

## ARTIFICIAL FLOOR PLAN GENERATION USING MACHINE LEARNING, CONFPROFIT: A PERFORMANCE PROFILING TESTING



### Amrendra Kumar Raushan

M.Phil, Roll No. :150112: Session : 2015-16

University Department of COMPUTER SCIENCE, B.R.A. Bihar University, Muzaffarpur, India.

E-mail:- eidamrendra@gmail.com

#### ABSTRACT

Sadly, builders often have no idea how the overall performance of a device suffers from configuration variables and the way they interact. The frequency of configuration errors inducing overall performance concerns has been studied in advance research. According to Han et al, configuration issues account for 59% of performance issues. Displays a performance flaw in Apache as a result of the configuration. When one enters a high value for the configuration parameter start servers (for example, 60), Apache restarts

more slowly than usual. A valuable method called dummy connection contained inside a for loop is the number one culprit in this problem. This dummy connection technology initiates Apache Baby Server strategies through calling device features such as pick and ballot. To overcome this mistake, an if clause is added to the for loop. ConfPro prefers the white-field method to black-container performance profiling so that builders can choose configuration-dependent, inefficient code locations.

**KEYWORDS:** Artificial, Machine Learning

## INTRODUCTION

There are amazingly configurable software solutions available now. To regulate the program's functionality, users can alternate a wide range of configuration parameters. On the primary web page. It is necessary to respect the copyright of any part of this work not owned by acm. Credit-assisted abstraction is appropriate. Republishing, posting to a server, or redistributing to lists requires a great deal of configuration settings in the past, which can easily become an overall performance concern.

Sadly, builders often have no idea how the overall performance of a device suffers from configuration variables and the way they interact. The frequency of configuration errors inducing overall performance concerns has been studied in advance research. According to Han et al, configuration issues account for 59% of performance issues. Displays a performance flaw in Apache as a result of the configuration. When one enters a high value for the configuration parameter start servers (for example, 60), Apache restarts more slowly than usual. A valuable method called dummy connection contained inside a for loop is the number one culprit in this problem.

This dummy connection technology initiates Apache Baby Server strategies through calling device features such as pick and ballot. To overcome this mistake, an if clause is added to the for loop. ConfPro prefers the white-field method to black-container performance profiling so that builders can choose configuration-dependent, inefficient code locations. Conf Prof is divided into two sections. Conf prof collects execution profiles with various configuration choice values in order to accomplish this, and then it estimates a complexity variant to estimate the . Developers can use the rank to decide which configuration options are likely to most significantly affect overall performance on the software program under consideration. We plan to use confprof in at least 3 situations. First, confprof can be used to rank configuration options taking into account their performance impact by means of a developer. Confprof is based on dynamic analysis, like different profiling strategies, and is consequently limited to monitoring performance performed using a positive set of inputs. Second, developers can use ConfProf to understand code segments in which configuration option costs affect performance. Confprof

helps developers pinpoint areas of code where configuration-related changes affect overall performance in areas of concern.

In contrast, conprof is a dynamic approach in which program execution profiles are used to identify configuration parameters that have an impact on overall performance. As a result, on a large scale, ConProf can scale with multiple software systems without source code. Third, developers and researchers who build overall performance prediction models for software program structures that are originally based solely on configuration parameters can benefit from ConProf. Through sampling the overall performance-impacting configuration parameters determined using ConProf, they can combine existing performance modeling methodologies with ConProf. We use the technique on 4 c/c++ real-world programs to assess the efficacy of ConfProf. Our findings show that conprof efficiently detects configuration settings that may affect overall performance. Finally, the follow-up contributions are made through this chapter:

- A white container, dynamic aggregate performance analysis methodology that automatically ranks the impact of configuration options on the aggregate performance of specially customizable software architectures.
- A method that hyperlinks specific code areas to configuration selections that have an impact on the high performance of the target program.
- A method that is realistic and uses open-source implementation toolkits for c/c++ programs inside the real global child.

## **OVERALL PERFORMANCE BUG ASSOCIATED WITH CONFIGURATION**

In advance research, the difficulties of solving aggregate performance problems in highly flexible software structures were tested. According to the test, configuration settings are the cause of more than half (59%) of the 193 performance issues tested. A misconfiguration often results in subpar software program performance. There is a root cause. Software defects caused by code errors fall under the first category. Parent 1's software bug is one such example. As an example, in Apache Trojan horse #45834, a firewall misconfiguration disrupts authentication communications, causing the device to be blocked. Previous studies show that only a small percentage of problems (8% to 17%) are due to machine environment-unique configuration overall performance defects. As a result, we focus on the first kind of mistakes in overall performance coding in this look.

## **APPROACH**

This section introduces a performance profiling technique called ConfProf, which describes how configuration variables can impact the general performance of a device and helps builders do the same. The method is damaged in Parent 2, which can be seen here. Conprof takes as input a program that can be configured as well as a usage scenario that tests the program's configuration. The confprof process is divided into two levels. During the first part of its operation, conprof examines the ways in which male or female code spaces, including loops and machine calls (such as write(), ballot(), and select()), are plagued by configuration parameters. who can be chosen. A nice way to better illustrate the technique is we can speak it entirely in phrases of loops. Conprof serves this purpose by collecting execution profiles for detail settings and then inferring the location-degree complexity fashion (nearby-degree model for short) of the code. A space-degree variant is a model that explains the amount of time spent executing a particular code website, either for a single configuration option or for different combinations of values for options that interact with each other.

This model describes that the execution time (with a loop) in a code region is proportional to the number of iterations through the loop. During the second one phase, ConProf will calculate the performance impact of each configuration option, as well as summarizing the location-of-phase fashions collected during the primary phase. The results of step ii are provided in the form of a ranked list that gives information on each individual option and the interactions between those options. The preference with a better rating has a greater impact on the performance of the software program machine. Take a look at the latter example to help demonstrate factorization. Profiling techniques have a recognized drawback, which is that they are dependent on a chosen set of entry values. This makes it difficult to effectively uncover overall performance issues within the problem being investigated. In order to overcome this difficulty, several special strategies for generating check instances were proposed. These methods aim to generate huge workloads, taking a look at the input with the intention of increasing the likelihood of uncovering performance issues. On the other hand, there are many shortcomings in the existing techniques for improving the performance of examinations. Many of these techniques deal with different values of specific input parameters while leaving the parameter values unchanged from their unique settings. For example, Barnim et al. Listen to their attempts to increase the workload steps of information input while keeping default values for configuration parameters. Due to the combined results of many input parameters, it is possible that those techniques fail when it comes to finding performance defects. For example, in Apache, performance concerns are best left uncovered while decisions are made on

configuration option maintainability and request read timeouts. This is the case in all other cases. If the default configuration is used, however, this performance drawback no longer matters how many tons of work is done on the machine (eg, how many requests are made).

The above problem can be solved by performing a comprehensive aggregate with each possible combination of entry parameters; But, this approach is not always sensible due to the tremendous amount of feasible permutations. This technique, on the other hand, modifies all the input parameters of the application, which is probably useless for the reason that different parameters may not contribute to the universal performance of the software.

With regard to performance testing, the methods that are used are needed in an attempt to decorate the opportunity to discover overall performance defects. If there is a large wide variety of possible configurations, the cost of testing software can become very high. While sample-based methods were suggested as a way to cut the cost of configuration testing, those techniques are not robust enough to detect overall performance defects. These strategies aim to achieve a high level of coverage; However, revealing performance flaws routinely requires the use of positive admixture values and configuration alternative mixtures. Furthermore, performance test prediction is difficult to fine-tune because the incremental time it takes to execute the test is not always an adequate standard. There are several signs and symptoms that can be used to diagnose a performance fault.

In an ideal world, programmers would discover every defect at some point in the check out process. When a performance issue emerges in the build (for example, a substantial slowness with http replies in the web server), device administrators or builders will need to change the system to find configuration parameters that result in improved performance. But, when a system has a wide variety of configuration options, finding the appropriate settings for that system can be a difficult and time-consuming endeavor.

## **RESEARCH METHODOLOGY**

The goodness of a software device is directly related to the performance of a software program. A primary overall performance degradation may be due to an overall performance defect, behind schedule response times, and coffee system throughput. However, practical flaws often cause the gadget to crash or provide incorrect effects. Those insects cause full size damage to the person's experience. Problems affecting the performance of a software are some distance more difficult than bugs affecting the functionality of this system. This is because performance bugs routinely make themselves known through wide inputs and relatively restrained execution

contexts. Because of this, traditional testing methods, including methods based on coverage, have not been successful. A good way to spot performance issues is to have full-size studies done, specializing in dynamic processes in general, to investigate, find, and correct overall performance defects. Although these processes can discover performance flaws within the benchmark apps that have been analyzed, it is in most cases unknown whether or not they are successful in the large scale software tasks that are used in real international, servers with applications.

Worm tracking frameworks, such as Bugzilla and the issue tracker on Github, are used in a wide range of modern software program development tasks. These tools make it easy for both closed customers and software engineers to register problems they are experiencing with the product. Researchers not only use malicious program reviews to resource developers in understanding and repairing problems, but they also use them to evaluate an inspired testing or debugging technique. This is due to the fact that researchers use worm reviews to help builders understand and repair problems. Bug reviews are expected to provide developers with the facts they need to higher understand and solve problems. Researchers have the ability to determine, based on the description of the performance disorder documented that has been stable, whether or not the performance disorder can be included in their evaluation. Researchers will speak of a performance difficulty as a failed-to-reproduce computer virus if the researchers themselves are not able to reproduce the problem. In most cases, this is due to loss of functionality within the relevant area or within environment limitations (eg, compilation, dependencies, and many others). Because of this, the method of identifying insects is reasonably difficult, which may also dissuade researchers from examining a vast range of competence defects that may be relevant to the proposed approach. "Worm replication is extremely time consuming and difficult due to restricted and usually misleading data," note the authors of the currently published research on the dynamic detection of composite performance defects. They keep announcing that it can take up to a month for them to properly demonstrate a single issue in their test environment. We reviewed over thirty specialized articles on aggregated performance testing and analysis, and none of them explained how performance flaws could be replicated in a test environment. It turned into a first-rate ditch in the literature.

There is neither a test document nor a disclosure record that, to the best of our understanding, shows that the performance defects are so difficult to explain and recreate, and we have searched extensively and less for this type of report. For a Trojan horse record to be acceptable, the inputs, movements can be reproduced, and it is very important to check the oracle. One of

the problems that composite performance testing tools need to deal with is the fact that in order to reveal performance defects, they usually require a large amount of work or need to meet precise ambient conditions. it occurs. But, based on our previous investigations and observations, we found that it is very possible to be unable to replicate performance defects even when the oracle is checked with the stated inputs, replication methods, and computer virus reviews. This turned out to be our find, and it is true that it is still capable of replicating display defects. Next to the first rate of the worm file, some other queries that are requested are, "What other variables besides the good of the document result in failed attempts to duplicate defects?" It's a valid question to invite, and it should be asked. In order to maximize your chances of successfully duplicating performance defect reviews, it can be of tremendous help if we seek to understand these capabilities and provide answers to problems that we have not been able to duplicate.

The reason for this work is to quantify our enjoyment in reproducing performance worm reviews through examining the effect of various factors on each of the overall performance bugs obtained from open-supplied project worm reviews. This work was carried out for the purpose of sharing our information. This work was finalized with the intention of providing some of my information to others. We provide a fairly large number of answers that will increase your chance of effectively duplicating the overall performance problem. The primary validation of our study is to gain knowledge of non-reproducible performance issues from the researchers' perspective, as an opportunity to try to identify non-reproducible defects from the developers' perspective. The Apache http server and the mysql database are both important examples of open-supplied server projects, so we spent most of our research on them. Server applications are where we place our maximum interest considering mistakes in their performance are more likely to arise inside programs that can be used at tremendous scale and that handle large amounts of data in the direction Keep an extended period of time. We randomly selected a few insects, analyzed them, and then tried to duplicate them among the remaining 93 insects we studied. The primary purpose of this study's findings is to provide researchers with a better understanding of the challenges involved in the technique of overall performance bug replication and to advocate solutions to ease the process of malicious program selection. 32 Next is a list of number one results and contributions made through our study:

- By carefully following the specific details provided in the worm reviews, we attempted to recreate the overall performance problems that were previously addressed by the developers. After working on it for a period of six months, we were able to reproduce

only 17 insects out of 93. Our research has shown that the vast majority of overall performance issues (81%) cannot be replicated.

- We examined and reproduced the characteristics of 17 specific composite performance problem actions. Worm ten reports showing temporary overall performance issues are also cited as difficulty reproducing. An impressively large percentage (sixty-nine percent) of repeated performance issue reports required additional methods to be completed.
- Only one of the seventeen performance defects stated can be replicated by following exactly the outline provided in the worm report. But, so someone could repeat the previous 15 troubles, we had to use positive measures.
- We looked at various possible causes of overall performance problems, but despite our best efforts, we were not able to duplicate them. These factors include lack of hints, hardware requirements, system dependencies, object dependencies, unavailability of supplied code, compilation errors, set up errors, missing steps, and absence of hints. The overwhelming majority (39%) had missing steps, machine dependence, and lack of symptoms.
- We continued our investigation into possible causes of performance problems that could not be replicated in the initial attempt. We've given a list of several methods that, if taken collectively, should increase your chances of effectively replicating performance problems.

## INSPIRING EXAMPLES

Programming defects and configuration issues that cause a drastic drop in overall performance are what we are talking about when we talk about software performance defects. They can have a negative impact on the device's velocity, throughput, and responsiveness, ultimately resulting in a bad experience for the user. There are also some larger phrases that can be used regularly, including "performance problem" and "performance problem". For the duration of this lesson these terms can be used interchangeably. In an attempt to answer the latter question, we'll use 3 examples taken from overall performance ailment reviews: 1) What are some of the limitations shown in the bug reviews that may be contributing to the performance issues that should be able to be reproduced Not possible ? 2) What are we able to do to improve our chances of reproducing the overall performance disorder correctly, and how can we do it?



Apache worm #58037 is a composite performance Trojan horse file that has not been enabled to reproduce. After upgrading the Apache server from version 2.2 to model 2.4, the person reporting the issue noticed that the login method protocol (LDAP) for the light-weight listing took longer to obtain access rights. But, in spite of our great efforts, we were not successful in reproducing this display defect for some reasons. To begin with, the Trojan horse file no longer included any information on the defunct Apache server's minor version. It's pretty much impossible to determine which model is faulty because it will take too long (up to twenty-eight hours in the worst case). We were able to adopt a model in the long run that is as close to the time factor as possible, while reporting performance issues; But, no matter what, we were unable to copy the bug for the reasons mentioned in the previous sentence. Second, the wrm docs state that the configuration parameter ldap connection poolttlinside the ldap module should have a charge of zero for the problem to be reproduced.

We believe that some configuration options must be set to appropriate values for the problem to appear, but the report does not mention any of these requirements, despite the fact that they are largely dependent on the ldap module. Can be Third, the symptom is described as "We have determined that testing the repository at scale will take longer." This is stated inside the Trojan horse documentation. It is not clear exactly what length is meant by the phrase "a large repository". However, facts are necessary that allow you to carefully simulate the input load that may be required to reflect the overall performance defect and be able to detect the expected symptom. Despite the duplicated procedures as nearly practicable, we were unable to note that the Were unable to find the exact one inside the bug record.

Apache computer virus number 27106 is a reproducible performance bug file attempted to be assembled. While testing with the Apache benchmark, the bug locator noticed a memory leak, which resulted in the device slowing down. Primarily, the amount of RAM required via the httpd method improved exponentially when testing using http requests on ports that support ssl encryption. To begin with, there has been a lack of clarity in the way the environment has been defined. There has been a lack of facts on the version of the Linux running machine (OS) in which the problem was encountered. In addition, dependent modules—such as the OpenSSL module—have to be enabled with Apache Server Model 2.0. 45 is not mentioned anywhere in this article. When Apache is compiled, its configuration has to be changed so that the openssl module is covered. 2D is the loss of detail within the detail of the input.

Reporters of the disorder recommend that the Apache Benchmark (AB) be used to replicate malicious programs; But, the parameters that can be sent to AB are not given. Hypertext

Transfer Protocol Restful (https) is not supported with the help of the Apache benchmark which includes version 2.0.45. (https). It is necessary for us to find an AB version that is of the same mind as https. Third, there was a lack of clarity in the description of the signs and symptoms observed. It was the duty of the person who suggested the problem to cause users to monitor memory usage on Apache's main thread (eg using Linux machine monitoring tools with PlayStation to reveal system memory). Instead, the author decided to offer a crude hint and left it up to the readers to decide which information is maximally relevant. We spent about ten hours researching viable components to fill in the missing facts that we needed to reproduce this Trojan horse file, and finally, we were able to reproduce the performance issue. First, we build Apache using its default configuration to check whether the desired model (v2.0.45) is usable. The release date of Apache model 2.0.45 is the one we use to decide whether a version of openssl is of the same ideology (ie, openssl 0.9.7a

There is a bug report that can be reproduced that affects performance: MySQL computer virus #74325. Mysql version five. suffers from this problem, which causes a drop in performance. Regarding the method of updating indexed columns, MySQL v5.zero.85 is much faster than MySQL model five.7.five. The person who observed the problem has precise facts on the inputs that caused the worm, the way the environment changed in the installation, and the symptoms that are being observed. It is reported that the mysqlslap benchmark is being provided to the device, this is the first step. The person who saw the fault also says that the overall performance worm can only be caused by the use of a certain combination of configuration settings (for example, query cache length). Second, the description of how the environment is set up, while brief, is unique. The reporter makes it clear exactly which version of MySQL (ie, v5.7.five) causes the performance drop as well as which software components and versions of these components depend on MySQL v5.7.5. In MySQL model 5.7.5, "updating on an indexed column" takes more than twice as long as it did in MySQL version 5.6.21, according to the symptom outline, which is specific enough to indicate an overall performance fault. Because this Trojan horse file provides more accurate data than previous problem reviews, it took us about five hours to correctly copy the Trojan horse.

## **RESULTS**

### **rq1: Reproducible Malicious Program Reviews and Their Characteristics**

In Table 1 with columns representative and #failed, we present the amount of issues that can be recreated in addition to those that fail to do so in all different versions of the two subjects.

**Table 1 characteristics**

Subject	init relay	last link	#Sample	#failed	# Representative	success rate
Apache 2.0	2002	2013	20	16	4	20%
Apache 2.2	2005	2017	31	26	5	16%
Apache 2.4	2012	2017	4	3	1	25%
MySQL5.0	2005	2012	19	16	3	15%
MySQL5.1	2008	2013	15	12	3	20%
MySQL5.5+	2010	2017	4	3	1	25%
Joint	,	,	93	76	17	,

Finding 1: The majority (82%) of the above overall exposure pests fail to reproduce. The fee

**Table 2 Reproducible bugs and their characteristics**

Subject	bugid	set	Investment	choose	Burden	Take action	Command	duration	workaround
aboriginal people	54852	12	0	1	Yes	4	Yes	momentary	Yes
aboriginal people	52914	9	2	2	No	3	Yes	permanent	No
aboriginal people	37680	6	1	2	No	2	Yes	permanent	Yes
aboriginal people	22030	12	1	0	No	2	Yes	permanent	Yes
aboriginal people	51714	11	1	0	Yes	7	Yes	permanent	Yes

```

Environment Setup
1. export INSTALL=$PWD/apache-install/
2. ./configure --prefix=$INSTALL --enable-sed --enable-proxy
...
12. python -m SimpleHTTPServer & #Start backend server
13. ./bin/apachectl start #Start proxy server
Data Input
A static file with 1M+ characters on a single line
Configuration Options
Header unset Content-Length
SetOutputFilter Sed
ProxyPass / http://127.0.0.1:8000/
Input Actions
HTTP request: http://localhost:8000/a.1
    
```

**Figure 1**Reproducing the Apache bug

**Table 3**Workload Types

workload type	Description	bug example
web traffic	concurrent web page requests	Apache bug #54852
web traffic	long http connection session	Apache bug #43081
my sql	large number of database tables	Mysql bug #15653
my sql	concurrent update on db table	Mysql bug #74325

Finding 4.4: To reiterate, the majority (fifty-three%) of overall performance defects want to be subjected to a certain amount of labor before they become apparent. Table 3.2 In the column labeled "Act", enter the different types of play that must be performed to be able to characterize the performance disturbance. After the initialization of the environment, we specify the entry movement as one of the logical steps if we want to trigger the overall performance fault. An example in action can be seen in Figure 3.1, which shows the sending of an http request. Locating 6: The overall performance problem The great majority (88%) of the reviews can be found repeating names for more than three entry operations.

The records included in the column sequence of Desk 3.2 show whether the input games need to be treated in a certain order, with the aim of reproducing the overall performance problem explained. In step with the findings, each of the 17 worm reports called for entering multiple activities to provoke performance issues. Due to the fact that the objects of our research are server programs, replication of these programs must begin with the process of starting the server. Mysql difficulty #26938 is a performance worm that manifests itself as the database server becoming unresponsive even though it is passing through the list of currently used

instructions for miles. Want to eliminate the latter processes as a way to reproduce this bug: 1) start the database server the way you want using `./bin/mysqld secure`; 2) Connect a class buyer to the database server with the help of running `./bin/mysql`; and iii) `display profile`;" Submit a class question with the help of walking. Despite this, the series in which Enter Sports ended has remained relevant 9 issues after the start of the server. For example, for the purpose of a spike in CPU usage in Apache Trojan horse #37680, a sequence of enter sports must be performed in the actual order shown in Figure 2.

```
1  Replace default port with "Listen 50000"  
2  Add another port option "Listen 50001"  
3  Restart sever  
4  Make a request on port 50000  
5  Delete option "Listen 50001"  
6  Restart server  
7  Make a request on port 50000
```

**Figure 2of input actions**

Reproducing the problem in 52.9% of cases requires multiple entry steps, in which the sequence of events is very important. In the column titled "Duration" of Desk 4.2, the period of time during which the performance problem was present is specified. A symptom is said to be permanent if it can be located at any time after it has come to light. Alternatively, a symptom is said to be transient if it appears for only a short period of time and then disappears.


Findings: A large proportion (47%) of all suggested bugs have safe hints. As an example, in Apache Malicious Program #48024, a dramatic jump to 100% CPU usage occurs After that, the CPU consumption indicator returns to its previous, regular state. The figures contained in the Workaround column of Desk 3.2 show whether or not it is important to take the time to bypass some problems (including an obscure description or model inconsistencies) that will reproduce the overall performance worm.

Finding Three.Five: The vast majority (88%) of problem reports that name can be repeated for some sort of solution. As an example, take a look at the do abi block in the Makefile.in, the change in behavior given in subsequent changes to the GCC compiler causes MySQL to fail to build issue #44723. After the block was removed, MySQL was able to compile without issue.

### **rq2: Main factors not reproducing overall performance bug reviews**

Before we can move on to increasing our chances of success in replicating the overall performance computer virus report, we want to pick out the primary targets that cause replication to fail. It may allow us to transform education in a way that increases the potential for fulfillment. In the case of assigning a bug to a class, the eight classes do not overlap with

each other. For example, a Trojan horse record may also have "deficient steps", but if we run into "compile errors" trouble, the malicious program reports not counting the number under "deficient steps". until we get to the "Compile Errors" phase. This also happens when we are able to stay away from the "compilation mistakes" phase. In this case, a single factor can be matched against a total of two times of each type. Parent 3 provides a graphical illustration of the breakdown of performance issue e reviews in each of the eight categories.



**Figure 3 bug breeding failure factor distribution**

**CONCLUSION**

In Bankruptcy 3, we completed an investigation on performance-related components of a bug we discovered in surprisingly customizable systems. We analyzed over three hundred

configuration-related aggregate performance issues arising from three of the most prominent open source initiatives. In the context of state-of-the-art highly flexible software, we examined a broad spectrum of performance parameters. To solve At some level in counseling devoted to discussion, we provide insight that is potentially useful to both students and practitioners in their respective fields. In Chapter Four, we mentioned our earlier enjoyment in copying the overall performance worm report and examined the effect of a number of factors on performance defects from open-supplied malicious program reports, both of which would be recreated or reimplemented. : Cannot be submitted. On the way to increase the chances of efficiently reproducing the overall performance defect, we offer an extension of the answers. We examined a very wide range of open-source server development initiatives. We chose several insects at random, examined them, and then attempted to breed as many of them as was viable. The reason for our research is to replicate overall performance defects from the researchers' point of view.

## REFERENCE

- [1] Apache http server benchmarking tool, 2019.
- [2] Apache bug 34508, 2005. <https://bz.apache.org/bugzilla/showbug.aspx.cgi?id=34508>.
- [3] Apache bug 42031, 2007. <https://bz.apache.org/bugzilla/showbug.aspx.cgi?id=42031>.
- [4] Apache bug 54852, 2013. <https://bz.apache.org/bugzilla/showbug.aspx.cgi?id=54852>.
- [5] TF Abdelzahar, K. Ji Shin, and N. Furnace. Performance guarantees for web server end-systems: A control-theoretic approach. *IEEE Transactions on Parallel and Distributed Systems*, 13(1):80–96, 2002.
- [6] Automated Combination Testing for Software, 2016. <http://csrc.nist.gov/groups/sns/acts/index.html>.
- [7] N Ali, W. Woo, Mr. Antoniol, M. D. Penta, Y. G Gueneuc, and J. H Hayes. Mothers: Multipurpose Miniaturization of Software. In *International Conference on Software Maintenance*, pages 153-162, 2011.
- [8] M. Attarian, M. Chou, and J. Flynn. X-ray: automating root-cause diagnosis of performance anomalies in production software. In *Proceedings of the 10th Unix Conference on Operating System Design and Implementation* , pages 307–320, 2012.
- [9] M. Attarian and J. Flynn. Automated configuration troubleshooting with dynamic information flow analysis. In *OSDI*, pages 1–11, 2010.
- [10] RA Beja-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley

Longman Publishing Company, Inc., Boston, MA, USA, 1999.

- [11] Beautiful Soup, 2017. <https://www.crummy.com/software/beautifulsoup/>.
- [12] C. Bird, A. Bachmann, E. Oun, J. Duffy, A. Bernstein, V. Filkov, and P. Devanbu. Fair and balanced?: Bias in bug-fix datasets. In Esec, pages 121-130, 2009.
- [13] SJ Bradke and M. Oh duff. Reinforcement learning methods for continuous time markov decision problems. Advances in Neural Information Processing Systems, pages 393–400, 1995.
- [14] E. Bruneton, R. Lenglet and T. Coupe. ASM: A code manipulation tool for implementing adaptive systems. Adaptable and Extensible Component Systems, 30:19, 2002.
- [15] Beatrice, 2016. <https://kenai.com/projects/btrace>.
- [16] X Boo, J. Rao, and C.-Z. Ju. A reinforcement learning approach to online web system auto-configuration. In Distributed Computing Systems, 2009. icdcs'09. 29th IEEE International Conference on, pages 2–11. IE, 2009.