

Automated Detection of Software Performance Antipatterns in Java-Based Applications

Catia Trubiani¹, Riccardo Pincirolì¹, Andrea Biaggi¹, and Francesca Arcelli Fontana¹

Abstract—The detection of performance issues in Java-based applications is not trivial since many factors concur to poor performance, and software engineers are not sufficiently supported for this task. The goal of this manuscript is the automated detection of performance problems in running systems to guarantee that no quality-based hinders prevent their successful usage. Starting from software performance antipatterns, i.e., bad practices (e.g., extensive interaction between software methods) expressing both the problem and the solution with the purpose of identifying shortcomings and promptly fixing them, we develop a framework that automatically detects seven software antipatterns capturing a variety of performance issues in Java-based applications. Our approach is applied to real-world case studies from different domains, and it captures four real-life performance issues of Hadoop and Cassandra that were not predicted by state-of-the-art approaches. As empirical evidence, we calculate the accuracy of the proposed detection rules, we show that code commits inducing and fixing real-life performance issues present interesting variations in the number of detected antipattern instances, and solving one of the detected antipatterns improves the system performance up to 50%.

Index Terms—Dynamic analysis, java-based applications, software performance antipatterns

1 INTRODUCTION

THE performance evaluation of Java-based applications is challenging due to many variabilities, such as software failures and workload fluctuation in requests [1], [2], [3], that may occur when the system is running and inevitably contribute to affect the overall service quality [4], [5], [6]. Understanding if an application can always meet the desired performance (e.g., the system response time must be shorter than 5 seconds or resource utilization must be lower than 80%) is of key relevance since it impacts the perception of end-users and their satisfaction while interacting with the system [7], [8].

In the literature, several approaches have been proposed for modeling, analyzing, and optimizing the performance of software applications [9], [10]. Two main directions have been pursued: (i) model-based performance analysis, i.e., performance models are built out of Java applications [11], [12] and used for predictions; (ii) application performance monitoring, i.e., tools that collect trace data for inspection [13], [14]. Motivated by the recent trend of integrating development (Dev) and operations (Ops) teams, processes,

and tools [15], [16], [17], it is necessary that software engineers are aware of the performance evolution of their applications. If performance issues are detected, then engineers must also be able to promptly fix such issues. To this end, several approaches emerged, e.g., automated performance tests [18] to guarantee the prompt identification and fixing of performance degradation, or performance load testing [19] to evaluate software refactorings that most likely lead to performance improvement. However, most of the approaches in the literature, e.g., [2], [20], [21], [22], act statically on the implementation code. A recent study [23] pointed out that static code analysis may fail in capturing complex root causes of real-life performance issues, e.g., the interactions between procedures that occur when executing the source code only. Hence, in this paper we aim to exploit dynamic information, which is fundamental for the detection of some performance issues, at the cost of deploying a testing environment for profiling applications and introducing runtime efforts.

We focus our attention on identifying performance issues in Java-based applications, i.e., the target systems under analysis are already in production and subject to multiple variations, such as changes in the execution environment. As a motivating example, let us consider a real-life case study [24] where a performance overhead of 17% is experienced for the continuous integration of a software release. The diagnosis of performance problems is indeed non-trivial, a study in [25] indicates the Apache project, i.e., relevant to our research since Java is used, with the longest average diagnosis time (194 days). Our goal is to improve the system performance by identifying the bad practices of software components and fixing them before the system becomes unusable. Let us consider as an example of bad practice a software component that monopolizes the processing (namely the *Blob* [26]). This implies a single and complex *controller* component that orchestrates the computation by extensively interacting with other components. As a consequence, the system response

- Catia Trubiani and Riccardo Pincirolì are with Gran Sasso Science Institute, 67100 L'Aquila, Italy. E-mail: {catia.trubiani, riccardo.pincirolì}@gssi.it.
- Andrea Biaggi and Francesca Arcelli Fontana are with the University of Milano-Bicocca, 20126 Milano, Italy. E-mail: a.biaggi1@campus.unimib.it, francesca.arcelli@unimib.it.

Manuscript received 6 May 2022; revised 14 December 2022; accepted 23 December 2022. Date of publication 10 January 2023; date of current version 18 April 2023.

This work was supported by MUR PRIN project under Grant SEDUCE 2017TWRCNB.

(Corresponding author: Catia Trubiani.)

Recommended for acceptance by N. Meng.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TSE.2023.3234321>, provided by the authors.

Digital Object Identifier no. 10.1109/TSE.2023.3234321

time can suffer by such a behaviour. To fix this problem, it would be beneficial to involve other software components and delegate them part of the overall computation. This way, the system response time can improve since it benefits from some processing running in parallel. To achieve this objective of identifying bad practices, we use software performance antipatterns [26], [27] since they include the description of both (i) the problems leading to performance flaws, and (ii) the best practices aimed to get performance improvements.

In the context of Java-based applications, we focus on the following seven software performance antipatterns: 1) Circuitous Treasure Hunt (CTH), 2) Extensive Processing (EP), 3) Wrong Cache Strategy (WCS), 4) Blob, 5) Tower of Babel (ToB), 6) Empty Semi Trucks (EST), and 7) Excessive Dynamic Allocation (EDA). Further details on these antipatterns and the motivation on the selection of these antipatterns are provided in Section 3.3. We develop a framework to automatically detect these seven performance antipatterns and we evaluate it on a variegated set of real-world applications. The conducted experimentation advocates the following main findings: (i) our framework is efficient, the detection of antipatterns is performed, on average, in less than a minute; (ii) our framework is accurate, the F1 score is larger than 85% in the considered cases; (iii) our framework advances state-of-the-art methodologies, it detects complex performance problems not recognized by other tools; (iv) antipattern-based refactoring can lead to system performance improvement up to 50%. The main contributions of our work are as follows:

- the specification of seven software performance antipatterns that are customized to verify a set of properties for Java-based applications;
- the development of JPAD, *Java Performance Antipattern Detector*, a framework that automatically detects the seven software performance antipatterns;
- the evaluation of JPAD efficiency and accuracy on five real-world applications with different complexity and representative of multiple domains;
- the comparison of JPAD with state-of-the-art approaches on the detection of real-life performance issues in nine code commits of two further systems;
- empirical evidence on the benefit of solving performance antipatterns.

In summary, our approach advocates the usage of software performance antipatterns as valuable support to automatically detect performance issues of Java-based applications. The benefit is that software engineers are promptly informed of software components showing specific bad practices and candidate of being refactored.

The rest of the manuscript is organized as follows. Section 2 reviews the related work. Section 3 describes our approach, and we discuss the key properties of software performance antipatterns, thus to motivate the choice of implementing some of them. Details on detection algorithms for the seven implemented software performance antipatterns are provided in Appendix A, available online, the rationale for not implementing some of them is explained in Appendix B, available in the online supplemental material. Research questions,

analyzed software systems, and the experimental evaluation are presented in Section 4. Threats to validity are argued in Section 5. A discussion on limitations of the approach is reported in Section 6. Concluding remarks and possible directions for future research are outlined in Section 7. Replication data are publicly available [28].

2 RELATED WORK

Our work mainly relates to three streams of research, i.e., architectural antipatterns, code smells, and Java-specific performance issues that are briefly reviewed in the following. This manuscript moves a step forward in the attempt of establishing synergies between architectural antipatterns and code smells for the performance evaluation of Java applications.

Architectural Antipatterns. In the broader context of (anti)patterns and quality attributes (e.g., reliability, security), there are several works that aim to match their connections, e.g., [29], [30], [31], [32]. When focusing on performance-related concerns there is much less work. Software performance antipatterns are studied first by Smith [26], [27] who provides the preliminary definitions based on her experience. The specification is expressed in natural language and is technology-independent, meaning that antipatterns can be customized in many different contexts. Other researchers redefine these natural language definitions using first-order logical predicates later applied to architectural design models [33] and recently adapted to further architectural formalism such as probabilistic model checking [34]. A first attempt of adopting architectural antipatterns in running systems is provided in [35], where problem root causes are isolated and a graph of dependencies is built to match problems with the specification of antipatterns.

Static and Dynamic Approaches for Code Smells. In the literature, extensive work is devoted to investigate code smells [36], [37], [38], and several investigations are performed, e.g., the analysis of (i) inter-smell interactions to understand their effects [39] and (ii) sequences of different kinds of bad smells to improve detection and solution [40]. Static analysis techniques are adopted to locate bugs in software, e.g., performance bugs that waste processing time due to superfluous loop iterations are detected in [21], and tools like FindBugs [41] can find potential root causes for performance antipatterns. Object-Relational Mapping (i.e., non-trivial database access) is exploited in [20] where static analysis can detect a huge number of performance antipattern instances. Li et al. [22] keep using static analysis, but focus on problematic duplicate logging code smells to emphasize that logging code is highly associated with both the structure and the functionality of the surrounding code. An approach to identify code changes that may potentially be responsible for performance regressions is proposed in [42], but it does not analyze root causes behind such regressions. An exploratory study on performance regressions is presented in [43] where six code level root-causes are identified, but they mainly refer to inner changes, e.g., function calls or parameter values. An attempt of using static and dynamic metrics is proposed in [44] where evolutionary algorithms are adopted to detect performance

regressions, however causes are not treated. Static and dynamic data is exploited also in [45] where a selection of benchmarks limits runtime overhead at the cost of missing the prediction of some performance issues. Preventing performance issues before the commit of code changes is pursued in [23] where random forest classifiers are trained on large datasets of performance regressions.

More recently, an experience report on locating the root causes of performance regressions is presented in [46] where web-access logs are exploited to tune the expected workload. However, this report targets only web-based systems, and machine learning techniques may lack to capture the relationship between problematic runtime activities and their impact on system performance. Summarizing, the main difference between static and dynamic approaches is that the former can detect some performance issues during the development process and cost fewer efforts, whereas the latter can benefit from runtime information probably capturing a larger set of issues (as confirmed in [23]), with the drawback of introducing monitoring efforts.

Java-Specific Performance Issues. Our work mainly relates to Java-specific performance antipattern definitions, and there are several approaches that identify different bad practices in Java platforms and APIs [47], [48], [49], [50]. The difficulty of evaluating the performance of Java applications is acknowledged also in the testing domain; Java open source projects are usually subject to a limited number of performance tests that are rarely updated and typically maintained by a small group of developers [51]. In [52] a rule-based approach is proposed to detect performance antipatterns from runtime traces while targeting Java EE antipatterns. In [53] load testing and profiling data are exploited to detect bad practices in Java applications. This approach leverages performance experts to identify problematic snapshot(s), and the antipattern detection is performed (similarly to approaches dealing with performance regressions [23], [45]) comparing snapshots with the recognized problematic one(s).

To the best of our knowledge, there are few attempts in the direction of bridging performance and code-related issues with performance antipatterns. Our approach relies on dynamic analysis and adopts software performance antipatterns to identify bad practices arising when the system is in execution. The main difference w.r.t. work more closely related to ours [35], [53] is that we consider a plethora of seven software performance antipatterns applied to a broad set of real-world case studies. Experimental results are promising to foster further research.

3 OUR APPROACH

This section introduces the approach proposed to detect software performance antipatterns in Java applications. The design of the approach is driven by the following key insight. Each performance antipattern describes a bad design practice that can be (partially) observed and checked by a combination of a particular set of system key properties including design characteristics (e.g., large number of calls) and performance metrics (e.g., long execution time). Our detection approach relies on monitoring these system key properties. A system component is recognized to be an antipattern when its design characteristics and/or performance metrics deviate most from the average

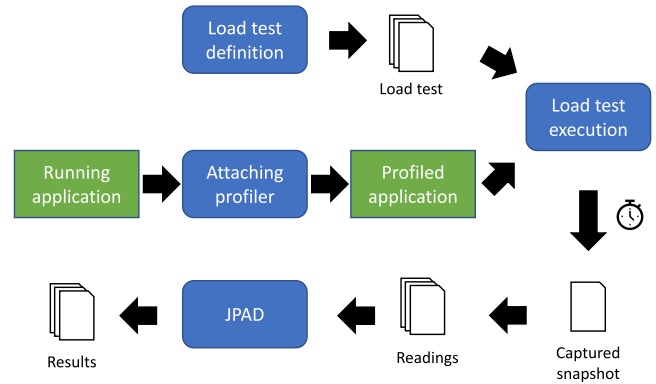


Fig. 1. High-level workflow of our approach.

values, as calculated considering all other components belonging to the same system.

Fig. 1 depicts the workflow of our approach to automatically detect software performance antipatterns in Java applications. First, a load test suite is defined for each system, and it is used to simulate the interaction of users with the system under analysis. Then, the system under analysis is launched and the profiler is attached. We are aware that runtime performance monitoring of Java applications is expensive and the profiling process can generate overhead, however this is also assessed as a necessary task to collect the performance characteristics of interest [54]. After attaching the profiler, test suites are executed. During their execution, we periodically capture snapshots that contain the readings related to CPU, threads, and memory, later exported in CSV or XML files, according to the represented data. These files are provided as input to JPAD, i.e., the tool we developed for the automatic detection of software performance antipatterns.

In the sequel of the section, we detail the operational tasks of the proposed framework, i.e., how the load testing is performed, the technology used, and the criteria adopted for defining the test suite. Then, we present the application profiler used to monitor the systems under analysis during the execution of the load tests. We also describe the seven software performance antipatterns that we implement into JPAD to enable their automatic detection. Starting from the natural language definition of antipatterns, the detection rules are expressed in terms of the data acquired through the profiler. JPAD takes as input the readings from the profiler and makes use of such data to automatically perform the detection of software performance antipatterns.

JPAD is a JavaFX application that allows the user to load the profiling data exported from YourKit (i.e., the input files) and contains an embedded console where the results are reported. The tool works with thresholds and offsets that establish when performance issues arise. A threshold value t represents a boundary on a specific performance metric value v , and JPAD verifies if $v > t$. An offset value o is instead a percentage addendum that allows to deviate from systems' properties extracted as average values, i.e., a property on a specific system component (p) is compared with the average across all components (\bar{P}), and JPAD verifies if $p > \bar{P} + \bar{P} \cdot o\%$. Note that both thresholds and offsets are used to evaluate a specific test input. The goal is to detect those system characteristics that deviate, under the same input assumptions, with a certain margin from

requirements and/or average values. For instance, 10% can be set as a threshold for the CPU usage, and JPAD considers suspicious all CPU resources showing larger values. As a further example, let us consider 5% as the offset on method number of calls. JPAD calculates the average value (considering the total number of calls across all the system methods) and augments such average by 5%. This way, JPAD detects as suspicious all the methods whose number of calls deviates by 5% from the average within a specific test input. Note that these threshold and offset settings can be modified by users in case further knowledge is available and other values are considered more appropriate. Sensitivity analysis on threshold values is discussed in Appendix C, available in the online supplemental material.¹

3.1 Load Testing

To define the required tests for each system under analysis, we explore the target application to understand key functionalities in execution. We are aware that designing load tests represents a threat to the internal validity of the approach. Indeed, we may miss some functional test cases, and there might be functionalities that are not monitored. Hence, there is a risk of not detecting some antipatterns and consequently not capturing performance issues. This threat is smoothed by delegating the design of load tests to software engineers that can decide which system functionalities require performance testing. To support software engineers in this task, as a rule of thumb, we foresee the selection of the software component(s) showing the highest CPU utilization as candidate to load testing. To support this guideline, in our experimentation (see Section 4) we show that such a selection leads to detect real-life performance issues.

3.2 Application Profiling

To profile the applications under analysis during their execution, we use the YourKit Java Profiler [55]. This decision is motivated by several reasons. First, YourKit is widely used for the performance evaluation of real-world applications both in industry [56] and in academia [53], [57]. Second, precise instructions about reducing or avoiding performance overhead (generated from the profiler) are provided [58]. Third, YourKit is a powerful application profiler, it monitors several aspects of the profiled system, and provides a view for multiple system features. Besides, the output of YourKit (i.e., the readings, see Fig. 1) can be easily exported in different data formats. For the purpose of this work, we monitor usage statistics: CPU usage, memory usage, CPU hotspots, call tree, blocked threads, and garbage collected objects. CPU and memory usage measure the percentage of used resources during the execution of the application. CPU hotspots are those methods that spend the longest time on the CPU. Call tree is divided in two different views: (i) a merged call tree that shows a top-down call tree of all application threads merged together into a single tree, and (ii) a call tree by thread that shows an individual top-down call tree for each application thread. From these files we can extract further information, e.g., callers and callees of each method. Blocked threads represent threads that fail to immediately enter the synchronized method/block. Garbage collected

objects allow estimating the load of the garbage collector. All these statistics are fed to JPAD that uses them to detect software performance antipatterns.

3.3 Antipattern Detection

The detection of software performance antipatterns relies on our interpretation of the natural language specification of software performance antipatterns [26]. Specifically, after extracting key properties, we implement the antipattern detection rule if a match with profiling data exists. Our effort is mainly devoted to match high-level guidelines and make them concrete for the inspection of Java applications. In this context, the difficulty relies on combining different sources of information (e.g., a class calling a high number of methods, high CPU usage), since the intrinsic nature of software antipatterns is to look for various performance problems that may arise when applications are running. Since the specification of antipatterns cannot be completely precise [26], the conditions and key properties we check as our detection rules approximate the antipatterns. These conditions are neither sufficient nor necessary, we empirically validate their relevance on spotting performance issues.

Table 1 reports a subset of seven software performance antipatterns proposed in [26], and we motivate in the following the choice of implementing them. The reason why some antipatterns cannot be automatically detected is argued in Appendix B, available in the online supplemental material. The first column reports the *name* of the performance antipatterns, and the second column describes the *problem* expressed in natural language. The third column lists the extracted *key properties* (along with the implemented helper functions) and provides a match with profiling data. The fourth column lists thresholds and offsets that are included in the detection algorithms. Helper functions are briefly described in Table 2, thresholds/offsets are presented in Table 3 along with the heuristics adopted in case users do not define their own preferences on perceived performance issues. Offsets are all set to 5% since we are interested to capture small deviations from average values. Thresholds are set to 10%, we refer to [59] where the CPU load in the *idle* phase is estimated to be 7% on average, and we are interested to exclude outliers and fluctuations that might be due to system's internal routines.

The detection algorithms and implementation details of these antipatterns are given in Appendix A, available in the online supplemental material, along with the match between textual descriptions and detection rules. In the following, we briefly discuss our interpretation of antipatterns, and their key properties.

Circuitous Treasure Hunt (CTH) occurs when a Java application must perform a large number of (database) queries to manage a request. This problem can be generalized considering a method that performs a chain of queries where the result of one query is used to build the next one, instead of writing a single, and more complex request. As key properties, we check the average number of calls performed by each thread (*call-tree-by-thread* view of YourKit), along with the average processor utilization (*chart-cpu-usage* view of YourKit) indicating if the system performance suffers.

Extensive Processing (EP) occurs when a long running job monopolizes the processor and creates a queue of processes

1. [Online]. Available: <https://doi.org/10.5281/zenodo.5878953>

TABLE 1
Software Performance Antipatterns - Key Properties Extracted From Natural Language Specification [26]

Name	Problem specification	Key properties (and helper functions)	Thresholds and offsets
Circuitous Treasure Hunt	Occurs when an object must look in several places to find the information it needs. If a large amount of processing is required for each look, performance will suffer.	Large number of calls (<i>getMethodCountMap</i> , <i>getAvgMethodCount</i>), high execution time (<i>getAvgTime</i>), and high resources utilization (<i>getAvgCpuUsage</i>).	countOffset, cpuTh
Extensive Processing	Occurs when extensive processing monopolizes the processors and impedes the overall response time.	Large number of blocked threads (<i>countThreads</i> , <i>countBlockedThreads</i>), high execution time (<i>getTotExecTime</i> , <i>getAvgExecTime</i> , <i>getHsExecTime</i>).	execTime-Offset
Wrong Cache Strategy	Caching too many objects (or objects that are rarely used) quickly changes the advantage of caching into a disadvantage due to higher memory usage.	High memory usage from methods (<i>getAvgMemUsage</i> , <i>getHsMemUsage</i>), i.e., larger than the average processor usage (<i>getHsCpuUsage</i>).	memUsage-Offset
Blob	Occurs when a single class either (i) performs all the work of an application or (ii) holds all the application data. Either manifestation results in excessive message traffic that degrades the performance.	Large number of calls performed and received by the methods (<i>getCallersMap</i> , <i>getCalleesMap</i> , <i>getAvg</i>) reflecting on high processor and/or memory utilization (<i>getAvgCpuUsage</i> , <i>getAvgMemUsage</i>).	msgOffset, cpuTh, memTh
Tower of Babel	Occurs when processes excessively convert, parse, and translate internal data into a common exchange format.	Large execution time of methods (<i>getMethodCallsMap</i> , <i>getTotExecTime</i> , <i>getAvgExecTime</i> , <i>getHsExecTime</i>) due to specific activities (<i>cpu-hotspots</i>).	execTime-Offset
Empty Semi Trucks	Occurs when an excessive number of requests is required to perform a task. It may be due to inefficient use of available bandwidth, an inefficient interface, or both.	Large number of method calls (<i>getCalleesMap</i> , <i>getAvg</i> , <i>countCalls</i>) whose execution time (<i>getCoTime</i>) follows a deterministic distribution.	msgOffset
Excessive Dynamic Allocation	Occurs when an application unnecessarily creates and destroys large numbers of objects during its execution. The overhead has a negative impact on performance.	Large number of objects stored by the Java garbage collector (<i>getGcObjs</i> , <i>getHsGcObjs</i>), high memory utilization (<i>getAvgMemUsage</i>).	memTh, gcObjects-Offset

TABLE 2
Software Performance Antipatterns - Helper Functions

Name	Description
<i>getMethodCountMap</i>	It calculates the minimum, maximum, or average (depending on the <i>option</i> input parameter) number of calls in a method
<i>getAvgMethodCount</i>	It calculates the average number of calls in a specific hotspot method, to be compared with the average of all hotspots
<i>getAvgTime</i>	It calculates the average execution time of all hotspots, to be compared with the executime time of a specific hotspot method
<i>getAvg[Cpu,Mem]Usage</i>	It calculates the average utilization values showed by all system hardware (CPU or memory) resources
<i>count[Blocked]Threads</i>	It counts the number of (blocked) threads, to be compared with active threads to understand delays in the execution of methods
<i>get[Tot,Avg,Hs]ExecTime</i>	It retrieves the (total, average, or related to a specific hotspot method) execution time, to be compared with other methods
<i>get[Callers,Callees]Map</i>	It retrieves all methods that are callers or callees, and it is used to quantify the number of calls performed/received by a method
<i>getAvg</i>	It takes as input parameter a map of methods (e.g., through <i>getMethodCallsMap</i> , see next helper function) and calculates the average
<i>getMethodCallsMap</i>	It retrieves all methods that are invoked by other methods, to be compared with average values for spotting most invoked methods
<i>countCalls</i>	It takes as input a callee method and it calculates the average number of times that hotspots invoke their callees
<i>getCvTime</i>	It takes as input a callee method and it retrieves the coefficient of variation of the callee's service time
<i>get[Hs]GcObjs</i>	It retrieves the average number of garbage collected objects in the (hotspot) methods, thus to spot if there are many unused objects

that cannot be executed until the computation of such job is completed. As key properties, we monitor the average number of blocked threads (*monitor-usage* view of YourKit), and

the execution time of Java methods (*call-tree-all-threads* view of YourKit), selecting only those methods that are showing a large number of the identified blocked threads.

TABLE 3
Software Performance Antipatterns - Thresholds and Offsets

Name	Description	Heuristic
countOffset	Addendum for methods' number of calls	Set to 5%, it augments the average value for the number of methods' calls
cpuTh	Upper bound for CPU utilization	Set to 10%, lower values are considered part of the system's internal routines
execTimeOffset	Addendum for methods' execution time	Set to 5%, it augments the average value for the execution time of methods
memUsageOffset	Addendum for memory utilization	Set to 5%, it augments the average value for the memory resources utilization
msgOffset	Addendum for number of exchanged messages	Set to 5%, it augments the average value for the methods' exchanged messages
memTh	Upper bound for memory utilization	Set to 10%, lower values are considered part of the system's internal routines
gcObjectsOffset	Addendum for number of garbage collected objects	Set to 5%, it augments the average value for the number of garbage collected objects

Wrong Cache Strategy (WCS) occurs when too many objects (or objects hardly ever used) are cached. This leads to generate performance overhead resulting in high memory usage. As key properties, we check the average memory usage of methods (*method-list-allocation* view of YourKit), and we are interested to verify whether the memory is more used than the processor (*method-list-cpu* view of YourKit).

Blob occurs in two different scenarios that contribute to distinguish two different types of antipattern instances in Java applications, i.e., *Blob-Controller* and *Blob-DataContainer*. The former is observed when a class centralizes many responsibilities, delegating minor roles to other classes. Classes affected by this problem usually are complex controllers which depend on simpler classes (with little to no computation). The latter case is observed when a class includes most of the application data and other functions need to access that class to retrieve/update the data. As key properties, we monitor the average number of calls performed and received by the methods (*call-tree-all-threads* view of YourKit) to determine if these methods can be classified as potential controllers or data containers, respectively. Besides, we also check if such large number of calls impacts on the average usage of processor and memory (*chart-cpu-usage* and *chart-heap-memory-usage* views of YourKit).

Tower of Babel (ToB) occurs when some data is translated into an exchange format, such as XML, by the sending process. This data is later parsed and translated into an internal format by the receiving process. This means that the system may spend most of its time processing the text when the translation and parsing of data formats are excessive. As key properties, we check all the methods showing a large execution time (*call-tree-all-threads* view of YourKit), and then we inspect the names of such methods (*cpu-hotspots* view of YourKit). We check if there is a match with some specific keywords (i.e., "converse", "parse", and "translate"), thus to associate the performance overhead to the processing of exchange format data.

Empty Semi Trucks (EST) occurs when an excessive number of requests is required to perform a task. As key properties, similarly to the Blob-controller, we consider the number of calls performed by the methods (*call-tree-all-threads* view of YourKit), and we verify if the execution time of such calls shows a small coefficient of variation, i.e., the

execution time follows a deterministic distribution (*call-tree-all-threads* view of YourKit). This way, we aim to capture the peculiarity of this antipattern when inefficiently using the bandwidth and/or interfaces.

Excessive Dynamic Allocation (EDA) occurs when an application unnecessarily creates and destroys objects. As key properties, we monitor the objects collected by the Java garbage collector (*method-list-garbage* view of YourKit), and we want to recognize those situations leading to a performance overhead, i.e., the average memory utilization (*chart-heap-memory-usage* view of YourKit) is high.

It is worth remarking that our detection algorithms may be *orthogonal*, i.e., they may return the simultaneous occurrence of multiple antipatterns since they share the verification of some performance indices, e.g., the CPU utilization. This is not a drawback of the approach since it may happen that a performance issue may be caused by the simultaneous presence of multiple bad practices [60].

4 EXPERIMENTATION

This section is organized as follows. We first present the research questions (Section 4.1), followed by the description of the analyzed real-world applications (Section 4.2). We describe the experiment design (Section 4.3), and we discuss the obtained experimental results (Section 4.4).

4.1 Research Questions

The purpose of our experimental evaluation is threefold: (i) it shows that JPAD is efficient and accurate when applied to real medium/large-sized systems; (ii) it compares JPAD with state-of-the-art approaches on the detection of real-life performance issues; (iii) it provides empirical evidence on the usefulness of detecting antipatterns. In particular, we aim to answer three research questions:

RQ₁ *Efficiency and accuracy of the antipatterns' detection rules:* Are the proposed detection rules efficient and accurate? Does JPAD efficiently detect antipattern instances in real-world case studies? *Motivation.* We want to provide developers with an efficient and accurate framework that points out performance

TABLE 4
Analyzed Systems

Name	Version	Domain	# of Classes
CloudStore	2	E-commerce	68
TeaStore	1.4.1	Microservices	138
WebGoat	8.0.0.M26	E-learning	301
TrainTicket	1.0	Microservices	584
OpenMRS	2.9.0	Medical Record System	1093

issues, if any. By evaluating the efficiency of JPAD, we can understand whether developers are motivated to apply our approach in practice.

RQ₂ *Comparison with state-of-the-art approaches:* Are the proposed detection rules comparable to other techniques in the literature? Does JPAD capture complex performance problems that are not recognized by available tools? *Motivation.* We want to compare our framework with state-of-the-art approaches that deal with the detection of performance issues. By comparing JPAD to other tools, we can study whether developers are motivated to use our framework.

RQ₃ *Implication of applying antipattern-based refactorings:* What happens when an antipattern is solved? How does the number of detected antipattern instances for the refactored system change? What is the effect on the overall system performance? **Motivation.** The goal of our approach is to spot performance problems, so that developers are aware of possible shortcomings in some portions of the code. In this research question, we aim to answer at which extent our approach benefits developers interested to know if code fixings improve the system performance.

To answer the first two research questions, we analyze real-world Java applications under different loads, i.e., varying the number of clients and the execution time. For each combination of (i) number of clients and (ii) duration (in minutes), the application is profiled, data is collected, and JPAD is used to detect the software performance antipatterns. Moreover, to further stress the benefit of the proposed framework and to answer the third research question, we provide empirical evidence of its impact by solving one instance of detected antipatterns (for one of the analyzed systems), and showing variations in the number of detected instances and the system performance.

4.2 Analyzed Systems

We select five Java applications that are highly concerned with their performance and have been studied in prior research [46], [61], [62]. These applications show a different complexity (in terms of number of classes) and belong to different domains, see Table 4 for their main characteristics. The rationale for selecting these five subject systems is that they provide evidence of all the seven performance antipatterns, as later described in the experimental results (see Section 4.4). Specifically, systems from [46] do not include ToB and EST antipattern instances. We select one system from [61] to provide evidence on the ToB antipattern. To strengthen the analysis of the EST antipattern, we borrow the subject system used in [62]. Hereinafter, a brief description of the analyzed systems is provided.

TABLE 5
Analyzed Issues Extracted From Chen et al. [23]

Name	Issue ID	Issue fixing commit	Issue inducing commit
Hadoop	YARN-4307	308d63f	e914220 7af5d6b
	YARN-7102 HDFS-12754	ff8378e 738d1a2	528b809 decf8a6
Cassandra	CASSANDRA-13794	f93e6e3	88d2ac4

- CloudStore [63] is a free software synchronizing files between multiple devices. Its primary focus is on preventing data loss and unauthorized access.
- TeaStore [64] emulates a basic web store for automatically generated tea and tea supplies. It has been published in [65], and later largely used as a micro-service reference system and test application.
- WebGoat [66] is a deliberately insecure web application maintained by OWASP and designed to teach web application security lessons.
- TrainTicket [67] is the largest benchmark for micro-service architectures in the literature. It provides train ticket booking functionalities and is used for fault analysis and error prediction [62], [68].
- OpenMRS [69] is a free medical record system for health care providers. It is a modular open-source web application used by over 40 countries to improve health care delivery in resource-constrained environments.

Moreover, to compare JPAD with state-of-the-art approaches [23], [45], [53], we consider four real-life performance issues (see details in Table 5) of two other systems:

- Hadoop [70] is a framework which performs data processing in a reliable, efficient, high fault tolerance, low cost, and scalable manner.
- Cassandra [71] is a distributed NoSQL database management system; fault-tolerance on commodity hardware makes it suitable for mission-critical data.

Table 5 reports the code commits (inducing and fixing the four real-life performance issues) that we analyze, according to the study on performance regressions presented in [23]. The motivation of selecting these specific four performance issues is the following. Hadoop commits are evaluated in [23] by means of multiple performance metrics (i.e., response time, CPU and memory utilization, I/O operations) and we focus on those issues that have been highlighted as particularly complex, since such issues are not predicted by any of the considered metrics. About Cassandra, there is only one issue that is not predicted by *PerffIT* in [23], and this is why we concentrate our effort on investigating that specific issue. Summarizing, our investigation includes those specific four issues (triggering the analysis of nine code commits) since they are more relevant to conduct a comparison.

4.3 Experimental Setup

In the following, we discuss the design choices taken to run experiments and avoid biases in results.

System Workload Specification. To avoid biases when profiling applications, we test several workloads acting in the

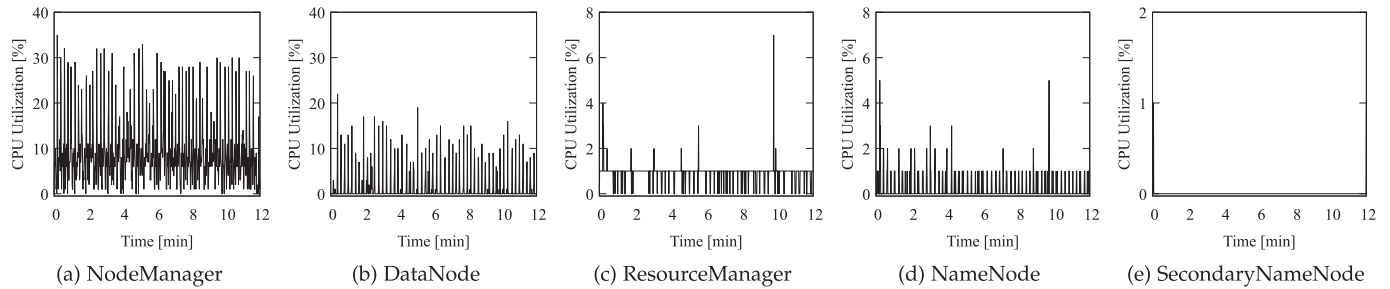


Fig. 2. CPU usage over 12 minutes for the five Hadoop components. The profiled Hadoop commit is *decf8a6*.

considered systems for a different duration. The choice on the number of clients and the duration of load testing is not trivial, however our assumption is that this is decided by software developers that are aware of load conditions and system’s dynamics. Subject systems in Table 4 run for 3, 6, and 12 minutes with 25, 50, 75, and 100 clients. A larger variation in the number of clients is considered for the issues listed in Table 5 and related to Hadoop and Cassandra; these systems are tested with 1, 10, 100, 500, and 1k clients running for 3, 6, and 12 minutes. Each combination of number of clients (C) and duration of the testing (D) represents an input we use in our detection. Such a combination leads to a system configuration that is labeled in the following as $C - D$, for instance the system configuration “1000-6” means to consider 1000 clients with 6-minute duration of testing. Overall, 300 different system configurations are analyzed to explore a variegated set of systems’ characteristics.

Load Test Definition. To avoid biases in the load testing, we explore benchmarks stressing different aspects (e.g., input/output operations, end-users services) of considered applications. For instance, we make use of available benchmarks for Hadoop and Cassandra (i.e., TestDFSIO [72] and Cassandra Stress [73]) stressing write and read operations. Locust [74] (i.e., a Python-based load testing tool) is adopted for systems reported in Table 4 since ready-to-use benchmarks are not available. We stress end-users services identified as crucial for the application, e.g., in TrainTicket we generate requests for monitoring the security service invoked when users make a reservation. As anticipated in Section 3.1, to support software engineers in the selection of system functionalities to be monitored, as a rule of thumb, we consider the CPU utilization of system software components. This analysis provides support to the component choosing stage of the approach that decides which component (s) to test, and they are given as input to our detection. As an example, let us consider Fig. 2 that shows the CPU utilization (observed when executing I/O operations) of all Hadoop components. Results show that two components have a non negligible utilization, see Figs. 2a and 2b with CPU utilization varying up to 40%: (i) the *NodeManager* which is in charge of monitoring the resource usage of a node and reporting such information to the *ResourceManager* and (ii) the *DataNode* that stores data. Hence, these two components are selected to be profiled. Figs. 2c, 2d, and 2e instead indicate that there are underutilized components: *ResourceManager* and *NameNode* show a CPU utilization lower than 8%, whereas the CPU used by *SecondaryNameNode* is even less than 2%. This denotes that it is not relevant to profile such components.

Load Test Execution. To avoid biases in the obtained results, all experiments are run on a single node cluster

deployed on a dedicated machine (with a 2.80 GHz quad-core CPU and 16 GB memory) to easily monitor used resources and avoid misleading performance results due to background activities. Readings extracted from CPU, memory, and garbage collector snapshots taken for each configuration are fed to JPAD that automatically detects performance antipatterns.

Antipattern Detection and Thresholds Setup. We recall that most of thresholds are calculated through heuristics that consider average values and offsets, see Table 3. This means that each system configuration includes thresholds that vary when changing number of clients and duration. In the following, as illustrative examples, we report the threshold values used with the lowest (25 clients, 3 minutes) and the largest (1k clients, 12 minutes) boundary values adopted in the proposed experimentation. Table 6 reports the threshold numerical values used to analyze the systems when setting the number of clients to 25 and the duration of the run to 3 minutes. In the first row of Table 6 we can notice that *countTh* threshold is calculated on average values, and it shows a large variation among the analyzed systems (e.g., 108 for OpenMRS, and 1198 for TrainTicket). There are other thresholds that vary less even if they are calculated with average values, e.g., *msgTh* varies between 1.40 and 1.58. Some other thresholds are instead fixed to values that usually are extracted from system requirements, e.g., *memTh* is fixed to 10% for all the systems. Table 7 reports thresholds

TABLE 6
Threshold Values Used for Analyzing Systems
With #clients = 25 and duration = 3 Minutes

Antipatterns	Thresholds	Systems				
		CloudStore	TeaStore	WebCoat	TrainTicket	OpenMRS
CTH	<i>countTh</i>	309	148	170	1198	108
	<i>cpuTh</i>	10%	10%	10%	10%	10%
	<i>option</i>	avg	avg	avg	avg	avg
EP	<i>execTimeTh</i>	5.29%	5.40%	5.88%	5.00%	5.10%
WCS	<i>memUsageTh</i>	16%	16%	13%	27%	76%
	<i>callersTh</i>	19.36	24.00	4.76	9.28	8.58
Blob	<i>calleeTh</i>	426.75	427.12	596.05	2637.25	1562.40
	<i>cpuTh</i>	10%	10%	10%	10%	10%
	<i>memTh</i>	10%	10%	10%	10%	10%
ToB	<i>execTimeTh</i>	5.29%	5.40%	5.88%	5.00%	5.10%
EST	<i>msgTh</i>	1.05	1.05	1.40	1.47	1.58
EDA	<i>gcObjectsTh</i>	8711349.8	4387604.0	1562411.2	3652923.5	3538209.5
	<i>memTh</i>	10%	10%	10%	10%	10%

TABLE 7

Threshold Values Used for Analyzing Hadoop (Two Components) and Cassandra With #Clients = 1k and Duration = 12 Minutes

Antipatterns	Thresholds	Hadoop – NodeManager							Hadoop – DataNode							Cassandra			
		decf8a6	e914220	7af5d6b	528b809	308d63f	738d1a2	ff8378e	decf8a6	e914220	7af5d6b	528b809	308d63f	738d1a2	ff8378e	88d2ac4	f93e6e3		
CTH	<i>countTh</i>	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	44	26
	<i>cpuTh</i>	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%
EP	<i>option</i>	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg	avg
	<i>execTimeTh</i>	5.06%	5.04%	5.06%	5.06%	5.04%	5.02%	5.02%	5.29%	5.21%	5.26%	5.25%	5.39%	5.32%	5.28%	5.24%	5.25%		
WCS	<i>memUsageTh</i>	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	5.00%	13.00%	26.00%		
	<i>callersTh</i>	11.12	11.95	11.89	11.33	10.15	10.26	9.56	2.42	2.91	2.51	2.59	2.04	2.28	2.31	4.61	4.51		
Blob	<i>callesTh</i>	1.05	1.05	0.70	1.05	0.70	–	–	–	–	–	–	–	–	–	82.35	44.28		
	<i>cpuTh</i>	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%		
	<i>memTh</i>	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%		
ToB	<i>execTimeTh</i>	5.06%	5.04%	5.06%	5.06%	5.04%	5.02%	5.02%	5.29%	5.21%	5.26%	5.25%	5.39%	5.32%	5.28%	5.24%	5.25%		
	<i>msgTh</i>	–	–	–	–	–	–	–	–	–	–	–	–	–	–	1.05	1.05		
EDA	<i>gcObjectsTh</i>	–	–	–	–	–	–	–	–	–	–	–	–	–	–	15980736.0	16730745.2		
	<i>memTh</i>	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%		

The symbol “–” indicates that a threshold is not calculated due to lack of data.

used for all commits of Hadoop (*NodeManager* and *DataNode*) and Cassandra, when systems are loaded with 1k clients for 12 minutes. Note that some commits do not provide data needed to calculate the thresholds, e.g., *738d1a2* and *ff8378e* lack information for deriving *callesTh* related to the Blob antipattern disabling its detection. Besides, the same threshold may vary differently based on the considered applications. For example, *callersTh* related to the Blob antipattern changes significantly when considering the two Cassandra commits (i.e., 82.35 for *88d2ac4* and 44.28 for *f93e6e3*) and shows slight variations when observing each Hadoop component separately. Instead, relevant variations of this threshold are observed also for Hadoop when the two components are compared, i.e., it varies in the range 9.56–11.95 for the *NodeManager* and in the range 2.04–2.91 for the *DataNode*. Threshold and offset values are given as input to JPAD, and these values can be easily modified in case system stakeholders express their own performance requirements.

Summarizing, our experimental setup includes the following inputs: (i) system workload specification, i.e., number of clients and duration of the testing, (ii) load test definition, i.e., which component(s) to test, and (iii) threshold and offset setup, i.e., numerical values for antipatterns’ thresholds and offsets. Our experimentation investigates the variation of these inputs and how they affect the detection accuracy and results, see more details in Tables 8, 9, 10, 11, and 12, and Appendix C, available in the online supplemental material.

4.4 Experimental Results

This section presents experimental results answering our three research questions (see Section 4.1).

4.4.1 Effectiveness of the Antipatterns’ Detection Rules

To answer *RQ₁*, Table 8 shows the software performance antipattern instances detected by JPAD for each system, when varying the number of clients and the monitoring duration. The last column of Table 8 shows the detection time (in seconds), i.e., the time required by JPAD to complete the analysis. This way, we aim to analyze the

scalability of the tool when exposed to medium- and large-sized systems.

We observe that the number of detected antipatterns may increase with the number of clients. For example, for the WebGoat case study, the number of detected antipatterns with 25 clients (i.e., 3, 6, 0, 5, 4, 0, 3) is smaller than (or equal to) the case with 100 clients (i.e., 4, 7, 2, 5, 4, 0, 4). The number of detected instances can also decrease for a larger number of clients, e.g., OpenMRS shows 2 instances of the CTH antipattern with 25 clients and 0 with 100 clients. However, we notice that this may lead to generate further instances of different antipatterns types, in fact with 100 clients we get 3 Blob instances instead of 2 (observed with 25 clients). Similarly, CloudStore shows 7 instances of the EP antipattern and 8 instances of the Blob antipattern with 25 clients. With 100 clients instead we get 6 instances of the EP antipattern and 7 instances of the Blob antipattern, at the cost of 2 instances of the EDA antipattern (not observed with 25 clients). This may be due to performance issues showing up to (previously not critical) parts of the system, i.e., affecting different system elements only when the number of clients increases. As a result, a larger workload may produce more instances of other antipattern types.

A longer experiment duration might reduce the number of detected instances due to the performance problems flattening/elevating in different parts of the system depending on the application peculiarities. In our experiments, OpenMRS benefits from a longer run time, whereas all other systems (especially TeaStore) generally show a constant or higher number of antipattern instances. The number of antipattern instances detected when OpenMRS is run for 3 minutes with 25 clients (i.e., 11) decreases when the same application is observed for 12 minutes (i.e., 3). For all other systems, it is unlikely that the number of detected instances decreases when the experiment duration increases. In a few cases, the opposite trend is observed, e.g., there are 21 antipattern instances when WebGoat is run with 25 clients for 3 minutes and 25 instances when the same application is run for 12 minutes. A further example is represented by the CloudStore system that shows 17 instances with 100 clients for 3 minutes and 19 instances when running for 12 minutes.

TABLE 8
Analysis of How System Workload Specification Inputs Affect the Detection. Number of Detected Antipattern Instances and Scalability of JPAD Across the Analyzed Systems.

	Config.	CTH	EP	WCS	Blob	ToB	EST	EDA	time (sec)
CloudStore	25-3	2	7	0	8	0	0	0	1.9
	25-6	2	7	0	8	0	0	2	0.9
	25-12	2	6	0	7	0	0	2	1.0
	50-3	2	8	0	8	0	0	0	1.1
	50-6	2	7	0	7	0	0	0	1.1
	50-12	2	8	0	9	0	0	0	1.5
	75-3	2	6	0	7	0	0	0	1.0
	75-6	2	7	2	8	0	0	0	1.2
	75-12	2	6	0	7	0	0	0	1.6
	100-3	2	6	0	7	0	0	2	1.1
	100-6	2	7	2	8	0	0	2	1.4
	100-12	2	6	2	7	0	0	2	1.8
TeaStore	25-3	6	9	0	6	0	0	2	144.1
	25-6	6	9	0	6	0	0	2	166.8
	25-12	6	9	0	6	0	0	2	183.6
	50-3	6	9	0	6	0	0	2	103.8
	50-6	6	9	0	6	0	0	2	158.2
	50-12	6	9	0	6	0	0	2	178.8
	75-3	5	9	0	5	0	0	2	172.8
	75-6	6	9	0	6	0	0	2	245.4
	75-12	6	9	0	5	0	0	2	218.3
	100-3	5	9	0	6	0	0	2	282.5
	100-6	6	9	0	5	0	0	2	257.6
	100-12	6	9	0	5	0	0	2	220.3
WebGoat	25-3	3	6	0	5	4	0	3	29.8
	25-6	3	6	0	5	4	0	3	30.6
	25-12	4	7	0	6	4	0	4	29.9
	50-3	4	7	2	5	4	1	4	33.6
	50-6	4	7	1	5	4	1	4	37.4
	50-12	4	7	2	5	4	0	4	34.3
	75-3	3	6	1	5	4	0	3	30.1
	75-6	4	7	2	5	4	0	4	34.4
	75-12	4	7	2	5	4	0	4	33.7
	100-3	4	7	2	5	4	0	4	32.6
	100-6	4	7	2	5	4	0	4	30.6
	100-12	4	7	2	5	4	1	4	30.2
TrainTicket	25-3	1	6	1	1	0	1	0	1.0
	25-6	1	4	2	1	0	1	0	1.1
	25-12	1	5	2	1	0	1	0	1.1
	50-3	1	4	0	1	0	1	0	1.1
	50-6	1	4	0	1	0	1	0	1.2
	50-12	1	5	1	1	0	2	0	1.4
	75-3	1	6	1	1	0	1	0	1.3
	75-6	1	5	1	1	0	1	0	1.2
	75-12	1	5	1	1	0	1	0	1.4
	100-3	1	5	0	1	0	1	0	1.2
	100-6	1	5	0	1	0	1	0	1.3
	100-12	1	5	0	1	0	1	0	1.7
OpenMRS	25-3	2	3	2	2	0	0	2	0.6
	25-6	0	0	0	1	0	0	1	0.5
	25-12	0	0	1	1	0	0	1	0.7
	50-3	0	1	0	1	0	0	0	0.5
	50-6	0	1	0	1	0	0	0	0.5
	50-12	0	1	0	1	0	0	0	0.6
	75-3	1	2	1	3	0	0	1	0.6
	75-6	1	2	1	3	0	0	1	0.7
	75-12	0	1	0	1	0	0	0	0.6
	100-3	0	2	0	3	0	0	1	0.4
	100-6	0	1	0	1	0	0	0	0.6
	100-12	0	1	0	1	0	0	0	0.6

From the collected results, no straightforward relationship between the number of clients and the number of detected antipatterns is observed. Similar observations are drawn when considering the experiment duration since such duration is not linearly related with the number of detected antipattern instances, thus to confirm the complexity in the performance analysis of Java applications.

Table 8 also reports the time required by JPAD to detect antipatterns, and we can notice that it varies between 0.4 and 282.5 seconds (i.e., 4.7 minutes). The average detection

time observed over all systems and loads is 45.9 seconds (i.e., less than a minute) with OpenMRS and TeaStore taking the shortest and longest average detection time (i.e., 0.6 and 194.4 seconds corresponding to 3.2 minutes), respectively.

Since TeaStore and WebGoat show the longest time to complete the detection, we further inspect these systems. For TeaStore we notice that Blob, ToB, and EST take longer than other antipatterns to be analyzed. This is due to the large size of the call tree (i.e., a key property analyzed to detect these antipatterns, see Table 1) of the TeaStore system. When Blob, ToB, and EST are excluded from the analysis, JPAD takes only 1.2 seconds on average to analyze TeaStore. For WebGoat, we observe that the detection of EST antipattern takes longer than others, i.e., the detection takes 3.5 seconds on average when EST is excluded. We think this is due to the large number of methods that are invoked (namely the *caltees*) in this application. In fact, these methods must be analyzed by the EST detection algorithm to check their coefficient of variation, i.e., if the execution time follows a deterministic distribution, thus to spot inefficiency on the usage of resources.

Inspired by [46], we build a *ground truth* to investigate the accuracy (along with precision, recall, and F1 score) of our detection rules. The ground truth is constituted of all the hotspot methods (i.e., over approximation of the detection outcome) since JPAD analyzes their design and performance characteristics to detect antipatterns. We exploit the variation in the system workload specification inputs to decide if a hotspot method is correctly detected as source of performance problems. We determine that a hotspot method has performance issues if it is reported as a violation of antipatterns while analyzing different system configurations. If a hotspot method is detected as an antipattern instance in at least half (i.e., 6 over a total of 12) analyzed configurations, then such a method is considered as a positive instance of that antipattern, otherwise it is classified as a negative instance. Let us consider as example `executeQuery()` that is a hotspot method in OpenMRS. It is detected as EP antipattern in 10 (out of 12) analyzed configurations, hence such method is considered as a positive instance of EP in OpenMRS. Another hotspot method of OpenMRS is `includeFragment(String, String)` but it is detected as EP antipattern in only 2 (out of 12) analyzed configurations. Consequently, this method is considered as a negative instance.

The *confusion matrix* is built for each system configuration as follows. True positives (TP) are the methods for which an antipattern instance is detected by both the ground truth and JPAD. True negatives (TN) are those hotspot methods that are detected as specific antipattern instances by neither the ground truth nor JPAD. False positives (FP) are defined by calculating hotspot methods that are detected as antipatterns by JPAD, but are not in the ground truth. False negatives (FN) are calculated by counting methods that are not detected by JPAD but are in the ground truth. The confusion matrix for the analyzed system configurations in all considered applications (when detecting the Blob antipattern) is reported in Table 9. Confusion matrices derived for other antipatterns are omitted for the sake of space, but they are reported as part of replication data [28].

TABLE 9
Confusion Matrix of All Considered Configurations and Applications When Detecting the Blob antipattern

	25-3	25-6	25-12	50-3	50-6	50-12	75-3	75-6	75-12	100-3	100-6	100-12																																																																																																
CloudStore	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1	<table border="1"><tr><td>8</td><td>21</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	8	21	0	0	0	0	1	1	<table border="1"><tr><td>7</td><td>21</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	7	21	0	1	0	0	1	1
	8	21	0	0																																																																																																								
	0	0	1	1																																																																																																								
	8	21	0	0																																																																																																								
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
8	21	0	0																																																																																																									
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
8	21	0	0																																																																																																									
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
8	21	0	0																																																																																																									
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
8	21	0	0																																																																																																									
0	0	1	1																																																																																																									
7	21	0	1																																																																																																									
0	0	1	1																																																																																																									
TeaStore	<table border="1"><tr><td>6</td><td>23</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	6	23	0	0	0	0	1	1	<table border="1"><tr><td>6</td><td>23</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	6	23	0	0	0	0	1	1	<table border="1"><tr><td>6</td><td>23</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	6	23	0	0	0	0	1	1	<table border="1"><tr><td>6</td><td>23</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	6	23	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1	<table border="1"><tr><td>6</td><td>23</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	6	23	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1	<table border="1"><tr><td>5</td><td>23</td><td>0</td><td>1</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	23	0	1	0	0	1	1									
	6	23	0	0																																																																																																								
	0	0	1	1																																																																																																								
	6	23	0	0																																																																																																								
0	0	1	1																																																																																																									
6	23	0	0																																																																																																									
0	0	1	1																																																																																																									
6	23	0	0																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
6	23	0	0																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
5	23	0	1																																																																																																									
0	0	1	1																																																																																																									
WebGoat	<table border="1"><tr><td>4</td><td>25</td><td>1</td><td>1</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr></table>	4	25	1	1	1	1	1	1	<table border="1"><tr><td>4</td><td>25</td><td>1</td><td>1</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr></table>	4	25	1	1	1	1	1	1	<table border="1"><tr><td>5</td><td>25</td><td>1</td><td>1</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr></table>	5	25	1	1	1	1	1	1	<table border="1"><tr><td>4</td><td>25</td><td>1</td><td>1</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr></table>	4	25	1	1	1	1	1	1	<table border="1"><tr><td>4</td><td>25</td><td>1</td><td>1</td></tr><tr><td>1</td><td>1</td><td>1</td><td>1</td></tr></table>	4	25	1	1	1	1	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1	<table border="1"><tr><td>5</td><td>26</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	5	26	0	0	0	0	1	1									
	4	25	1	1																																																																																																								
	1	1	1	1																																																																																																								
	4	25	1	1																																																																																																								
1	1	1	1																																																																																																									
5	25	1	1																																																																																																									
1	1	1	1																																																																																																									
4	25	1	1																																																																																																									
1	1	1	1																																																																																																									
4	25	1	1																																																																																																									
1	1	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
5	26	0	0																																																																																																									
0	0	1	1																																																																																																									
TrainTicket	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>16</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	16	0	0	0	0	1	1									
	1	16	0	0																																																																																																								
	0	0	1	1																																																																																																								
	1	16	0	0																																																																																																								
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
1	16	0	0																																																																																																									
0	0	1	1																																																																																																									
OpenMRS	<table border="1"><tr><td>1</td><td>11</td><td>1</td><td>0</td></tr><tr><td>1</td><td>1</td><td>0</td><td>0</td></tr></table>	1	11	1	0	1	1	0	0	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>10</td><td>2</td><td>0</td></tr><tr><td>1</td><td>10</td><td>2</td><td>0</td></tr></table>	1	10	2	0	1	10	2	0	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1	<table border="1"><tr><td>1</td><td>12</td><td>0</td><td>0</td></tr><tr><td>0</td><td>0</td><td>1</td><td>1</td></tr></table>	1	12	0	0	0	0	1	1									
	1	11	1	0																																																																																																								
	1	1	0	0																																																																																																								
	1	12	0	0																																																																																																								
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	10	2	0																																																																																																									
1	10	2	0																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
1	12	0	0																																																																																																									
0	0	1	1																																																																																																									
	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN	<table border="1"><tr><td>TP</td><td>TN</td><td>FP</td><td>FN</td></tr></table>	TP	TN	FP	FN																																																					
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									
TP	TN	FP	FN																																																																																																									

TABLE 10
JPAD Detection Performance

		CTH	EP	WCS	Blob	ToB	EST	EDA
CloudStore	A	1.00	0.98	0.98	0.98	1.00	1.00	0.97
	P	1.00	0.98	-	0.99	-	-	-
	R	1.00	0.93	-	0.94	-	-	-
	F1	1.00	0.95	-	0.96	-	-	-
TeaStore	A	0.99	1.00	1.00	0.99	1.00	1.00	1.00
	P	1.00	1.00	-	1.00	-	-	1.00
	R	0.97	1.00	-	0.94	-	-	1.00
	F1	0.99	1.00	-	0.97	-	-	1.00
WebGoat	A	0.99	0.99	0.98	0.98	1.00	0.99	0.99
	P	1.00	1.00	-	0.94	1.00	-	1.00
	R	0.94	0.96	0.67	0.95	1.00	-	0.94
	F1	0.97	0.98	-	0.94	1.00	-	0.97
TrainTicket	A	1.00	0.98	0.97	1.00	1.00	1.00	1.00
	P	1.00	0.97	-	1.00	-	0.96	-
	R	1.00	0.95	0.58	1.00	-	1.00	-
	F1	1.00	0.96	-	1.00	-	0.98	-
OpenMRS	A	0.97	0.96	0.97	0.96	1.00	1.00	0.96
	P	-	-	-	0.79	-	-	-
	R	-	0.83	-	1.00	-	-	-
	F1	-	-	-	0.88	-	-	-
Avg.	A	0.99	0.97	0.97	0.98	1.00	1.00	0.98
	P	1.00	0.98	-	0.93	1.00	0.96	1.00
	R	0.98	0.92	0.62	0.97	1.00	1.00	0.94
	F1	0.99	0.97	-	0.95	1.00	0.98	0.97

The symbol “-” means that the metric is undefined.

Accuracy, precision, and recall are defined for the detection of each antipattern and calculated for the analyzed system configurations as follows: $A_{conf} = \frac{TP+TN}{TP+TN+FP+FN}$, $P_{conf} = \frac{TP}{TP+FP}$, and $R_{conf} = \frac{TP}{TP+FN}$, respectively. P_{conf} and R_{conf} are undefined when their denominators are equal to zero. Accuracy, precision, and recall (i.e., A , P , and R , respectively) are also derived for the considered applications by averaging the results calculated for each configuration (M_{conf}), i.e.,

$$M = \begin{cases} \frac{\sum_{conf} M_{conf}}{\#conf} & \text{if } M_{conf} \neq \text{undefined } \forall \text{ conf} \\ \text{undefined} & \text{otherwise} \end{cases}$$

where $M = \{A, P, R\}$ and $\#conf = 12$, i.e., the number of analyzed system configurations in our experimentation. F1 score is defined as the harmonic mean of precision and recall, i.e., $F1 = 2 * \frac{P * R}{P + R}$, and it is not computed in case P or R are undefined.

Table 10 reports accuracy (A), precision (P), recall (R), and F1 score (F1) of the proposed detection rules. It is worth remarking that our calculations leverage the variation in system workload specification inputs since the ground truth (by construction) is an over approximation and TP, TN, FP, and FN values keep into account how these inputs affect the detection. The average accuracy (across the five systems) is well above 90% for all antipatterns. The average precision is also above 90%, the lowest value (i.e., 79%) is observed in OpenMRS for the Blob antipattern. The average recall is mostly larger than 90% except for WCS that shows a lower value (62%); the

TABLE 11

Analysis of How *Load Test Definition* Inputs Affect the Detection. Number of EP and Blob Antipatterns Detected Using JPAD in Two Components (i.e., *NodeManager* and *DataNode*) and Seven Commits of Hadoop. EP and Blob Instances are Reported only, JPAD does not Detect Other Antipatterns. JPAD Takes 0.31 Seconds on Average for the Detection of Antipatterns in These Configurations.

(a) YARN-4307. Commits *e914220* and *7af5d6b* introduce the issue, commit *308d63f* solves it [23].

Config.	<i>NodeManager</i>						<i>DataNode</i>					
	<i>e914220</i>		<i>7af5d6b</i>		<i>308d63f</i>		<i>e914220</i>		<i>7af5d6b</i>		<i>308d63f</i>	
	EP	Blob	EP	Blob	EP	Blob	EP	Blob	EP	Blob	EP	Blob
1-3	1	0	1	0	0	0	1	0	1	0	1	0
1-6	0	0	0	0	0	0	0	0	0	0	0	0
1-12	0	0	0	0	0	0	1	0	1	0	1	0
10-3	1	0	1	0	0	0	1	1	1	0	1	0
10-6	1	0	1	0	0	0	1	1	1	1	1	1
10-12	1	0	1	0	0	0	1	1	1	1	1	1
100-3	1	0	1	0	0	1	1	1	1	1	1	1
100-6	1	0	1	0	0	1	1	1	1	1	1	1
100-12	1	0	1	0	0	0	1	1	1	1	1	0
500-3	1	0	1	0	1	0	1	1	1	1	1	0
500-6	1	0	1	0	1	0	1	1	1	1	1	0
500-12	1	0	0	0	0	0	1	1	1	1	1	0
1000-3	1	0	1	0	1	0	1	0	1	0	1	0
1000-6	1	0	1	0	1	0	1	0	1	0	1	0
1000-12	1	0	1	0	0	0	1	0	1	0	1	0

(b) YARN-7102. Commit *528b809* introduces the issue, commit *ff8378e* solves it [23].

Config.	<i>NodeManager</i>				<i>DataNode</i>			
	<i>528b809</i>		<i>ff8378e</i>		<i>528b809</i>		<i>ff8378e</i>	
	EP	Blob	EP	Blob	EP	Blob	EP	Blob
1-3	1	0	0	0	1	0	0	0
1-6	0	0	0	0	0	0	0	0
1-12	0	0	0	0	0	0	0	0
10-3	1	0	0	0	1	1	0	0
10-6	1	0	0	0	1	1	1	0
10-12	1	0	0	1	1	1	1	0
100-3	1	0	0	0	1	1	1	0
100-6	1	0	0	1	1	1	1	0
100-12	0	0	0	0	1	1	1	0
500-3	1	0	1	0	1	1	1	1
500-6	1	0	0	0	1	1	1	0
500-12	0	0	0	0	1	1	1	0
1000-3	1	0	0	1	1	0	1	0
1000-6	1	0	1	0	1	0	1	0
1000-12	0	0	0	0	1	0	1	0

(c) HDFS-12754. Commit *decf8a6* introduces the issue, commit *738d1a2* solves it [23].

Config.	<i>NodeManager</i>				<i>DataNode</i>			
	<i>decf8a6</i>		<i>738d1a2</i>		<i>decf8a6</i>		<i>738d1a2</i>	
	EP	Blob	EP	Blob	EP	Blob	EP	Blob
1-3	0	0	0	0	0	0	0	0
1-6	0	0	0	0	0	0	0	0
1-12	0	0	0	0	0	0	0	0
10-3	1	0	0	0	1	1	0	0
10-6	1	0	1	0	1	1	1	0
10-12	1	0	0	0	1	1	1	0
100-3	1	0	0	0	1	1	1	0
100-6	1	0	0	0	1	1	1	0
100-12	0	0	0	0	1	1	1	0
500-3	1	0	1	0	1	1	1	1
500-6	1	0	0	0	1	1	1	0
500-12	1	0	0	0	1	1	1	0
1000-3	1	0	1	0	1	0	1	0
1000-6	1	0	1	0	1	0	1	0
1000-12	1	0	0	0	1	0	1	0

lowest recall (58%) is observed in TrainTicket for WCS. The F1 score shows the lowest value (i.e., 88%) for Blob in OpenMRS as a reflection of the previous result on precision, even if average values are larger than 90% for all antipatterns.

RQ₁: efficiency and accuracy

The proposed detection rules efficiently and accurately capture performance issues of medium- and large-sized systems. CTH, EP, and Blob are the antipattern types that occur in all the analyzed systems. EP shows the largest number of instances across all the system configurations. WCS is the antipattern type with the smallest number of instances. JPAD efficiently detects the instances of the presented software performance antipatterns, in fact the system configurations are analyzed, on average, in less than a minute. TeaStore shows a longer detection time, but in the worst case it is less than 5 minutes.

The F1 score, derived from precision and recall metrics, is larger than 85% in all the considered cases, denoting accurate detection rules.

4.4.2 Comparison With State-of-the-Art Approaches

The goal of this section is to investigate if JPAD is able to detect a variation on the number of antipattern instances when comparing the code commits that are known from [23] to induce and fix real-life performance issue, respectively.

Table 11 reports the performance antipattern instances that have been found in Hadoop across 7 different code commits. JPAD takes 0.31 seconds on average to analyze these configurations, and it detects EP and Blob antipatterns only, other antipatterns are not captured.

Table 11 focuses on the issue identified by YARN-4307 that is not predicted by state-of-the-art approaches [23], [45]. Interestingly, we can notice that the selected software component(s) show a different number of antipattern instances. The column reporting the fixing of the issue (i.e., *308d63f*) always shows an equal or lower number of antipattern instances in all the cases for both the considered *NodeManager* and *DataNode* components. For the *NodeManager* we can notice that there are some configurations (e.g., 100-3

TABLE 12

Number of Detected Antipatterns Using JPAD in Two Commits of Cassandra. Commit 88d2ac4 Introduces the CASSANDRA-13794 Issue, Commit f93e6e3 Solves it [23]. JPAD Takes 1 Second on Average for the Detection of Antipatterns in These Configurations.

Config.	88d2ac4				f93e6e3			
	CTH	WCS	Blob	EDA	CTH	WCS	Blob	EDA
1-3	0	0	0	0	0	0	0	0
1-6	0	0	0	0	0	0	0	0
1-12	0	0	0	0	0	0	0	0
10-3	1	0	2	0	1	0	2	0
10-6	1	0	2	0	1	0	2	0
10-12	1	0	2	0	1	0	2	0
100-3	2	1	3	2	2	1	3	1
100-6	2	1	4	2	2	1	3	1
100-12	2	1	5	2	2	1	4	2
500-3	2	1	3	1	2	1	3	1
500-6	3	1	5	2	2	1	3	1
500-12	3	1	4	2	2	1	3	1
1000-3	3	2	4	2	2	1	3	1
1000-6	3	1	4	2	2	1	4	1
1000-12	3	1	4	2	2	1	4	2

and 100-6) where EP instances are not detected in the commit fixing the issue, at the cost of emerging Blob antipattern instances. There are some further configurations (e.g., 500-3 and 500-6) for the *NodeManager* where we can notice no variation for the EP, one instance is detected in all commits. Overall, commits inducing the issue (*e914220* and *7af5d6b*) show a total of 13 and 12 EP instances across all the analyzed configurations, respectively. The commit fixing the issue manifests less antipattern instances, i.e., 4 EP and 2 Blob instances. The *DataNode* component instead has a different number of EP and Blob instances. Specifically, JPAD detects a total of 14 EP instances summing up all configurations independently of the commit. 9 and 8 Blob instances are detected for commits inducing the issue (i.e., *e914220* and *7af5d6b*, respectively), and 4 Blob instances are found in the commit solving the issue.

Table 11 b reports the results for the YARN-7102 issue. The commit fixing the issue (i.e., *ff8378e*) shows a behavior similar to the previous case, i.e., there are some configurations (i.e., 10-12, 100-6, 1000-3) of *NodeManager* for which EP is solved at the cost of a new Blob instance. EP instances in the *DataNode* rarely change (i.e., only for 1-3 and 10-3), whereas Blob instances are reduced in eight configurations (e.g., 10-3 and 10-6). Similarly to the previous issue, the *DataNode* component shows more variations for Blob instances than for EP instances. Summing up all the analyzed configurations there are 13 EP instances and 9 Blob instances for the code commit inducing the issue (*528b809*), against 11 EP instances and 1 Blob instance for the code commit fixing the issue (*ff8378e*).

Table 11 c presents the results for the HDFS-12754 issue, and also here the number of antipattern instances is equal or lower when considering the commit fixing the issue, i.e., *738d1a2*. Differently from previous cases, if EP instances are not detected then Blob instances do not arise in the *NodeManager* component. Overall, 11 and 4 EP instances across all configurations are observed for *decf8a6* and *738d1a2*, respectively. For the *DataNode* component, we get 1 Blob instance and 11 EP instances in the commit solving the issue, whereas in the commit inducing the issue we found 12 EP instances and 9 Blob instances. Summarizing, JPAD

TABLE 13

Detection Capability of JPAD and State-of-the-Art Tools w.r.t. Hadoop and Cassandra Real-Life Performance Issues. Selected Issues, Systems Commits, and Detection Results of PerfJIT and Perphecy are Extracted from [23].

Issue fixing commit	Issue inducing commit	PADprof [53]	PerfJIT [23]	Perphecy [45]	JPAD
308d63f	e914220	NO	NO	NO	YES
	7af5d6b	NO	NO	NO	YES
ff8378e	528b809	NO	NO	YES	YES
738d1a2	decf8a6	NO	NO	YES	YES
f93e6e3	88d2ac4	NO	NO	YES	YES

detects the variation across different commits when calculating the total number of detected antipattern instances. Overall, we can notice that detected instances significantly decreases when comparing code commits which induce and fix issues. Besides, software components impact on such influence, the *NodeManager* shows, on average, less EP and Blob instances when issues are fixed. Instead, for the *DataNode*, only Blob instances are observed to reduce after fixing the issue, the number of EP instances slightly varies. This is due to the nature of the analyzed issues, in fact both YARN-7102 and HDFS-12754 are indicated in [23] as complicated performance issues (like deadlock), and EP captures that there is a large number of blocked threads (i.e., a symptom of a deadlock) leading to long execution time (see Table 1).

Table 12 reports the performance antipattern instances that are found in Cassandra across two different code commits. JPAD takes 1 second on average to analyze these configurations, and it detects CTH, WCS, Blob, and EDA antipatterns, other antipatterns are not reported since no instances are detected. Between the two commits we can notice that all detected antipatterns show some decrease in their numbers when considering the commit fixing the issue, i.e., *f93e6e3*. Blob is the antipattern showing a larger number of instances, in fact commit *88d2ac4* shows 42 instances summing up all configurations, whereas commit *f93e6e3* includes 36 instances. About WCS, we can notice that there is one configuration only (i.e., 1000-3) showing a decrease of antipattern instances, no major variation is observed for this specific antipattern. Both CTH and EDA show a considerable variation; looking at the total number of detected instances across all analyzed configurations of *88d2ac4* and *f93e6e3* commits, we get 29 and 21 (17 and 11) CTH (EDA) instances, respectively. Hence, JPAD effectively detects a remarkable difference across the analyzed code commits (inducing and fixing real-life performance issues).

To answer RQ_2 , Table 13 summarizes results of comparing JPAD with state-of-the-art approaches. The last four columns of this table indicate if the specified tool can detect the considered performance issue. The column named *PADprof* refers to the framework presented in [53] which we test providing problematic snapshots (i.e., commits inducing the issue) and comparison snapshots (i.e., commits fixing the issue). All the analyzed snapshots show that no antipatterns are detected. Results for PerfJIT [23] and Perphecy [45] are instead extracted from [23] when investigating the detection of real-life performance issues.

RQ₂: comparison with state-of-the-art

JPAD overcomes state-of-the-art approaches [23], [45], [53] in the detection of some real-life performance issues. The proposed detection rules effectively capture complex performance problems that are not recognized by available tools. This consolidates the adoption of our framework as an alternative approach to support software engineers in understanding performance issues in Java applications.

4.4.3 Implication of Applying Antipattern-Based Refactorings

To answer *RQ₃*, we refactor OpenMRS (i.e., the largest application among those considered in Table 4) to understand if the detection information provided by JPAD can support software engineers in solving performance issues. The selection of OpenMRS as target system for investigating the refactoring is also motivated by a recent paper [46] that analyzes the same system to locate performance regression root causes. Specifically, we focus on the `OwaFilter` method that JPAD detects as a Blob-Controller instance. Such a method is responsible for filtering the requests directed to protected endpoints, i.e., access is granted for authenticated requests only. Due to the modular nature of the application, requests come from different modules and the filter must check the URL of all incoming requests before granting access to authenticated users and forward their requests.

Listing 1 reports a code excerpt of the `OwaFilter`. We can notice that there are several requests to be managed, for example:

- `getRequestURL()` (see line 9),
- `getServletPath()` (see line 12),
- `getAdministrationService()` (see line 14)

to mention a few. There is indeed a match with the textual description of the antipattern (see Table 1) indicating that the Blob-controller occurs in case of a single class performing all the work of an application.

As specified in the literature [26], when solving a Blob-Controller antipattern, the refactoring consists of moving computation from the affected instance to a different one.

We delegate the verification of URLs to a centralized authentication system that forwards requests to the correct endpoint after the authentication process is completed. This way, the `OwaFilter` method must only check that users are authenticated. After refactoring OpenMRS, we evaluate its performance under all loads and compare the obtained results with those observed from the original OpenMRS version. To quantify performance improvements, we consider these metrics of interest: (i) the number of detected software performance antipatterns, (ii) the CPU utilization, and (iii) the system response time. It is worth remarking that our focus is on showing empirical evidence on the benefit of solving antipatterns, i.e., possible performance improvements that can be derived by detecting and removing antipatterns, and this is why we do not investigate further refactoring types or solutions.

Table 14 reports the antipattern instances detected for the refactored system. Overall, compared to the original system, the number of instances mostly decreases, see the *Diff.* column where negative numbers indicates that the number of instances is decreased after the refactoring. For example, with 75 clients and 3 minutes of load tests running, JPAD detects 1 CTH, 2 EP, 1 WCS, 3 Blob, and 1 EDA instances in the OpenMRS original system. When applying the antipattern-based refactoring, we remove 1 instance of CTH, EP, WCS, and EDA, and 2 Blob instances. A similar improvement (see -2 entries in Table 14) is observed for CTH and WCS when there are 25 clients and the load test runs for 3 minutes. However, it is worth noting that some configurations (i.e., 25-6, 25-12, and 100-3) show more antipattern instances in the refactored system. For instance, for the configuration 25-3, the number of Blob instances increases in the refactored case. After further investigation, we find that interestingly this is due to the introduction of a *Blob-DataContainer* instance while solving the *Blob-Controller* antipattern. However, this is the only case for which Blob instances increase. Generally, the number of Blob instances is constant and decreases with 75 or 100 clients and 3 minutes of load tests running. The increment of antipattern instances (see +1 entries in Table 14) is observed for EP in two configurations and for CTH, WCS, and Blob in only one configuration. The number of EDA instances does not increase in any of the considered configurations.

Listing 1. Code excerpt of the `OwaFilter` method detected by JPAD as Blob-Controller antipattern instance.

```

1 public void doFilter(final ServletRequest req, final ServletResponse res, final FilterChain chain) throws IOException,
   ServletException {
2     final HttpServletRequest request = (HttpServletRequest)req;
3     String owaBasePath = Context.getAdministrationService().getGlobalProperty("owa.appBaseUrl", "/owa");
4     if (StringUtils.isBlank((CharSequence)owaBasePath)) {
5         owaBasePath = "/owa";
6     }
7     String requestURL = null;
8     if (isFullPath(owaBasePath)) {
9         requestURL = request.getRequestURL().toString();
10    }
11    else {
12        requestURL = request.getServletPath();
13    }
14    final String loginUrl = Context.getAdministrationService().getGlobalProperty("login.url", "login.htm");
15
16    ...
17    }
18 }

```

TABLE 14
Number of Detected Instances in OpenMRS After Applying Antipattern-Based Refactoring.
ToB and EST are Omitted Since no Instances have been Detected.

Config.	CTH		EP		WCS		Blob		EDA	
	Refact.	Diff.	Refact.	Diff.	Refact.	Diff.	Refact.	Diff.	Refact.	Diff.
25-3	0	-2	2	-1	0	-2	3	+1	1	-1
25-6	0	0	1	+1	0	0	1	0	0	-1
25-12	0	0	1	+1	0	-1	1	0	1	0
50-3	0	0	1	0	0	0	1	0	0	0
50-6	0	0	1	0	0	0	1	0	0	0
50-12	0	0	1	0	0	0	1	0	0	0
75-3	0	-1	1	-1	0	-1	1	-2	0	-1
75-6	0	-1	2	0	0	-1	3	0	1	0
75-12	0	0	1	0	0	0	1	0	0	0
100-3	1	+1	2	0	1	+1	2	-1	1	0
100-6	0	0	1	0	0	0	1	0	0	0
100-12	0	0	1	0	0	0	1	0	0	0

The impact of the antipattern-based refactoring is also observed on performance indices of interest, i.e., CPU utilization and system response time of the OpenMRS Java application, results are shown in Table 15. The two performance indices are reported for both the original and the refactored OpenMRS system. The *Variation* column shows the observed performance change and is computed as: $Variation = [(Original - Refactored) / Original] \cdot 100$. When $Variation > 0$, the considered index is smaller for the refactored system than for the original one, meaning that the system performance has improved. Table 15 highlights with bold entries all those for which *Variation* is larger than 10%. The relevance of antipattern-based refactoring is shown by the general enhancement of the OpenMRS application performance (up to 50.47%, observed for the system response time). The configuration with 25 clients and 3 minutes of load tests running is the only exception to this observation. In this case, the response time of the refactored system is 25% longer than the one of the original system even if the CPU utilization is lower for the refactored system. This may be due to performance issues that are generated in different (and previously not critical) parts of the system. We already observed that the 25-3 configuration introduces a *Blob-Data-Container* instance, when solving the *Blob-Controller* antipattern. This might be the reason for the longer response time. As future work, we plan to further investigate the solution of antipatterns and possible implications in generating new instances.

RQ₃: antipattern-based refactoring

Antipattern-based refactoring does not guarantee in advance neither a reduction of the total number of detected instances nor an improvement in the system performance. However, our experimentation shows that usually less antipattern instances are detected, and most performance indicators of interest improve. By refactoring OpenMRS, we find empirical evidence on the benefit of solving one antipattern: 12 less antipattern instances are detected; on average (across all the analyzed configurations), the CPU utilization is 18.90% lower and the system response time is 23.81% shorter. Maximum improvements for CPU utilization and system response time are 47.25% and 50.47%, respectively.

5 THREATS TO VALIDITY

Besides inheriting all the limitations related to the performance evaluation of Java-based applications [2], our approach exhibits the following main (construct, conclusion, internal, and external) threats to validity [75].

Construct Threats. This type of threat is observed when metrics deviate from the focus of the investigation. To smooth it, we provide a quantitative evaluation of the approach motivated by the research questions. We show that (i) detection rules work on real-world case studies, (ii) real-life performance issues are captured, and (iii) solving one antipattern instance improves the system performance.

Conclusion Threats. A threat of this type is related to the reliability of collected measures. To smooth these threats, we run all experiments on the same machine. Moreover, the profiling of the Java applications under analysis is delegated to the YourKit Java Profiler, a well-assessed and widely-used tool for this scope [76].

Internal Threats. We thoroughly test JPAD to spot errors in its implementation. For each experiment (whose setup can be easily changed by users), when an antipattern instance is detected, we verify if thresholds are violated. We recall that JPAD is publicly available [28] for inspection and to replicate experiments of this paper.

External Threats. We are aware that findings from our experiments may not transfer to different Java applications. To increase the external validity, we select software systems from different domains whose class number ranges from 68 to more than 1k. We also inspect code commits related to four real-life performance issues that are considered rather complex to be predicted by state-of-the-art approaches [23]. This way, we evaluate our approach against diverse applications so that our results may generalize to other case studies.

6 DISCUSSION

In this section, we discuss limitations of our approach that we consider as open issues paving the way for future research investigations.

Soundness and Completeness. Our approach currently detects seven software performance antipatterns experimented on five Java applications belonging to different domains, and nine specific commits of two further subject systems used in prior research to extract performance

TABLE 15
Performance Variation (%) Obtained by Applying the Antipattern-Based Refactoring to OpenMRS

Config.	Utilization			Response Time		
	Original [%]	Refactored [%]	Variation [%]	Original [s]	Refactored [s]	Variation [%]
25-3	27.89	15.00	46.22	3.913	4.874	-24.56
25-6	23.03	12.38	46.24	8.706	8.075	7.25
25-12	20.38	10.75	47.25	18.295	17.299	5.45
50-3	55.37	53.54	3.31	5.919	4.111	30.55
50-6	57.22	52.77	7.78	13.743	9.042	34.20
50-12	56.23	51.26	8.84	30.180	21.446	28.94
75-3	63.29	62.36	1.47	5.545	3.138	43.42
75-6	66.84	60.30	9.78	14.839	7.350	50.47
75-12	66.96	56.80	15.17	32.977	17.902	45.71
100-3	55.92	47.21	15.58	3.262	3.038	6.88
100-6	57.37	49.23	14.19	11.546	6.747	41.57
100-12	57.49	51.20	10.94	28.883	24.292	15.89

data [77]. Even if we demonstrate that our approach is able to recognize some performance issues that are not detected by other approaches in the literature (see Table 13), soundness and completeness are not guaranteed. To partially address this issue, a preliminary investigation is conducted experimenting (i) a set of four commits known to fix performance issues (see Tables 11 and 12) and (ii) an antipattern-based refactoring along with the consequent performance variation on utilization and response time indices (see Table 15). As future work we plan to strengthen this investigation involving practitioners in the evaluation of JPAD.

Antipattern Specification. Detection algorithms reflect our interpretation of the textual description of software performance antipatterns provided in [26]. We are aware that further interpretations can be provided by different stakeholders (e.g., practitioners), and we leave as part of our future work the possibility of customizing detection rules and to provide a flexible framework that reflects multiple interpretations. More in general, we plan to introduce a domain-specific language for software performance antipatterns as support for users that may define their own detection rules. This way, we aim to strengthen the specification of antipatterns and to collect the experience of different stakeholders, possibly even discovering new antipatterns.

Profiling Overhead. The performance monitoring of Java applications is known to generate overhead [54], a comparison of different profiling tools and their overhead is presented in [78], [79]. In this paper, we use YourKit since both academia [53], [57] and industrial partners, such as Apple and Google, employ it as support for evaluating the performance of industrial and real-world applications. To partially cope with the overhead introduced by YourKit, all our detection algorithms include at least one threshold derived from offsets and average values. Offsets are independent of the absolute value of considered metrics, and they allow specifying thresholds based on values that already include the profiling overhead. This way, JPAD compares performance metrics and thresholds that are both affected by the profiling overhead. We leave as future work the investigation on the usage of other monitoring tools to compare (and possibly smooth) the profiling overhead.

Antipattern Thresholds. As argued in [80], thresholds must be set in software performance antipatterns to express performance requirements (when available), or to establish boundaries which represent the perception of different

system stakeholders. In fact, users can differently judge the importance of performance requirements, e.g., the hardware utilization may be associated to monetary costs and more relevant for system administrators, whereas the execution time of a service is taken into account mainly by software developers. Therefore, JPAD provides the possibility to specify such thresholds, and this task is intentionally transferred to users that can decide which numerical values are more suitable for their purposes. We leave as future work the possibility of exploring further strategies, possibly synthesizing the need of different stakeholders.

Software Performance Testing. Test cases are often very important for an effective dynamic analysis [81]. Our approach delegates the test design to software engineers that may focus on general requirements and miss the relevant ones (from a performance-based perspective). Our experimentation highlights the importance of designing test suites (see Table 11), and demonstrates that such a selection can be guided by a preliminary analysis of the CPU utilization of software components. However, as future work we plan to investigate if approaches in the literature dealing with an efficient design of performance tests [82], [83] can be integrated in JPAD.

Antipattern-Based Refactorings. This is a very complex activity, and it is not guaranteed that the number of detected antipatterns decreases or the system performance improves. Our experimentation shows that solving one antipattern may generate other antipattern instances; more in general, antipattern instances can increase and the system response time can worsen (see Tables 14 and 15). Besides, the complexity is exacerbated by the possibly large number of detected antipattern instances, each matching with multiple code refactorings, and it is very difficult to understand which changes should be prioritized. In our previous work [33] we proposed a ranking methodology for the evaluation of architectural alternatives. We leave as part of future work to experiment ranking strategies on code refactorings thus to better investigate this aspect. Large systems may show the additional difficulty of being more sensitive to the impact of code changes, probably due to dependencies (among components) that need to be propagated when implementing refactorings. Antipattern-based solutions might be enriched with information about their effect (e.g., the involvement of dependent components) to identify which subsystems are involved in the refactoring process

and may trigger new antipatterns. To automatically fix the detected antipatterns, it is necessary that code refactorings undergo a verification process that guarantees their functional correctness.

Guidelines for Developers. When adopting JPAD in practice, we encourage developers to consider two different dimensions on the results they get as output. First, one can determine that a hotspot method shows performance issues if it is reported as a violation of some antipatterns. Second, when considering the different system configurations, the presence of the very same hotspot method (across many system configurations) contributes to the decision that such a method is indeed relevant for the performance issues under analysis. Both these two cases may indicate that such a hotspot method includes several design flaws and it indeed contributes to poor system performance.

7 CONCLUSION AND FUTURE WORK

In this paper we present JPAD, a tool-based approach to automatically detect software performance antipatterns in Java applications. The experimentation is performed on real-world Java applications from different domains, and JPAD captures four real-life performance issues that are not predicted by state-of-art approaches [23], [45], [53]. Results show the efficiency and accuracy of the proposed approach. The antipattern detection is executed on 300 configurations and we exploit such extensive experimentation to build a ground truth, thus to quantify JPAD accuracy. Overall, the accuracy is larger than 95% and the F1 score, derived from precision and recall metrics, is larger than 85% in the considered cases, leading to assess accurate detection rules. About efficiency, system configurations are analyzed, on average, in less than a minute, some configurations require more time and JPAD always takes less than 5 minutes to complete the detection of antipatterns. Besides, the number of detected antipattern instances substantially vary when experimenting software code commits known to induce and fix real-life performance issues. Antipattern-based refactoring turns out to be beneficial, the system performance improves up to 47% and 50% when measuring two specific metrics of interest, i.e., CPU utilization and system response time, respectively. JPAD points out system characteristics (e.g., number of times a method is invoked) that lead to performance issues, and its report includes quantitative information. This way, we aim to support software engineers in the task of taking decisions on which methods require more attention than others from a performance-based perspective.

Several research directions have been identified for future research. First, we want to extend the specification of antipatterns and make them *flexible*, i.e., users can add and modify detection rules to provide their own interpretation of antipatterns, possibly by introducing a domain-specific language. Second, we plan to extend the set of analyzed systems, possibly including case studies from the industrial domain to further assess both efficiency and accuracy. Third, we plan to extend JPAD to point out possible directions for antipattern-based refactorings, but the actual implementation of code fixings is delegated to software engineers who can assure the

preservation of the business logic of applications. Moreover, the solution process is complex due to the number of detected antipatterns that may be large, as demonstrated in this paper, and it is difficult to select which antipattern to solve first. Hence, we want to investigate concurrent (or prioritized) resolution of multiple antipatterns, this may lead to inconsistencies due to conflicting solutions for which ad-hoc methodologies need to be defined.

ACKNOWLEDGMENTS

The authors would like to thank the Editor and the anonymous reviewers for their constructive comments and valuable feedback. We are also grateful to the Computing and Network Service for their support in our experiments on the U-LITE cluster at INFN, LNGS, L'Aquila, Italy.

REFERENCES

- [1] C. Hunt and B. John, *Java Performance*. Englewood Cliffs, NJ, USA: Prentice Hall Press, 2011.
- [2] A. Georges, D. Buytaert, and L. Eeckhout, "Statistically rigorous Java performance evaluation," *ACM SIGPLAN Notices*, vol. 42, no. 10, pp. 57–76, 2007.
- [3] M. Harkema, D. Quartel, B. Gijzen, and R. D. van der Mei, "Performance monitoring of Java applications," in *Proc. Int. Workshop Softw. Perform.*, 2002, pp. 114–127.
- [4] R. Rabiser, S. Guinea, M. Vierhauser, L. Baresi, and P. Grünbacher, "A comparison framework for runtime monitoring approaches," *J. Syst. Softw.*, vol. 125, pp. 309–321, 2017.
- [5] R. Calinescu, C. Ghezzi, M. Z. Kwiatkowska, and R. Mirandola, "Self-adaptive software needs quantitative verification at runtime," *Commun. ACM*, vol. 55, no. 9, pp. 69–77, 2012.
- [6] H. Zhang and S. Kim, "Monitoring software quality evolution for defects," *IEEE Softw.*, vol. 27, no. 4, pp. 58–64, Jul./Aug. 2010.
- [7] R. Capilla, M. A. Babar, and O. Pastor, "Quality requirements engineering for systems and software architecting: Methods, approaches, and tools," *Requirements Eng.*, vol. 17, no. 4, pp. 255–258, 2012.
- [8] J. L. De La Vara, K. Wnuk, R. Berntsson-Svensson, J. Sánchez, and B. Regnell, "An empirical study on the importance of quality requirements in industry," in *Proc. Int. Conf. Softw. Eng. Knowl. Eng.*, 2011, pp. 438–443.
- [9] A. Aleti, B. Buhnova, L. Grunske, A. Koziolok, and I. Meedeniya, "Software architecture optimization methods: A systematic literature review," *IEEE Trans. Softw. Eng.*, vol. 39, no. 5, pp. 658–683, May 2013.
- [10] H. Koziolok, "Performance evaluation of component-based software systems: A survey," *Perform. Eval.*, vol. 67, no. 8, pp. 634–658, 2010.
- [11] F. Brosig, S. Kounev, and K. Krogmann, "Automated extraction of palladio component models from running enterprise Java applications," in *Proc. Int. Conf. Perform. Eval. Methodol. Tools*, 2009, pp. 1–10.
- [12] S. Voneva, M. Mazkatli, J. Grohmann, and A. Koziolok, "Optimizing parametric dependencies for incremental performance model extraction," in *Proc. Eur. Conf. Softw. Archit.*, 2020, pp. 228–240.
- [13] C. Heger, A. V. Hoorn, D. Okanovic, S. Siegl, and A. Wert, "Expert-guided automatic diagnosis of performance problems in enterprise applications," in *Proc. IEEE Eur. Dependable Comput. Conf.*, 2016, pp. 185–188.
- [14] A. Van Hoorn, J. Waller, and W. Hasselbring, "Kieker: A framework for application performance monitoring and dynamic software analysis," in *Proc. Int. Conf. Perform. Eng.*, 2012, pp. 247–248.
- [15] L. Leite, C. Rocha, F. Kon, D. Milojevic, and P. Meirelles, "A survey of devops concepts and challenges," *ACM Comput. Surveys*, vol. 52, no. 6, pp. 1–35, 2019.
- [16] L. Bass, I. Weber, and L. Zhu, *DevOps: A Software Architect's Perspective*. Reading, MA, USA: Addison-Wesley Professional, 2015.
- [17] M. Hüttermann, *DevOps for Developers*. New York, NY, USA: Apress, 2012.

- [18] J. Waller, N. C. Ehmke, and W. Hasselbring, "Including performance benchmarks into continuous integration to enable DevOps," *ACM SIGSOFT Softw. Eng. Notes*, vol. 40, no. 2, pp. 1–4, 2015.
- [19] H. Schulz, D. Okanovic, A. van Hoorn, and P. Tuma, "Context-tailored workload model generation for continuous representative load testing," in *Proc. Int. Conf. Perform. Eng.*, 2021, pp. 21–32.
- [20] T. Chen, W. Shang, Z. M. Jiang, A. E. Hassan, M. N. Nasser, and P. Flora, "Detecting performance anti-patterns for applications developed using object-relational mapping," in *Proc. Int. Conf. Softw. Eng.*, 2014, pp. 1001–1012.
- [21] A. Nistor et al., "CAMEL: Detecting and fixing performance problems that have non-intrusive fixes," in *Proc. IEEE Int. Conf. Softw. Eng.*, 2015, pp. 902–912.
- [22] Z. Li, T. P. Chen, J. Yang, and W. Shang, "DLFinder: Characterizing and detecting duplicate logging code smells," in *Proc. IEEE Int. Conf. Softw. Eng.*, 2019, pp. 152–163.
- [23] J. Chen, W. Shang, and E. Shihab, "PerfJIT: Test-level just-in-time prediction for performance regression introducing commits," *IEEE Trans. on Softw. Eng.*, vol. 48, no. 5, pp. 1529–1544, May 2022.
- [24] M. Harman and P. O'Hearn, "From start-ups to scale-ups: Opportunities and open problems for static and dynamic program analysis," in *Proc. IEEE Int. Work. Conf. Source Code Anal. Manipulation*, 2018, pp. 1–23.
- [25] L. Song and S. Lu, "Statistical debugging for real-world performance problems," *ACM SIGPLAN Notices*, vol. 49, no. 10, pp. 561–578, 2014.
- [26] C. U. Smith and L. G. Williams, "More new software performance antipatterns: Even more ways to shoot yourself in the foot," in *Proc. Int. Conf. Comput. Meas. Group*, 2003, pp. 717–725.
- [27] C. U. Smith, "Software performance antipatterns in cyber-physical systems," in *Proc. Int. Conf. Perform. Eng.*, 2020, pp. 173–180.
- [28] C. Trubiani, R. Pinciroli, A. Biaggi, and F. Arcelli Fontana, "Replication package: Automated detection of software performance antipatterns in Java-based applications," 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.5878953>
- [29] A. Shokri, J. C. S. Santos, and M. Mirakhorli, "Arcode: Facilitating the use of application frameworks to implement tactics and patterns," in *Proc. IEEE Int. Conf. Softw. Archit.*, 2021, pp. 138–149.
- [30] D. Feitosa, A. Ampatzoglou, P. Avgeriou, A. Chatzigeorgiou, and E. Y. Nakagawa, "What can violations of good practices tell about the relationship between GoF patterns and run-time quality attributes?," *Inf. Softw. Technol.*, vol. 105, pp. 1–16, 2019.
- [31] G. Hecht, B. Jose-Scheidt, C. D. Figueiredo, N. Moha, and F. Khomh, "An empirical study of the impact of cloud patterns on quality of service (QoS)," in *Proc. IEEE Int. Conf. Cloud Comput. Technol. Sci.*, 2014, pp. 278–283.
- [32] M. Galster and P. Avgeriou, "Qualitative analysis of the impact of SOA patterns on quality attributes," in *Proc. IEEE Int. Conf. Qual. Softw.*, 2012, pp. 167–170.
- [33] C. Trubiani, A. Kozirolek, V. Cortellessa, and R. H. Reussner, "Guilt-based handling of software performance antipatterns in palladio architectural models," *J. Syst. Softw.*, vol. 95, pp. 141–165, 2014.
- [34] R. Calinescu, V. Cortellessa, I. Stefanakos, and C. Trubiani, "Analysis and refactoring of software systems using performance antipattern profiles," in *Proc. Int. Conf. Fundam. Approaches Softw. Eng.*, 2020, pp. 357–377.
- [35] A. Wert, J. Happe, and L. Happe, "Supporting swift reaction: Automatically uncovering performance problems by systematic experiments," in *Proc. IEEE Int. Conf. Softw. Eng.*, 2013, pp. 552–561.
- [36] B. Walter, F. A. Fontana, and V. Ferme, "Code smells and their collocations: A large-scale experiment on open-source systems," *J. Syst. Softw.*, vol. 144, pp. 1–21, 2018.
- [37] F. A. Fontana, M. V. Mäntylä, M. Zanoni, and A. Marino, "Comparing and experimenting machine learning techniques for code smell detection," *Empirical Softw. Eng.*, vol. 21, no. 3, pp. 1143–1191, 2016.
- [38] F. A. Fontana, V. Ferme, A. Marino, B. Walter, and P. Martenka, "Investigating the impact of code smells on system's quality: An empirical study on systems of different application domains," in *Proc. IEEE Int. Conf. Softw. Maintenance*, 2013, pp. 260–269.
- [39] A. Yamashita and L. Moonen, "Exploring the impact of inter-smell relations on software maintainability: An empirical study," in *Proc. IEEE Int. Conf. Softw. Eng.*, 2013, pp. 682–691.
- [40] H. Liu, Z. Ma, W. Shao, and Z. Niu, "Schedule of bad smell detection and resolution: A new way to save effort," *IEEE Trans. Softw. Eng.*, vol. 38, no. 1, pp. 220–235, Jan./Feb. 2012.
- [41] D. Hovemeyer and W. Pugh, "Finding more null pointer bugs, but not too many," in *Proc. Int. Workshop Prog. Anal. Softw. Tools Eng.*, 2007, pp. 9–14.
- [42] Q. Luo, D. Poshyvanyk, and M. Grechanik, "Mining performance regression inducing code changes in evolving software," in *Proc. IEEE Work. Conf. Mining Softw. Repositories*, 2016, pp. 25–36.
- [43] J. Chen and W. Shang, "An exploratory study of performance regression introducing code changes," in *Proc. IEEE Int. Conf. Softw. Maintenance Evol.*, 2017, pp. 341–352.
- [44] D. Alshoabi, K. Hannigan, H. Gupta, and M. W. Mkaouer, "Price: Detection of performance regression introducing code changes using static and dynamic metrics," in *Proc. Int. Symp. Search Based Softw. Eng.*, 2019, pp. 75–88.
- [45] A. B. De Oliveira, S. Fischmeister, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Perphycy: Performance regression test selection made simple but effective," in *Proc. IEEE Int. Conf. Softw. Testing, Verification Validation*, 2017, pp. 103–113.
- [46] L. Liao et al., "Locating performance regression root causes in the field operations of web-based systems: An experience report," *IEEE Trans. Softw. Eng.*, vol. 48, no. 12, pp. 4986–5006, Dec. 2022.
- [47] B. A. Tate, *Bitter Java*. Shelter Island, NY, USA: Manning Publications, 2002.
- [48] B. Dudney, S. Asbury, J. K. Krozak, and K. Wittkopf, *J2EE antipatterns*. Hoboken, NJ, USA: Wiley, 2003.
- [49] B. Tate, M. Clark, and P. Linskey, *Bitter EJB*. Shelter Island, NY, USA: Manning Publications, 2003.
- [50] H. Hallal, E. Alikacem, W. Tunney, S. Boroday, and A. Petrenko, "Antipattern-based detection of deficiencies in Java multithreaded software," in *Proc. IEEE Int. Conf. Qual. Softw.*, 2004, pp. 258–267.
- [51] P. Leitner and C. Bezemer, "An exploratory study of the state of practice of performance testing in Java-based open source projects," in *Proc. Int. Conf. Perform. Eng.*, 2017, pp. 373–384.
- [52] T. Parsons and J. Murphy, "Detecting performance antipatterns in component based enterprise systems," *J. Object Technol.*, vol. 7, no. 3, pp. 55–90, 2008.
- [53] C. Trubiani, A. Bran, A. van Hoorn, A. Avritzer, and H. Knoche, "Exploiting load testing and profiling for performance antipattern detection," *Inf. Softw. Technol.*, vol. 95, pp. 329–345, 2018.
- [54] V. Horký, J. Kotrc, P. Libic, and P. Tuma, "Analysis of overhead in dynamic Java performance monitoring," in *Proc. Int. Conf. Perform. Eng.*, 2016, pp. 275–286.
- [55] YourKit GmbH, YourKit java profiler features. Accessed: Dec. 14, 2022. [Online]. Available: <https://www.yourkit.com/features/>
- [56] YourKit GmbH, YourKit Customers. Accessed: Dec. 14, 2022. [Online]. Available: <https://www.yourkit.com/customers/>
- [57] Y. Zhao, L. Xiao, X. Wang, Z. Chen, B. Chen, and Y. Liu, "Butterfly space: An architectural approach for investigating performance issues," in *Proc. IEEE Int. Conf. Softw. Archit.*, 2020, pp. 202–213.
- [58] YourKit GmbH, Profiling overhead: How to reduce or avoid. Accessed: Dec. 14, 2022. [Online]. Available: <https://www.yourkit.com/docs/java/help/overhead.jsp>
- [59] G. Canfora, F. Martinelli, F. Mercaldo, V. Nardone, A. Santone, and C. A. Visaggio, "LEILA: Formal tool for identifying mobile malicious behaviour," *IEEE Trans. Softw. Eng.*, vol. 45, no. 12, pp. 1230–1252, Dec. 2019.
- [60] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, "Understanding and detecting real-world performance bugs," in *Proc. ACM SIGPLAN Conf. Program. Lang. Des. Implementation*, 2012, pp. 77–88.
- [61] J. Thome, L. K. Shar, D. Bianculli, and L. Briand, "An integrated approach for effective injection vulnerability analysis of web applications through security slicing and hybrid constraint solving," *IEEE Trans. Softw. Eng.*, vol. 46, no. 2, pp. 163–195, Feb. 2020.
- [62] X. Zhou et al., "Fault analysis and debugging of microservice systems: Industrial survey, benchmark system, and empirical study," *IEEE Trans. Softw. Eng.*, vol. 47, no. 2, pp. 243–260, Feb. 2021.
- [63] CloudStore synchronizes your data. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/cloudstore/cloudstore>
- [64] TeaStore: A micro-service reference and test application. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/DescartesResearch/TeaStore>

- [65] J. von Kistowski, S. Eismann, N. Schmitt, A. Bauer, J. Grohmann, and S. Kounev, "TeaStore: A micro-service reference application for benchmarking, modeling and resource management research," in *Proc. IEEE Int. Symp. Modelling, Anal., Simul. Comput. Telecommun. Syst.*, 2018, pp. 223–236.
- [66] WebGoat 8: A deliberately insecure Web Application. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/WebGoat/WebGoat>
- [67] Train Ticket: A Benchmark Microservice System. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/FudanSELab/train-ticket>
- [68] X. Zhou et al., "Latent error prediction and fault localization for microservice applications by learning from system trace logs," in *Proc. Joint Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2019, pp. 683–694.
- [69] OpenMRS: Medical Record System. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/openmrs/openmrs-core>
- [70] The Apache Software Foundation, Apache Hadoop. Accessed: Dec. 14, 2022. [Online]. Available: <https://hadoop.apache.org/>
- [71] The Apache Software Foundation, Apache Cassandra: Open source NoSQL database. Accessed: Dec. 14, 2022. [Online]. Available: https://cassandra.apache.org/_/index.html
- [72] TestDFSIO. Accessed: Dec. 14, 2022. [Online]. Available: <https://github.com/apache/hadoop/blob/master/hadoop-mapreduce-project/hadoop-mapreduce-client/hadoop-mapreduce-client-jobclient/src/test/java/org/apache/hadoop/fs/TestDFSIO.java>
- [73] The Apache Software Foundation, Cassandra Stress. Accessed: Dec. 14, 2022. [Online]. Available: https://cassandra.apache.org/doc/4.0/cassandra/tools/cassandra_stress.html
- [74] J. Heyman, C. Byström, J. Hamrén, and H. Heyman, "Locust: An open source load testing tool," Accessed: Dec. 14, 2022. [Online]. Available: <https://locust.io/>
- [75] C. Wohlin, P. Runeson, M. Hst, M. C. Ohlsson, B. Regnell, and A. Wessln, *Experimentation in Software Engineering*. Berlin, Germany: Springer, 2012.
- [76] T. Mytkowicz, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Evaluating the accuracy of Java profilers," *ACM SIGPLAN Notices*, vol. 45, no. 6, pp. 187–197, 2010.
- [77] M. D. Syer, W. Shang, Z. M. Jiang, and A. E. Hassan, "Continuous validation of performance test workloads," *Automated Softw. Eng.*, vol. 24, no. 1, pp. 189–231, 2017.
- [78] P. Su, Q. Wang, M. Chabbi, and X. Liu, "Pinpointing performance inefficiencies in Java," in *Proc. Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2019, pp. 818–829.
- [79] F. David, G. Thomas, J. Lawall, and G. Muller, "Continuously measuring critical section pressure with the free-lunch profiler," *ACM SIGPLAN Notices*, vol. 49, no. 10, pp. 291–307, 2014.
- [80] V. Cortellessa, A. Di Marco, and C. Trubiani, "An approach for modeling and detecting software performance antipatterns based on first-order logics," *Softw. Syst. Model.*, vol. 13, no. 1, pp. 391–432, 2014.
- [81] L. Mariani, F. Pastore, and M. Pezze, "Dynamic analysis for diagnosing integration faults," *IEEE Trans. Softw. Eng.*, vol. 37, no. 4, pp. 486–508, Jul./Aug. 2011.
- [82] S. Mostafa, X. Wang, and T. Xie, "Perfranker: Prioritization of performance regression tests for collection-intensive software," in *Proc. Int. Symp. Softw. Testing Anal.*, 2017, pp. 23–34.
- [83] D. G. Reichelt and S. Kühne, "Better early than never: Performance test acceleration by regression test selection," in *Proc. Int. Conf. Perform. Eng.*, 2018, pp. 127–130.
- [84] A. Ali, R. Pincirolì, F. Yan, and E. Smirni, "Batch: Machine learning inference serving on serverless platforms with adaptive batching," in *Proc. IEEE Int. Conf. High Perform. Comput., Netw., Storage Anal.*, 2020, pp. 1–15.
- [85] K. Nguyen et al., "Yak: A high-performance big-data-friendly garbage collector," in *Proc. USENIX Symp. Operating Syst. Des. Implementation*, 2016, pp. 349–365.



Catia Trubiani is an associate professor with the Gran Sasso Science Institute (GSSI), Italy. Previously she collaborated with the Karlsruhe Institute of Technology in Germany, and the Imperial College of London in UK. Her research interests include software performance modeling and analysis, software quality optimization, uncertainty propagation, model-based testing and software architectures. For more information, please visit <https://cs.gssi.it/catia.trubiani>



Riccardo Pincirolì received the MS and PhD degrees in computer engineering from Politecnico di Milano, in 2014 and 2018, respectively. He is currently a postdoc fellow in computer science with the Gran Sasso Science Institute. His research interests include stochastic modeling, performance evaluation, energy efficiency, and uncertainty propagation applied to cloud computing, data-centers, and cyber-physical systems.



Andrea Biaggi received the master's degree with the Evolution and Reverse Engineering Lab of University of Milano Bicocca, Italy, with the supervision of Prof. Francesca Arcelli Fontana. He had a fellowship from the same Lab and currently works as a software engineer in a company in Milano, Italy.



Francesca Arcelli Fontana is full professor with the University of Milano Bicocca, Italy. Her research activity principally concerns software engineering, in particular software evolution and reverse engineering, architectural smell detection, program comprehension, software quality assessment and managing technical debt. She is the head of the Software Evolution and Reverse Engineering Lab with the University of Milano Bicocca.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.

Open Access funding provided by 'Gran Sasso Science Institute - GSSI' within the CRUI CARE Agreement