

PERUN: Performance Version System

Tomáš Fiedor*, Jiří Pavela*, Adam Rogalewicz*, and Tomáš Vojnar*

*Brno University of Technology, Faculty of Information Technology, Brno, Czech Republic

Abstract—In this paper, we present PERUN: an open-source tool suite for profiling-based performance analysis. At its core, PERUN maintains links between project versions and the corresponding stored performance profiles, which are then leveraged for automated detection of performance changes in new project versions. The PERUN tool suite further includes multiple profilers (and is designed such that further profilers can be easily added), a performance fuzz-tester for workload generation, methods for deriving performance models, and numerous visualization methods. We demonstrate how PERUN can help developers to analyze their program performance on two case studies: detection and localization of a performance degradation and generation of inputs forcing performance issues to show up.

Supplementary materials—git repository [1], demo video and replication package [2], and a longer version of the paper [3].

Index Terms—version system, profiling, performance analysis, fuzz-testing, performance testing.

I. INTRODUCTION

Performance bugs are a common problem in software development, often encountered during new software releases. Despite such bugs need not crash the software, neglecting them can lead to losing the customers’ trust. Moreover, the need of efficient software is nowadays becoming more important than ever due to the rising scale of software and growing pressure on saving resources such as, e.g., battery power on mobile devices [4] or the huge power consumption of computing centers [5]. Finally, software performance is also a matter of security as performance bugs can be exploited to effectively kill a system by some form of a denial-of-service attack (DoS), making it run exceptionally long on some inputs.

Profiling-based analysis is arguably one of the most widely adopted techniques used when analysing software performance. Here, *profiling* means monitoring a running program and collecting performance-related data, e.g., function call counts, function run times, etc. These then need to be analysed and possibly compared with previous profiling results to identify the most costly functions and to see if some performance degradation did happen and what its root cause may be.

The described process is laborious and moreover repeating. Indeed, before every new software release (at the latest), one should profile the new version, analyze the results, compare them with past profiling, and store the results for future profiling. Hence, it is highly desirable to automate the process. However, while there exist various widely-used open-source *profilers* (e.g., GNU `gprof` or `CALLGRIND`), we are not aware of any actively-maintained open-source *solution* that would cover profiling-based performance analysis in a *complex* way. Here, by complex, we mean solutions not focusing just

on collecting various performance metrics (on time, memory, energy, etc.) but also on other needed steps (degradation detection, workload generation, etc.). To address this gap, we designed PERUN: *a performance version system*.

PERUN stores *performance profiles* for all chosen project versions and maintains links between the project versions and their profiles. These can then be used for automated detection of performance changes, including creeping ones, in new project versions. The PERUN tool suite also includes a highly configurable runtime profiler called TRACER capable of collecting various data about profiled C/C++ functions, several other profilers (and a support for easily adding profilers on further metrics and/or languages), a performance fuzz-tester for generating suitable inputs for the program under performance analysis (i.e., the so-called *workloads*), methods for deriving performance models from saved profiles, as well as numerous degradation detection and visualization methods.

Below, we first describe the architecture of PERUN and the methods implemented in it, and then discuss various ways in which it can be used. We illustrate the usefulness of PERUN on two concrete tasks. First, we show how it can be used to automatically detect and locate a degradation between two software versions, namely a recent performance degradation introduced in the 3.11.0a7 version of the CPYTHON interpreter [6]. Second, we present how PERUN can be used to automatically generate inputs that force performance issues to manifest, e.g., to test different hashing algorithms or vulnerability against ReDoS (regex denial of service) attacks.

II. ARCHITECTURE OF PERUN

PERUN consists of a tool suite and a wrapper over *Version Control Systems (VCS)*, such as *git*, that keeps track of performance profiles for different project versions. Figure 1 shows the architecture of PERUN. It builds on a generic JSON-based format for storing performance profiles and other artifacts (e.g., performance models). To implement a flexible and extensible automation of its various supported tasks, Perun uses so-called *runners* that manage *jobs*: sequences of calls of individual tools from its tool suite, i.e., data collectors, methods deriving performance models, detection methods, or tools used for interpretation of the results (e.g., various visualisation methods). Jobs may produce profiles or other results (such as performance models). PERUN’s tools share a simple API, and developers can therefore easily register their own tools. Moreover most PERUN’s tools are highly configurable, allowing their fine-tuning based on domain knowledge of the profiled program. PERUN also contains workload generators:

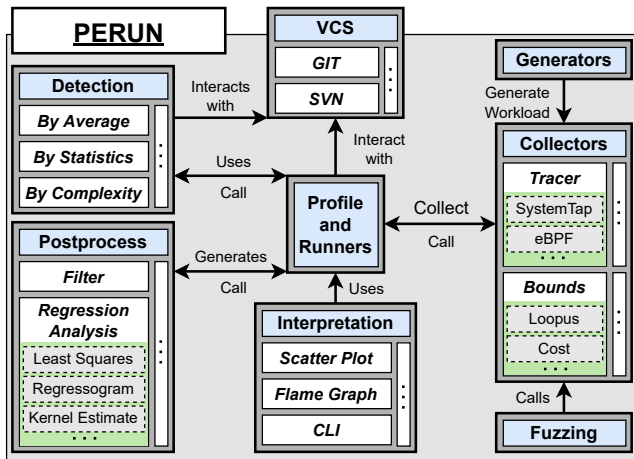


Fig. 1: An overview of the PERUN's architecture.

a fuzz-tester focused on generating time-consuming workloads or automatic generators of gradually scaled workloads (e.g., longer and longer files) that are fed to the profiled program.

PERUN's current profilers mainly focus on profiling the runtime of functions or programs. The main runtime profiler is TRACER (described in more detail below). PERUN further includes a lightweight runtime profiler with minimum overhead that is simply a wrapper over the Unix `time` utility as well as a memory profiler, which is, however, so far experimental and not too scalable. Other profilers (such as CALLGRIND) can be easily added to PERUN in a similar fashion as `time`.

Below, we describe selected tools of PERUN in more detail.

a) *Tracer*: TRACER leverages multiple instrumentation frameworks to monitor programs and measure the runtime of C/C++ functions while also keeping track of the call hierarchy (i.e., the *caller-callee* relations). TRACER currently supports the `SYSTEMTAP` [7] and the `eBPF` [8], [9] frameworks, and can be extended by other frameworks. `SYSTEMTAP` injects instrumentation probes into the profiled program through an on-demand-compiled kernel module; while `eBPF` applies a more dynamic approach based on a lightweight in-kernel virtual machine allowing better dynamic probe runtime manipulation (e.g., probe activation/deactivation). Due to its rich parameterization, TRACER is highly configurable, e.g., by choosing the instrumentation framework, the functions to be monitored, or the frequency with which they should be monitored.

TRACER works as follows: 1) It is invoked with the target executables (including libraries). The functions to be instrumented can be specified by the user or extracted from the executables automatically. 2) For every instrumented function, TRACER assembles and compiles *handlers* that generate the performance data. 3) It instruments the target executables, invokes the profiled program and collects performance data until the program terminates or a timeout is reached. More precisely, the collected data are in the form of *time stamps* recorded at the entry and exit of each instrumented function. 4) Finally, the raw performance data are parsed into a profile.

b) *Models and Detection Techniques*: PERUN allows one to summarize profiles into *performance models* mostly having the form of functions mapping input sequence numbers to

the expected runtime. By suitably constructing the workload, the users can then get, e.g., the classical *complexity models* (this is, *Big-O classes* such as linear, quadratic, etc.) showing the dependence of runtime on the size of the input (by the workload exercising the program/function for larger and larger inputs), or to see how the performance is changing over time (by providing inputs as they would arrive in time). Such workloads can be generated by workload generators coming with PERUN or implemented by the users. The derived models are then used in PERUN's performance degradation detection methods and/or visualization methods.

PERUN currently supports four kinds of methods of deriving performance models: the least-square regression line, regressogram, moving average, and kernel regression line (for their details, see, e.g., [10], [11] or the PERUN's documentation). For each function/profile we can, naturally, have multiple models.

PERUN implements multiple methods of *detecting changes in the performance*, ranging from simple heuristics to more elaborate techniques, building on well-known principles from the complexity theory and statistics. Each method can be applied to each pair of profiles (or models obtained from them) for the same functions in any pair of versions of the program under analysis.

Here are three examples of the detection methods: 1) The *best model-order equality* builds models representing various asymptotic complexity classes of the involved functions (linear, quadratic, etc.), selects those that are statistically the most precise for the respective functions, and compares the orders of the models considered the most precise for different versions of a given function. 2) The *integral comparison* approach is a heuristic based on an assumption that the area under the performance models viewed as mathematical functions should stay approximately the same: the technique thus computes and compares the integrals of the most precise models. 3) The *exclusive-time outliers* checks whether the function's exclusive runtime (not including the runtime of the functions it calls) changes so much that it becomes an outlier w.r.t. other functions based on a selected statistic (the z-score or interquartile range).

c) *Performance Fuzzing*: PERUN's tool suite includes PERUN-FUZZ: an automated fuzz-testing generator of workloads focused on triggering performance issues. Fuzz-testing is an approach of automated testing that feeds a program with randomised or mutated inputs to force unexpected/untested behaviour. However, most of the existing approaches focus only on manifesting program crashes, memory leaks, or failed assertions; not on manifesting performance issues.

Well-established fuzzers, such as `afl` [12], often build on mutation rules that work on the level of bits in the binary encoding of program inputs, for which, in our experience, it may be difficult to derive inputs that could make performance issues to surface. For PERUN-FUZZ, we developed mutation rules inspired by various known performance issues such as the following ones developed specifically for text-file inputs: "*double the size of a line*" (inspired by the issues of the

gedit editor with long lines), “repeat a random word on a line” (to force issues, e.g., in hash tables), “sort words or numbers on a line” (to force issues in algorithms expecting randomly sorted inputs), or “prepend whitespaces” (inspired by the issues of Stack Overflow [13] that froze this web site).

Whenever a new workload is derived, the `gcov` tool that measures code coverage (and counts also multiple visits of a location) is used for a quick check whether the mutation led to a coverage increased by more than some threshold. If so, PERUN is used to profile the program and to check whether some performance degradation did indeed happen and where.

III. USAGE OF PERUN

The most common usage of PERUN is to analyze newly released program versions. The usual scenario goes as follows: 1) The user makes a set of code changes and commits them into the repository of the program of interest. 2) The commit triggers a new build of the program, and TRACER is used over the resulting binary to collect new performance profiles of the involved functions. 3) New performance models of the functions are derived and stored. 4) Profiles and models of the functions from the previous version are retrieved from the persistent storage of PERUN. 5) For each pair of corresponding previous and new profiles (or performance models, depending on the detection method used), a selected performance degradation detection method is run. 6) The method returns a set of located performance changes with their *severity* (i.e., how good or bad the change is), *location* (i.e. where the change has happened), and *confidence* (i.e., how likely the change is real).

Another common usage scenario for PERUN is in *debugging*. In its case, one can use the rich parameterization of TRACER to switch between quick and thorough profiling while, e.g., trying to locate some suspected bottlenecks or opportunities for optimisation. To help this process, some of the workload generators (generating, e.g., random files) or the performance fuzzer PERUN-FUZZ may be used. Further, to help understanding what happens with performance in a given program, various well-established *visualizations* of profiles and models implemented in PERUN can be used. These include, e.g., scatter plots, bar plots, flamegraphs, or flowgraphs (for some samples, see [2]).

PERUN can be used through its rich command-line interface, but it can also be easily integrated with *git hooks* or other *continuous integration* solutions so that checks for performance degradation are triggered automatically upon a commit. Moreover, adapters based on OSLC [14] and the UNITE and UNIC tools [15], [16] are available for using PERUN as a *web service* and for its integration with the *Eclipse IDE*.

IV. CASE STUDIES

We now demonstrate how PERUN and its tool suite can aid developers with two common performance analysis tasks.

a) Finding Performance Degradations: First, we will focus on showing how PERUN can help with identifying a performance degradation and its root cause. In particular, we will illustrate this on the CPYTHON project where a recent release

of the version 3.11.0a7 introduced a performance degradation compared to the version 3.10.4. The problem appeared in the `ctypes` module¹. The developer who discovered the bug has already proposed a fix; and the fix has been merged into the upstream. In the following, we will replicate the issue and try to detect it using PERUN.

CPYTHON comes with a benchmarking suite `pyperformance` that is used to evaluate the performance of various modules. However, this particular issue has not been detected by these benchmarks. Since the bug was reported and fixed, a new benchmark targeting the `ctypes` module has been implemented and is currently pending as a pull request at the time of writing of this paper. Although the new benchmark now retrospectively detects the presence of the `ctypes` performance degradation, finding the exact source of the issue can still pose a challenge as it requires a manual inspection of all relevant code changes since the last version.

PERUN and its version-sensitive performance analysis can nicely complement the benchmarking approach by automatically flagging the offending function(s), thus aiding the developers in promptly hotfixing the issue. In our experiment, we did the following steps, simulating what the developers could have done with PERUN when going between the mentioned versions: 1) We collected a *baseline profile* of the `ctypes` benchmark, linked to the 3.10.4 release (normally, it would already be stored in PERUN). 2) For the CPYTHON version 3.11.0a7, we ran `perun collect` with the same profiling configuration of the `ctypes` benchmark as for the 3.10.4 release and obtained a *target profile* for the release 3.11.0a7. 3) Next, we ran `perun check profiles` to perform an automated detection of performance changes between the *baseline* and *target* profiles. We chose the *exclusive time outliers* method for detecting the degradation. 4) The check successfully identified the two functions that were responsible for the performance degradation as shown in the below table:

Location	Result	Δ [ms]	Δ [%]
<code>._ctypes_init_fielddesc</code>	NotInBaseline	77.95	5.23
<code>._ctypes_get_fielddesc</code>	SevereDegradation	52.9	3.55
<code>._ctypes_callproc</code>	Degradation	2.84	0.19
...			
<code>._ctypes.cpython-311</code>	TotalDegradation	136.92	9.19

The Δ [ms] (Δ [%]) is the absolute (relative) exclusive time change w.r.t. the total duration of the program or library which contains the function. Note that the total degradation we measured is comparable to the degradation of around 8% reported in the original issue. 5) We created a hotfix in a new VCS branch `hotfix-issue-92357` and repeated the profiling to obtain a new *target profile* for the branch `hotfix-issue-92357`. 6) Finally, we checked that the fix was successful: we ran `perun check profiles` using the original *baseline* and the new *target* profiles. This time, the check report shows that both of the previously degraded functions have comparable performance with the release 3.10.4 (see the table below). Note that the `ctypes` module still reports a minor degradation, corresponding to the `pyperformance` benchmarking results.

¹<https://github.com/python/cpython/issues/92356>

Location	Result	Δ [ms]	Δ [%]
<code>_ctypes_callproc</code>	SevereDegradation	9.7	0.7
...			
<code>_ctypes_get_fielddesc</code>	MaybeDegradation	0.89	0.06
<code>_ctypes_init_fielddesc</code>	NotInBaseline	0.02	0.00
<code>_ctypes.cpython-311</code>	TotalDegradation	23.45	1.70

In the usual profiling scenario, should the developer wish to continue with optimizations, the next most severe degradation would be targeted next. The *exclusive time outliers* method aids the developer with this effort and reports a new severe degradation, relative to the other found degradations: `_ctypes_callproc`. To define a stopping point for this workflow, a *cut-off* threshold for Δ can be set, so that no degradations or improvements below this threshold are reported.

For detection of *creeping performance degradations*, one can utilize the same approach as above with a slight modification only. Namely, the *baseline* profile should be selected such that it corresponds to an older project version or release.

b) *Inferring Interesting Workloads*: In our second case study, we will demonstrate how to automatically generate workloads suitable for performance analysis. In this study, we will use minimalistic single-purpose projects. We believe that such programs are enough at least if our goal is to find workloads for testing specific code patterns.

In particular, we will generate workloads to (1) identify potentially expensive regex matching and (2) to compare the efficiency of various hash functions used to implement a hash table. For each experiment, we list the size of each workload denoted as w_i (with the starting, manually prepared workload denoted as *seed*), the runtime of the program under test on the given workload, the slowdown w.r.t. the runtime on the seed, and the number of mutation rules that led to the workload.

First, we focus on regex matching that can be very expensive for some matchers and some regexes, potentially admitting a ReDoS (Regex Denial of Service) attack [17], [13], [18]. These were reported in various applications, e.g., the following regex used for Java classname validation [17]: $\hat{((([a-z]))+.)+[A-Z]([a-z]))+\$$. The issue here is that the developer did not escape the `.` character, which makes some matchers to perform an excessive number of backtracking steps on some non-matching words.

We tested the `std::regex_search` function with several offending regexes that can cause ReDoS attacks. Below, we present results for the above Java Classname validation regex. For other regexes and results, see [2].

Java Class validation regex [17]				
input	size [B]	time [s]	slowdown	used rules
<i>seed</i>	19	0.005	-	-
w_2	19	0.016	3.2	9
w_3	36	1.587	317.4	6
w_5	78	timeout (13h)	∞	5

Clearly, at least for the considered regexes, we can find problematic workloads with only a few applications of our mutation rules. Indeed, some generated inputs were only 4x bigger than the seed, yet leading to a timeout of thirteen hours.

Finally, we tested a simple word frequency counter that uses a hash table. We compared its performance using two different hash functions: (1) the hash function used in the Java

1.1 string library, which examined only 8–9 evenly spaced characters, and, (2) the DJB hash function [19], one of the most efficient hash functions. The results show that with increasing sizes of the workload, the DJB hash function does indeed perform in a much more stable way compared to the inefficient implementation of Java 1.1.

input	Work		Java 1.1.		DJB	
	size [kB]	time [s]	slowdown	time [s]	slowdown	
<i>seed</i>	210	0.026	-	0.013	-	
w_6	458	0.115	4.4	0.027	2.1	
w_7	979	0.187	7.19	0.043	3.3	

V. RELATED WORK

Perhaps the closest solution to PERUN is PERFREPO [20] whose idea is similar to what PERUN aims to achieve, but it covers profiling-based analysis in a less complex way (e.g., it does not include a profiler, workload generation, etc.) and it lacks a tight VCS integration (this is, it lacks built-in linkage between profiles and project VCS versions). Moreover, its development seems to be discontinued, based on the GitHub repository activity. KIEKER [21] provides a tool suite for performance data collection and results interpretation but lacks any form of project version context. VALGRIND [22], a widely-adopted profiling tool suite, lacks support for results comparison, and VCS/CI integration.

A range of individual profilers, either sampling or event-based, for heap, memory or CPU profiling exist. We can list, e.g., OPROFILE [23], PERF [24], GNU gprof [25], or gperftools [26]. Furthermore, we can mention various commercial solutions such as, e.g., the *Intel V-Tune Profiler* [27], *AMD μ Prof* [28], *Arm MAP* [29], or the *Oracle Developer Studio Performance Analyzer* [30]. Additionally, there are other performance analysis tools based on machine learning [31], code patterns [32], learning from performance bug trackers [33], automated and compositional load test generation [34], [35], checking correlation between performance testing results [36], clustering and comparing performance counters [37], or using statistical process control techniques [38]. However, these tools do either not provide as complex coverage of profiling-based performance analysis as PERUN does (being often restricted to profiling only) and/or they are not open source, often quite expensive and/or coming with some additional hardware or software requirements for unlocking their full potential.

Finally, as for performance fuzzing, the only tools of this kind we are aware of are PERFFUZZ [39] and BADGER [40]. Unlike PERUN-FUZZ, which is, however, just one part of PERUN, they use other fuzzing rules, working on the binary level as in afl (BADGER combines fuzzing with symbolic execution too). As an alternative to performance fuzzing, causal profiling [41] aims at finding particular functions to optimise (by slowing down other functions to virtually simulate the speed-up). PERFBLOWER [42] forces memory-related errors by amplifying their effects on performance. A support of such techniques could be added to PERUN in the future.

Acknowledgement. The work was supported by the project 20-07487S of the Czech Science Foundation and the Brno Ph.D. Talent Scholarship Programme.

REFERENCES

- [1] T. Fiedor and J. Pavela, “Perun: Lightweight performance version system.” [Online]. Available: <https://github.com/TFiedor/perun>
- [2] —, “Perun tool demonstration website.” [Online]. Available: <https://www.fit.vutbr.cz/research/groups/verifit/tools/perun-demo/>
- [3] T. Fiedor, J. Pavela, A. Rogalewicz, and T. Vojnar, “Perun: Performance version system,” 2022. [Online]. Available: <https://arxiv.org/abs/2207.12900>
- [4] D. Ferreira, A. Dey, and V. Kostakos, “Understanding human-smartphone concerns: A study of battery life,” in *Proc. of Pervasive Computing*, 2011.
- [5] M. Dayarathna, Y. Wen, and R. Fan, “Data center energy consumption modeling: A survey,” *IEEE Commun. Surv. Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.
- [6] “Cpython: The python programming language.” [Online]. Available: <https://github.com/python/cpython>
- [7] V. Prasad, W. Cohen, F. Egler, M. Hunt, J. Keniston, and B. Chen, “Locating system problems using dynamic instrumentation,” in *Proc. of Linux Symposium’05*, 2005.
- [8] S. McCanne and V. Jacobson, “The bsd packet filter: A new architecture for user-level packet capture,” in *Proc. of USENIX’93*, 1993.
- [9] J. Corbet, “Bpf: the universal in-kernel virtual machine.” [Online]. Available: <https://lwn.net/Articles/599755/>
- [10] J. L. Devore, *Probability and Statistics for Engineering and the Sciences*, 8th ed. Brooks/Cole, 2011, ISBN-13: 978-0-538-73352-6.
- [11] K. Takezawa, *Introduction to Nonparametric Regression*, 1st ed. Wiley, 2006, ISBN-13: 978-0471745839.
- [12] M. Zalewski, “American Fuzzy Lop (AFL) Fuzzer—Technical Details,” 2018. [Online]. Available: http://lcamtuf.coredump.cx/afl/technical_details.txt
- [13] “Malformed stack overflow post chokes regex, crashes site – adtmag,” online. [Online]. Available: <https://adtmag.com/Blogs/Dev-Watch/2016/07/stack-overflow-crash.aspx>
- [14] OASIS, “Open Services for Lifecycle Collaboration,” <https://open-services.net/>, 2022.
- [15] O. Vašíček, “Unite – GitLab,” 2022. [Online]. Available: <https://pajda.fit.vutbr.cz/verifit/unite>
- [16] J. Fiedor, “Unic – GitLab,” 2022. [Online]. Available: <https://pajda.fit.vutbr.cz/verifit/oslc-client>
- [17] “Owasp validation regex repository — owasp foundation,” online. [Online]. Available: https://owasp.org/www-community/OWASP_Validation_Regex_Repository
- [18] “Regular expression library,” online. [Online]. Available: https://regexlib.com/REDetails.aspx?regex_id=1757
- [19] “::: General purpose hash function algorithms - by arash partow :::” online. [Online]. Available: <http://www.partow.net/programming/hashfunctions/#DJBHashFunction>
- [20] “Meet perfrepo,” online. [Online]. Available: <https://github.com/PerfCake/PerfRepo>
- [21] A. van Hoorn, M. Rohr, W. Hasselbring, J. Waller, J. Ehlers, S. Frey, and D. Kieselhorst, “Continuous monitoring of software services: Design and application of the kieker framework,” Department of Computer Science, University of Kiel, Germany, Tech. Rep. TR-0921, 2009.
- [22] N. Nethercote and J. Seward, “Valgrind: A framework for heavyweight dynamic binary instrumentation,” in *Proc. of PLDI’07*. ACM, 2007.
- [23] “Oprofile - a system profiler for linux,” online. [Online]. Available: <https://oprofile.sourceforge.io/news/>
- [24] “Perf: Linux profiling with performance counters,” online. [Online]. Available: https://perf.wiki.kernel.org/index.php/Main_Page
- [25] “Gnu gprof,” online. [Online]. Available: <https://sourceware.org/binutils/docs/gprof/index.html>
- [26] “Gperftools,” online. [Online]. Available: <https://github.com/gperftools/gperftools>
- [27] “Intel v-tune profiler,” online. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/tools/oneapi/vtune-profiler.html>
- [28] “Amd μ prof,” online. [Online]. Available: <https://developer.amd.com/amd-uprof/>
- [29] “Arm map,” online. [Online]. Available: <https://www.arm.com/products/development-tools/server-and-hpc/forge/map>
- [30] “Oracle developer studio performance analyzer,” online. [Online]. Available: https://docs.oracle.com/cd/E77782_01/html/E77798/index.html
- [31] S. Tsakiltisidis, A. V. Miranskyy, and E. Mazzawi, “Towards automated performance bug identification in python,” *CoRR*, vol. abs/1607.08506, 2016.
- [32] A. Nistor, L. Song, D. Marinov, and S. Lu, “Toddler: Detecting performance problems via similar memory-access patterns,” in *Proc. of ICSE’13*. IEEE CS, 2013.
- [33] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, “Understanding and detecting real-world performance bugs,” in *Proc. of PLDI’12*. ACM, 2012.
- [34] P. Zhang, S. G. Elbaum, and M. B. Dwyer, “Automatic generation of load tests,” in *Proc of ASE’11*, 2011, pp. 43–52.
- [35] P. Zhang, S. Elbaum, and M. B. Dwyer, “Compositional load test generation for software pipelines,” in *Proc. of ISSTA’12*, 2012, pp. 89–99.
- [36] K. C. Foo, Z. M. Jiang, B. Adams, A. E. Hassan, Y. Zou, and P. Flora, “Mining performance regression testing repositories for automated performance analysis,” in *Proc. of QSI’10*, 2010.
- [37] W. Shang, A. E. Hassan, M. N. Nasser, and P. Flora, “Automated detection of performance regressions using regression models on clustered performance counters,” in *Proc. of ICPE’15*, 2015.
- [38] T. H. D. Nguyen, B. Adams, Z. M. Jiang, A. E. Hassan, M. N. Nasser, and P. Flora, “Automated detection of performance regressions using statistical process control techniques,” in *Proc. of ICPE’12*. ACM, 2012.
- [39] C. Lemieux, R. Padhye, K. Sen, and D. Song, “Perffuzz: Automatically generating pathological inputs,” in *Proc. of ISSTA’18*. ACM, 2018.
- [40] Y. Noller, R. Kersten, and C. S. Pasareanu, “Badger: Complexity analysis with fuzzing and symbolic execution,” in *Proc. of ISSTA’18*. ACM, 2018.
- [41] C. Curtsinger and E. D. Berger, “Coz: Finding code that counts with causal profiling,” *Communications of the ACM*, vol. 61, 2018.
- [42] L. Fang, L. Dou, and G. H. Xu, “Perfblower: Quickly detecting memory-related performance problems via amplification,” in *Proc. of ECOOP’15*, 2015.