | 0%  | +0% | 1%   | -0% | 6%  |
| serial               | +0% | 94% | +2% | 0%  | +1% | 4%      | +4%   | 0%  | +1%  | 5%  | -1% | 11%  | +0% | 39% |
| sunflow              | +1% | 32% | +2% | 1%  | +1% | 19%     | +1%   | 17% | +2%  | 1%  | +1% | 29%  | +1% | 16% |
| xml.transform        | +0% | 73% | +2% | 0%  | -0% | 60%     | +3%   | 0%  | +0%  | 24% | +0% | 83%  | +0% | 54% |
| xml.validation       | -1% | 0%  | +1% | 6%  | -1% | 10%     | -1%   | 13% | -1%  | 1%  | -1% | 2%   | -1% | 5%  |

 Table 15. Optimization impact – SPECjvm2008 benchmarks.

| optimization                           | compilation time change |
|--|-------------------------|
| Atomic-Operation Coalescing            | 0.6%                    |
| Dominance-Based Duplication Simulation | 19.6%                   |
| Loop-Wide Lock Coarsening              | 6.7%                    |
| Method-Handle Simplification           | 7.2%                    |
| Speculative Guard Motion               | 5.8%                    |
| Loop Vectorization                     | 5.1%                    |
| Escape Analysis with Atomic Operations | 6.9%                    |

**Table 16.** Compilation time associated with individual optimizations.