# TRIPLEAGENT: Monitoring, Perturbation and Failure-obliviousness for Automated Resilience Improvement in Java Applications

Long Zhang and Martin Monperrus[1]

[1]KTH Royal Institute of Technology, Sweden

*Abstract*—In this paper, we present a novel resilience improvement system for Java applications. The unique feature of this system is to combine automated monitoring, automated perturbation injection, and automated resilience improvement. The latter is achieved thanks to the failure-oblivious computing, a concept introduced in 2004 by Rinard and colleagues. We design and implement the system as agents for the Java virtual machine. We evaluate the system on two real-world applications: a file transfer client and an email server. Our results show that it is possible to automatically improve the resilience of Java applications with respect to uncaught or mishandled exceptions.

*Index Terms*—fault injection, dynamic analysis, exception-handling, software resilience

## I. INTRODUCTION

In modern software, resilience capabilities are engineered through error-handling code, in particular exception handling code in managed languages such as Java and C#. This resilience capability is manually engineered by developers, who write the error-handling code. For example, part of their coding activity is to write try-catch blocks to handle exceptions. The problem is that exception-handling code is notably hard to write and to test [10]. As a result, there often exists corner-cases where resilience is not provided by developer-written code. In production, when those corner cases are activated, the software system may simply stop providing its function because it crashes after an unhandled exception [36].

In order to improve error handling, two kinds of techniques are being researched: fault injection [24], [38] and failure-oblivious computing [30]. Fault injection is about injecting failures to trigger a system's error handling code and to analyze the abnormal behaviour [32]. Failure-oblivious computing is about adding fully generic error-handling code with automated code transformation [30]. In the context of exceptions, failure-oblivious computing means automatically adding catch blocks with a default exception-handling strategy [9].

In this paper, *our goal is to automatically improve the exception-handling code of software applications*. This is made by first finding weaknesses in resilience and then instrumenting the application with automated exception handling. To achieve our goal of automatically improving resilience, we design a novel system, called TRIPLEAGENT, made of three components, called "agent" in this paper[1]. Those three agents, automated monitoring, automated perturbation injection, and automated failure-oblivious method validation [30] are orchestrated by an agent controller. The controller analyzes all the monitored data and reveals both weaknesses and suggested improvements in the resilience capabilities.

To the best of our knowledge, TRIPLEAGENT is the first system which actively injects exceptions during execution in order to, after analysis, automatically detect failure-oblivious methods.

We evaluate TRIPLEAGENT on two real-world Java applications. One is TTorrent, a file transfer client which implements the BitTorrent protocol. The other one is HedWig, an email server. In both cases, we consider a production workload: respectively downloading a large file from the Internet, and sending and receiving emails from the server. By applying TRIPLEAGENT, we observe that exceptions thrown from 257(21%) perturbation points do not lead to failures anymore, which shows an automatic resilience improvement.

To sum up, our contributions are the following.

- The concept of joint usage of fault injection and failure-oblivious code instrumentation to evaluate and improve resilience against uncaught or mishandled exceptions. We propose a corresponding novel algorithm for automatic improvement of software resilience.
- A system called TRIPLEAGENT that combines monitoring, perturbation injection and failure-oblivious computing in Java, implemented with agents for the Java Virtual Machine. The system is publicly-available for future research in this area at http://bit.ly/tripleagent-repo.
- An empirical evaluation of TRIPLEAGENT on two real-world applications of $20.3K$ lines of code in total. By performing 9968 fault injection experiments under a realistic, production-like workload, it shows TRIPLEAGENT's effectiveness for improving software resilience.

The rest of the paper is organized as follows: Section II introduces the background. Section III and Section IV present the design and evaluation of TRIPLEAGENT. Section VI discusses the related work, and Section VII summarizes the paper.

---

[1]Here, an agent refers to the Java terminology, where it is a component that is attached to the Java Virtual Machine [14].

## II. Background

TRIPLEAGENT is founded on techniques from the fault injection and failure-oblivious computing [30]. This section presents a basic introduction to the core concepts.

### A. Fault Injection

Fault injection is a popular research topic in software testing and dependability evaluation. Fault injection techniques actively inject different kinds of errors into a target system in order to assess its dependability [24], [32], [38]. This can happen in several phases: 1) during unit testing, fault injection generates more test cases so that corner cases are detected, and the coverage of testing is improved. 2) during integration, fault injection can trigger different failure scenarios so that developers gain more confidence in their system's error-handling design. 3) when done in production, fault injection is usually called "chaos engineering" [2].

The kinds of failures that can be injected vary depending on the considered dependability aspect. For example, injecting processor errors or hardware-based errors is often done for evaluating the dependability of operating systems [16], [25]. Injecting connection errors between different micro-services supports to test a service's retry logic and its robustness of interacting with other services [6]. Injecting an exception in a certain method is useful for validating an application's error-handling capability [37].

In this paper, we focus on fault injection in the context of Java applications, which means rather high-level, application-level fault injection.

### B. Failure-oblivious Computing

In order to improve the resilience of an application, different techniques can be applied to prevent the application from crashing when an error occurs [23]. Failure-oblivious computing [30] is one of these approaches to overcome software failures at runtime. The main idea of failure-oblivious computing is to discard certain failures in a principled way. For example, if a method tries to write data into an invalid memory address, with failure-oblivious computing, the writing operation would be ignored. It has been shown that failure-oblivious computing is able to increase availability [9], [29], [30], eg to serve requests to more users despite errors.

In this paper, we use the concept of failure-oblivious computing for the Java programming language. In Java, there is no invalid memory addresses, but the biggest reason for crashing are exceptions thrown at runtime. Thus, we do failure-oblivious computing for uncaught and mishandled exceptions.

### III. Design of TRIPLEAGENT

This section presents the design of TRIPLEAGENT, including relevant definitions, algorithms and its architecture.

### A. Definitions

*a) Exception:* All major programming languages provide a way to signal problems through so-called exceptions. In some statically typed languages such as C# and Java, exceptions are typed, and some of them are statically verified at compile-time (e.g., checked exceptions in Java). For these checked exceptions, developers need to either handle them at the call site or explicitly declare them in the method signature (with the keyword $throws$ in Java [15])

*b) Perturbation point:* In this paper, a "perturbation point" is a unique location in code where a fault can be injected. In TRIPLEAGENT, the considered perturbations are injected exceptions, which means that a perturbation point is defined as a statement that potentially throws an exception. A perturbation point is noted $< m, l, e >$. $m$ describes the method where this point is located. $l$ is the line number before which the exception is thrown. $e$ is the type of this exception.

*c) Fault model:* In this paper, we consider two fault models: 1) injecting only one exception, when the perturbation point is reached for the first time and 2) always injecting exceptions when the perturbation point is reached.

*d) Perturbation search space:* We define the "perturbation search space" as the Cartesian product of all possible perturbation points and all fault models with respect to a workload [8], [9]. The size of the search space is the number of workload executions required to have an exhaustive picture of the behavior under perturbation.

*e) Exception handling method:* When an exception $e$ is handled in a method, this method is called the "exception handling method" for $e$. In a tuple <exception source, exception type, method>, the source refers to the location where the exception was thrown. In this paper, we make a distinction between "default exception handling methods" and "failure-oblivious method": the former refers to methods with manually written catch blocks while the latter refers to method with automatically instrumented catch blocks.

*f) Acceptability oracle:* An acceptability oracle is a mechanism for determining whether an application's behaviour remains acceptable under perturbation.

In order to evaluate and improve an application's resilience, we use oracles to describe acceptable behaviour, hence we call them "acceptability oracle". In this paper, an acceptability oracle is a combination of generic oracles (like the absence of crash) and domain-specific ones. For example, in the context of a file downloading client, an acceptability oracle could be that 1) the client does not crash and exits normally 2) the client successfully downloads the file with a correct checksum.

*g) Failure-oblivious method:* A method $fo$ is said to be failure-oblivious with respect to a perturbation point $p$ if and only if: 1) when an exception is thrown from $p$, it is possible to catch and stop its propagation in $fo$ that is upper in the call stack; 2) the behaviour of the application verifies the acceptability oracle if the exception thrown from $p$ is caught in $fo$. This is noted $< m, l, e > \mapsto fo$ (here $\mapsto$ denotes "thrown exception $e$ in method $m$ before location $l$ and caught at $fo$").

*h) Fault injection experiment:* Given an application $a$ and a workload $w$, injecting an exception during the execution of $a$ under $w$ is called an "experiment". TRIPLEAGENT is

designed to conduct experiments in order to evaluate and improve an application's error-handling capability.

*Example.* Let us assume an invocation chain across three methods: $m2 \rightarrow m1 \rightarrow m0$. Method $m0$ can throw an IOException before line number $l$, and the developers write a try-catch block to handle it in $m2$ (the method upper in the stack). Consequently, $m2$ is the default exception handling method for this exception. If the exception is caught and silenced in method $m1$ and the application behaviour is still acceptable according to the oracle, method $m1$ is considered as a failure-oblivious method for $< m0, l, IOException >$.

### B. Goals of TRIPLEAGENT

TRIPLEAGENT aims at improving the exception-handling capabilities of Java applications. The main goals of TRIPLEAGENT are: 1) to give developers feedback about the effectiveness of their exception-handling design; 2) to automatically identify improvements of exception handling.

The former is about detecting the weakness points of the system under consideration and the latter is about finding new failure-oblivious methods that improve the application's resilience.

**Input to TRIPLEAGENT:** TRIPLEAGENT takes arbitrary software written in Java as input and a workload. No manual change is required from the developer. Neither source code nor test suite are required for improving an application's resilience. TRIPLEAGENT also takes as input an acceptability oracle, which will be explained below in Section III-A.

**Output for the developer:** The output of TRIPLEAGENT is a report for developers. The report gives three pieces of information: 1) the perturbation points and their classification as defined next; 2) the verified failure-oblivious methods, i.e. the resilience improvements; 3) a log file which contains all the monitored information for the purpose of further analysis.

TRIPLEAGENT classifies the perturbation points into three categories as follows:

**Definition 1.** Fragile points: A fragile point is a statement in a method before which injecting one exception results in the application crashing or freezing.

**Definition 2.** Sensitive points: A sensitive point is a statement in a method before which injecting one single exception does not influence the application in the workload under consideration. But, continuously injecting exceptions results in the application to crash or freeze.

**Definition 3.** Immunized points: An immunized point is a statement in a method before which no matter how many exceptions are injected, the application still behaves acceptably.

### C. Core Algorithm

The whole procedure of TRIPLEAGENT to identify failure-oblivious methods could be split into 3 steps:

1) Define acceptability oracles and detect perturbation points. TRIPLEAGENT executes the application normally, in order to monitor and record the application's normal behaviour. With this execution, the perturbation agent goes

---

**Algorithm 1** Detection of Perturbation Points

**Input:**
    An application $A$;
    A repeatable workload for this application $W$;
**Output:**
    $P$, a set of perturbation points;
1: Execute the application $A$ normally under $W$;
2: **for** each method $m$ loaded into the JVM **do**
3:     **for** each checked exception $e$ thrown from $m$ **do**
4:         **for** each location $l \in m$ **do**
5:             $P \leftarrow P \cup \{< m, l, e >\}$;
6:         **end for**
7:     **end for**
8: **end for**
9: **return** $P$;

---

**Algorithm 2** Automated Classification of Perturbation Points

**Input:**
    An application $A$;
    A repeatable workload for this application $W$;
    A set of perturbation points $P$;
    An acceptability oracle $O$;
**Output:**
    $F$ a set of fragile points points;
    $S$ a set of sensitive points;
    $I$ a set of immunized points;
1: **for** each point $< m, l, e >\in P$ **do**
2:     Execute the application under $W$;
3:     **if** $m$ is executed for the first time **then**
4:         Inject the exception $e$ before location $l$;
5:         $b1 \leftarrow$ application behaviour;
6:     **end if**
7:     Execute the application under $W$;
8:     Always inject exceptions $e$ before location $l$;
9:     $b2 \leftarrow$ application behaviour;
10:     **if** $b1$ not meet O and $b2$ not meet $O$ **then**
11:         $F \leftarrow F \cup \{< m, l, e >\}$
12:     **else if** $b1$ meet $O$ and $b2$ not meet $O$ **then**
13:         $S \leftarrow S \cup \{< m, l, e >\}$
14:     **else if** $b1$ meet $O$ and $b2$ meet $O$ **then**
15:         $I \leftarrow I \cup \{< m, l, e >\}$
16:     **end if**
17: **end for**
18: **return** $F, S, I$;

---

through all classes loaded into JVM and locates every perturbation point, based on method signature, which is described in Algorithm 1.

2) Classify all the perturbation points into fragile, sensitive or immunized ones (as defined in Section III-B). For each perturbation point, TRIPLEAGENT conducts two experiments: only injecting one exception when the point is reached for the first time, and always injecting exceptions when the point is reached. Based on the observation of the application behaviour

**Algorithm 3** Detection of Failure-oblivious Methods

---

**Input:**
    An application $A$;
    A repeatable workload for this application $W$;
    A set of perturbation points $P$;
    An acceptability oracle $O$;
**Output:**
    A set of failure-oblivious methods $R$;

1: // Find candidate failure-oblivious methods to be assessed

2: $Q \leftarrow \emptyset$ // Worklist for candidate methods;
3: **for** each point $< m, l, e > \in P$ **do**
4:     Execute the application under $W$;
5:     **if** $m$ is executed for the first time **then**
6:         Inject the exception $e$ before location $l$;
7:         **for** each method $n$ in the call stack **do**
8:             $Q \leftarrow Q \cup \{< m, l, e > \mapsto n\}$;
9:         **end for**
10:     **end if**
11: **end for**
12:
13: // Assess all candidate failure-oblivious methods
14: **for** each perturbation point $< m, l, e > \in P$ **do**
15:     **for** $< m, l, e > \mapsto n \in Q$ **do**
16:         Execute the application under $W$ twice as specified in Algorithm 2;
17:         When $e$ is thrown, catch the exception $e$ in $n$;
18:         **if** the behaviour meets $O$ **then**
19:             $R \leftarrow R \cup (< m, l, e > \mapsto n)$;
20:         **end if**
21:     **end for**
22: **end for**
23: **return** $R$;

---

under perturbation, the perturbation point is classified using Algorithm 2.

3) Identify candidate failure-oblivious methods and evaluate each of them as described in Algorithm 3. TRIPLEAGENT detects candidate failure-oblivious methods with call stack analysis: every method in the stack before the default handling method is identified as a candidate failure-oblivious method. Then two fault injection experiments are conducted (only inject one exception, inject several exceptions). The difference is that all the thrown exceptions are caught in a catch block instrumented in candidate failure-oblivious methods. By analyzing the behavior once the exception is caught in this catch block, TRIPLEAGENT confirms whether the method under evaluation is indeed failure-oblivious or not.

### D. Architecture of TRIPLEAGENT

Figure 1 shows the general architecture of TRIPLEAGENT. TRIPLEAGENT considers a Java application in a JVM, such as a backend web application or a Java micro-service.

When an application is loaded into the JVM, TRIPLEAGENT attaches to it three different agents: a monitoring agent, a perturbation agent and a failure-oblivious agent. The monitoring agent is responsible for collecting the information needed by TRIPLEAGENT to evaluate the system's resilience capabilities. The perturbation agent injects exceptions into the application in order to trigger its error-handling logic. The failure-oblivious agent tries to improve the application's resilience by catching and silencing exceptions before they are handled by default exception handling methods.

All the agents are controlled by a controller which makes two kinds of decisions: 1) given an application under some specific workload, which perturbation point should be activated, 2) whether the point's corresponding failure-oblivious method should be switched on. Finally, the controller generates a report for the developer based on data gathered from a series of fault injection experiments.

*1) Monitoring agent:* In order to study the influence of perturbations and evaluate all possible failure-oblivious methods in a software system, it is necessary to collect different kinds of monitoring information. For this, we propose to use a monitoring agent that is attached to the runtime process.

Our monitoring agent works as follows. For each method in the code loaded in the JVM, the agent collects static and dynamic information.

The static information is: 1) its position in the code, 2) whether it declares checked exceptions to be thrown.

The collected dynamic information is: 3) the number of method executions over an fault injection experiment, 4) each time an exception is caught, the agent collects the stack information, including the stack distance between the method raising the exception and the method catching it. This includes both exceptions caught in default exception handlers and in failure-oblivious methods (as defined in Section III-A).

The TRIPLEAGENT monitoring agent also collects the following information:

- The set of classes that have been loaded into the JVM.
- Whether the application has exited normally or crashed due to an unhandled exception.

*2) Perturbation agent:* The perturbation agent injects specific perturbations at a specific point in time. The perturbation commands come from the agent controller.

The perturbation agent detects every method with a throws keyword and attaches itself into this method by rewriting the bytecode. In order to explore the entire perturbation search space, the agent injects different perturbation points before each statement in the method. In this way, the agent is able to throw such an exception anywhere in the method and compare the difference.

```
1  // a perturbation point in Class1
2  void exampleMethod() throws ExceptionA, ExceptionB {
3    // code injected with code transformation
4    PAgent.throwExceptionPerturbation(key1);
5    PAgent.throwExceptionPerturbation(key2);
6    // a statement
7    PAgent.throwExceptionPerturbation(key3);
8    PAgent.throwExceptionPerturbation(key4);
9  }
```

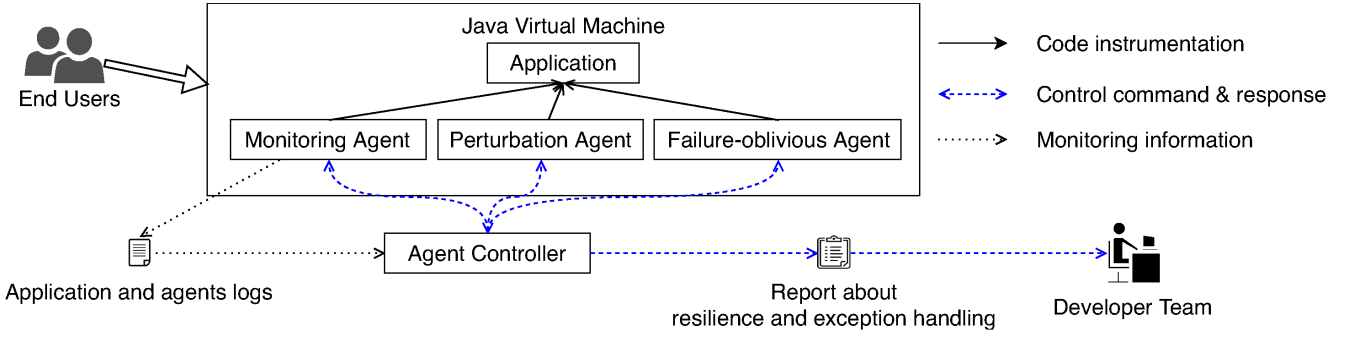Listing 1. The Perturbation Strategy in TRIPLEAGENT

Fig. 1. The components of TRIPLEAGENT

Listing 1 gives an example of how this perturbation agent works. When a method like exampleMethod() throws multiple exceptions, corresponding perturbation points are automatically injected with code transformation. The perturbation agent controls every perturbation point separately. When a specific point is activated, it throws an exception at the beginning of the method.

*3) Failure-oblivious agent:* The failure-oblivious agent instruments the code with try-catch blocks during a fault injection experiment. For reasoning about resilience with respect to uncaught exceptions, the failure-oblivious agent injects a try-catch wrapper in all methods. Basically, the whole method body is wrapped with a try-catch block which handles all types of exceptions (`catch Exception` in Java). By default, the catch block simply throws again the exception which makes it semantically equivalent to the original code. When the failure-oblivious method is activated, the injected catch block silences the exception and prevents it from propagating (note that the exception may come from this method or from other methods transitively called from this method).

When an exception is caught by the injected catch block, there are three possible outcomes: 1) the application runs normally; 2) the application runs in a gracefully degraded mode; 3) the application crashes.

```
1  // a candidate failure-oblivious method in Class2
2  void callExampleMethod() throws ExceptionA,
        ExceptionB {
3    try {
4      new Class1().exampleMethod();
5    } catch (Exception a) {
6      if (!FOAgent.modeIsOn(key2)) {
7        throw a;
8      } else { // nothing, the exception is silenced }
9    }
10 }
```

Listing 2. Automated Code Instrumentation for Identifying Failure-oblivious Methods in TRIPLEAGENT

Listing 2 illustrates how this is done. In method $callExampleMethod$, $exampleMethod$ is invoked. The failure-oblivious agent detects it as a possible failure-oblivious method. So the whole method body of $callExampleMethod$ is wrapped with a try catch block. When the agent controller activates this failure-oblivious method, it silences all exceptions coming from $exampleMethod$. Otherwise it throws the caught exception so that it is propagated as usual.

*4) Agents controller:* The agent controller is responsible for conducting a series of experiments (see Section III-A). It controls every agent and gathers all the information to analyze the system resilience. Additionally, the controller is configurable. For example, developers can define a filter to focus on resilience improvement for a specific package.

*E. Implementation*

There are different kinds of agents in the JVM. The monitoring agent is implemented on top of the JVM Tool Interface (JVMTI) [2]. The perturbation agent and failure-oblivious agent are implemented as JVM agents, using the ASM library for binary code transformation [3]. The agents controller is a standalone service, it communicates with the JVM and the agents through local files.

For sake of open-science, the code is made publicly available at http://bit.ly/tripleagent-repo.

## IV. EVALUATION

For evaluating this contribution, we apply a case-based evaluation methodology: this methodology consists of an in-depth analysis of relevant cases selected in a principled way [11]. In our research domain, it has been shown appropriate in Rinard et al's original paper on failure-oblivious computing [30].

We select two case studies according to the following three criteria: 1) the case should be a real-world application (i.e., not a toy example) 2) it should be medium-sized in order to be appropriate for the computing power available in the laboratory 3) it is possible to define a production-like workload. Those criteria yield two cases: TTorrent and HedWig. TTorrent is a file transferring tool which implements the BitTorrent protocol. HedWig is an email server for the IMAP, SMTP and POP3 protocols. They are also exemplary of applications with high resilience requirements: an email server must not crash, a file download on the dynamic internet must succeed regardless of unexpected network events, peer failures, or local machine issues.

The analysis of TRIPLEAGENT requires several executions. For each perturbation point, its classification requires 2 executions under the workload (as discussed in Algorithm 2). For
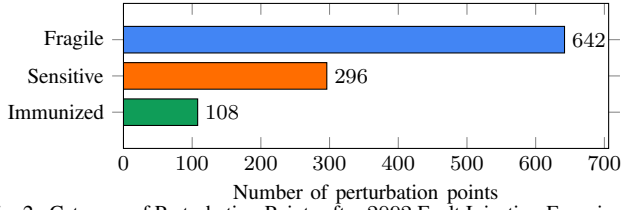
---

[2]See https://docs.oracle.com/javase/8/docs/platform/jvmti/jvmti.html
[3]See http://asm.ow2.org

5

Fig. 2. Category of Perturbation Points after 2092 Fault Injection Experiments on TTorrent



a) Fragile stays fragile, b) Fragile to sensitive, c) Fragile to immunized

d) Sensitive stays sensitive, e) Sensitive to immunized, f) Immunized stays immunized

Fig. 3. Resilience Improvement on TTorrent: Fragile points, sensitive points and immunized points are respectively shown in blue, orange and green. The area of bubbles are corresponding to the numbers of perturbation points under consideration.

each candidate failure-oblivious method, its evaluation also needs 2 executions under the workload (see Algorithm 3). For both cases, an execution takes no more than 1 minute in our testing environment. In total, the cost of the experiments presented in this section is upper-bounded by $2 \times 1 \times (1046 + 2844 + 372 + 722) = 9968$ minutes. Note that some experiments lead applications to a crash. It actually takes TRIPLEAGENT around 3 days to finish all the experiments.

### A. Evaluation on TTorrent

*1) Experiment Protocol:* We apply Section III-C to TTorrent, version 2.0.

The workload $W$ for TTorrent consists of downloading a large file (debian-9.9.0-amd64-netinst.iso, a Debian distribution installer of 292.0MB). Since BitTorrent is a peer-to-peer protocol, this workload involves other machines on the internet which serve (aka "seed") the file. To that extent, the workload is a production one. We perform a series of fault injection experiments, as described in Section III-A.
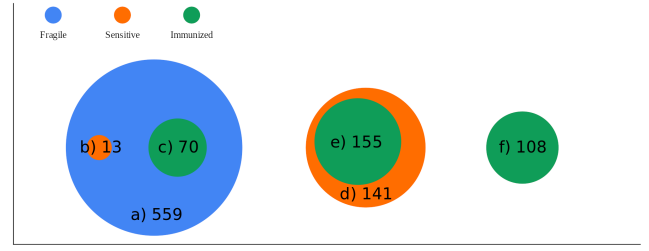
For TTorrent, we consider the following definition of acceptable behaviour to evaluate candidate failure-oblivious methods: the behaviour is considered acceptable if an end-user can successfully download a file with a correct checksum.

*2) Experimental Results:* Per Section III-C, the first step of TRIPLEAGENT is to execute TTorrent normally and to monitor all possible perturbation points. It detects 1046 points in total within the package com/turn/ttorrent.

Then, TRIPLEAGENT performs two series of experiments: 1) it injects one exception per perturbation point and compares the behaviour between these experiments and the normal execution, 2) it always injects exceptions when a perturbation point is reached and also compares the behaviour against the reference one.

As a result, all perturbation points get classified in the 3 categories defined in Section III. In total, TRIPLEAGENT identifies 642 fragile points, 296 sensitive points and 108 immunized points. Figure 2 shows the distribution of these perturbation points, which are used as a base line for the following experiments.

The next step of TRIPLEAGENT's main algorithm is to compute and assess the possible failure-oblivious methods. As explained in Algorithm 3, for a perturbation point, a set of failure-oblivious methods is identified. In our experiment, TRIPLEAGENT detects 2844 possible failure-oblivious methods, summed over all the perturbation points. The minimum, median and maximum number of candidate failure-oblivious methods per perturbation point is respectively 0, 2, 10.

Then, $2844 \times 2 = 5688$ executions are made to assess the failure-obliviousness of the candidate points (one per injection mode). Hopefully, the added catch blocks inserted by TRIPLEAGENT will increase the number of immunized points.

Let us now consider Figure 3. The fragile, sensitive and immunized perturbation points are respectively shown in blue, orange and green. The area of bubbles corresponds to the numbers of perturbation points under consideration. For example, the bubble e represents the 155 sensitive points transformed into immunized ones with failure-oblivious computing. Overall, TRIPLEAGENT successfully transforms 13 fragile points into sensitive ones, 70 fragile points into immunized ones and 155 sensitive points into immunized ones. The original 108 immunized points remain immunized. This means that resilience of the TTorrent has been automatically improved.

Table I presents a sample of perturbation points. Every row describes 1) a perturbation point (the class name, method name and its line number), the thrown exception type, and the corresponding default exception handler written by developers; 2) the failure-oblivious improvement (failure-oblivious method and concrete change of the perturbation point's category). For example, row 1 and row 2 show that TRIPLEAGENT verifies failure-oblivious methods which improve the original fragile perturbation points into sensitive ones. Row 7 and row 8 also describe the case that TRIPLEAGENT is able to detect multiple failure-oblivious methods for the same perturbation point. For original immunized perturbation points, alternative failure-oblivious methods which provide the same resilience are verified as well, which is shown in the last two rows.

*3) Case Studies:* In the following, we detail 3 case studies where the resilience is improved.

```
1  void send(PeerMsg m) throws IllegalStateException {
2    if (this.isConnected()) {
3      ...
4    } else {
5      // perturbation point here
6      unbind(true);
7    }
8    ...
9  }
```

Listing 3. IllegalStateException in SharingPeer/send

*a) Failure-oblivious Method as Alternative to Normal Resilience:* First, Listing 3 shows a failure-oblivious method

| 0 | Perturbation Point | Exception Type | Default Handling Method | Failure-oblivious Method | Improvement |
|---|---|---|---|---|---|
| 1 | BEValue/getNumber@122 | InvalidBEnException | ClientMain/main | BEValue/getLong | fragile - sensitive |
| 2 | HTTPTrackerC/encodeAnnoToURL@187 | AnnoException | TrackerClient/annoAllInterfaces | HTTPTrackerC/announce | fragile - sensitive |
| 3 | CommuManager/addTorrent@229 | IOException | ClientMain/main | CommuManager/addTorrent | fragile - immunized |
| 4 | TorrentParser/getStringOrNull@121 | InvalidBEnException | ClientMain/main | TorrentParser/getStringOrNull | fragile - immunized |
| 5 | HTTPTrackerC/sendAnnounce@235 | ConnectException | HTTPTrackerClient/announce | HTTPTrackerC/sendAnnounce | fragile - immunized |
| 6 | SharedTorrent/init@226 | InterruptedException | SharedTorrent/initIfNecessary | SharedTorrent/init | sensitive - immunized |
| 7 | SharedTor/handlePieceCompleted@671 | IOException | SharingPeer/handleMessage | SharedTor/handlePieceCompleted | sensitive - immunized |
| 8 | SharedTor/handlePieceCompleted@671 | IOException | SharingPeer/handleMessage | SharingPeer/firePieceCompleted | sensitive - immunized |
| 9 | WorkingReceiver/processAndGetNext@64 | IOException | ConnWorker/processSelectedKeys | ReadableKeyProcessor/process | sensitive - immunized |
| 10 | SharingPeer/send@352 | IllegalStateException | CommuManager/validatePieceAsync | SharingPeer/send | alternative resilient method |
| 11 | PeerMessage/parse@176 | ParseException | ConnWorker/processSelectedKeys | PeerMessage/parse | alternative resilient method |

TABLE I
SAMPLE OF PERTURBATION POINTS AND THE CORRESPONDING FAILURE-OBLIVIOUS METHODS IN TTORRENT

with respect to exception `IllegalStateException`. This method is executed only 1 time during normal download of the file. Under perturbation, TRIPLEAGENT identifies that if one single exception is thrown from this method, the application is still able to download the file correctly. By analyzing the stack, TRIPLEAGENT detects another two methods as candidate failure-oblivious methods: `SharingPeer/send` and `SharingPeer/notInteresting`.

By activating a failure-oblivious try-catch block in these two methods (i.e. the method body is wrapped with a try-catch block which blocks the exception), TTorrent still succeeds in downloading the file. It means that TRIPLEAGENT successfully detects 2 alternative methods in the stack that provide the same resilience as the original manually-written catch block.

```
1 DataProcessor processAndGetNext(ByteChannel sc)
      throws IOException {
2 ...
3 // perturbation point here
4 if (this.pstrLength > MAX_MESSAGE_SIZE) {
5   return new ShutdownAndRemovePeerProcessor(...).
        processAndGetNext(socketChannel);
6 }
7 }
```

Listing 4. IOException in WorkingReceiver/processAndGetNext

*b) Improving Resilience under a High Number of Exceptions:* Listing 4 shows method `processAndGetNext` in class `WorkingReceiver`. This method is executed 34304 times during the reference execution. If TRIPLEAGENT injects one single exception in this method when it is called for the first time, the application still downloads the file correctly. However, when the perturbation agent keeping injecting exceptions when downloading the file, the application gets stalled.

After analyzing the call stack, TRIPLEAGENT detects 4 candidate failure-oblivious methods, namely `WorkingReceiver/processAndGetNext`, `OutgoingConnectionListener/onNewDataAvail`, `ReadableKeyProcessor/process` and `ConnectionWorker/processSelectedKey`. After two fault injection experiments per candidate failure-oblivious method, TRIPLEAGENT observes that the last three methods are failure-oblivious, the application downloads the file successfully, no matter how many exceptions are thrown in `processAndGetNext`.

In this case, TRIPLEAGENT succeeds in detecting 3 failure-oblivious methods that provide better resilience compared to the normal error-handling code written by the developer.

```
1 String getStringOrNull(Map<...> m, String k) throws
      InvalidBEncodingException {
2 // perturbation point here
3 BEValue v = dictionaryMetadata.get(key);
4 if (v == null) return null;
5 return v.getString();
6 }
```

Listing 5. InvalidBEncodingException in TorrentParser/getStringOrNull

*c) Improving Resilience from Crashing to Resilient:* Let now us consider Listing 5. With a fault injection experiment in method `getStringOrNull` before line 3, TRIPLEAGENT identifies that an exception `InvalidBEncodingException` thrown at this location crashes the whole process. Hence, the perturbation point is a fragile one. After analyzing the stack information, `ClientMain/main` is the default handling method. There are 7 methods including `getStringOrNull` itself before this default handling method, which are all considered as candidate failure-oblivious methods by TRIPLEAGENT.

Then, TRIPLEAGENT performs 2 fault injection experiments for each method according to Algorithm 3, that is $2 \times 7 = 14$ experiments in total. The first experiment assesses whether the candidate failure-oblivious methods could handle only one injected exception. The second assesses whether they could handle as many as injected exceptions. Indeed, TRIPLEAGENT observes that when a catch block is automatically injected in `getStringOrNull`, the application does not crash anymore, and even better, the resulting behaviour is correct (the file is correctly downloaded, its content is the expected one, bit-per-bit). In this case, TRIPLEAGENT has automatically transformed a crashing exception into acceptable behaviour.

---

**Insights from the TTorrent experiment**

Under a realistic workload of downloading a 200MB+ file from the internet, TRIPLEAGENT performs 7780 experiments to evaluate 1046 perturbation points spread over 6.5kLOC. TRIPLEAGENT identifies 642 fragile points, 296 sensitive points and 108 immunized points. After analyzing all 2844 candidate failure-oblivious methods, TRIPLEAGENT confirms that there are 238 failure-oblivious methods in the application. This shows that it is feasible to automatically improve resilience by combining perturbation injection and failure-obliviousness analysis.
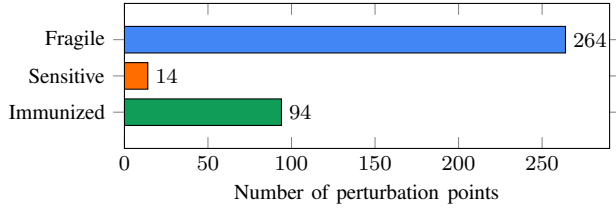
Fig. 4. Category of Perturbation Points under 744 Fault Injection Experiments on HedWig



a) Fragile stays fragile, b) Fragile to sensitive, c) Fragile to immunized

d) Sensitive stays sensitive, e) Sensitive to immunized, f) Immunized stays immunized

Fig. 5. Resilience Improvement on HedWig

## B. Evaluation on HedWig

*1) Experiment Protocol:* HedWig is an email server written in Java, a typical server side application. The main process of HedWig is a perpetual loop, which creates sub-threads to handle different user requests. HedWig relies on a MySql database to store email metadata and saves the email contents as files on the disk. In this experiment, we consider the latest version of HedWig (v0.7).
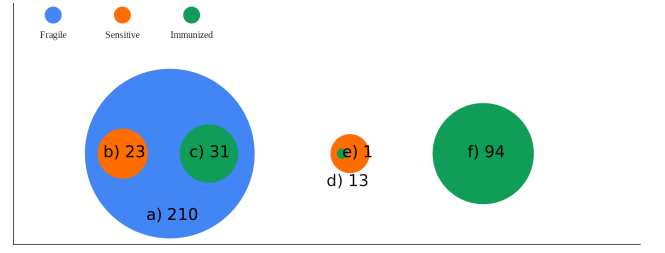
The considered workload is as follows. First, HedWig is deployed on a server. Then TRIPLEAGENT sends a specific email with a unique content to a testing email address using the SMTP protocol. Finally, TRIPLEAGENT logs in with the corresponding account and fetches the same email to do the comparison. The acceptable behaviour is that TRIPLEAGENT both successfully sends and fetches the email, and that the content of this email after final fetching is totally correct.

The experiments are performed sequentially. We note that some of the perturbation experiments crash the email server. In this case, the server is automatically restarted. Some perturbation experiments put it in a corrupted state: to detect this, TRIPLEAGENT adds a checking point after each experiment. All the perturbation agents are switched off and an email is sent and fetched as usual. If the server works correctly the next experiment goes on, otherwise TRIPLEAGENT runs a restart script to bring the server back to normal state.

*2) Experimental Results:* Within the package com/hs/mail TRIPLEAGENT detects 372 perturbation points. Each perturbation point is evaluated by two fault injection experiments: 1) only one exception is injected during the email sending and fetching process, when the point is reached for the first time and 2) exceptions are always injected when the point is reached. Based on these 744 experiments TRIPLEAGENT classifies all the perturbation points using the classification algorithm described in Algorithm 2. Overall, TRIPLEAGENT finds in Hedwig 264 fragile points, 14 sensitive points and 94 immunized points, which are shown in Figure 4

The next step for TRIPLEAGENT is to identify the candidate failure-oblivious methods. By summing over all perturbation points, TRIPLEAGENT detects 722 candidate methods. The minimum, median, maximum number of candidate failure-oblivious methods per perturbation point is respectively 0, 2, 10.

Similar to classifying perturbation points, each candidate failure-oblivious method also needs two fault injection exper-

iments to be evaluated. Finally, 1444 executions are made to evaluate all the candidate failure-oblivious methods based on Algorithm 3. By silencing exceptions in the candidate failure-oblivious methods, TRIPLEAGENT shows that 23 fragile perturbation points can be improved into sensitive ones. 31 fragile points are transformed to immunized ones. It upgrades 1 sensitive perturbation point to an immunized one as well. All those improvements are shown in Figure 5.

Table II shows a sample of interesting perturbation points. It shows different levels of automatic resilience improvement. Similar to Table I, every row describes one perturbation point and one of its corresponding failure-oblivious methods. For example, the first row gives details about perturbation point `queryForLong` in class `AbstractDao`, line number 100. When a `DataAccessException` is thrown from this point, by default it is handled by a try-catch block written by developers in `TransactionTemplate/execute`. But this catch block does not prevent the exception from failing user requests. If the same exception is caught earlier in the stack, in `AnsiMessageDao/getHeaderNameID`, the server is able to bear at least one exception. Note that it is possible to have multiple failure-oblivious methods for the same perturbation point. Such as row 3 and row 4, row 8 and row 9 in the table.

*3) Case Studies:* We now discuss two interesting case studies.

```
1  mailboxRowMapper = new RowMapper<Mailbox>() {
2    public Mailbox mapRow(ResultSet r, int n) throws
         SQLException {
3      // perturbation point here
4      Mailbox mb = new Mailbox();
5      mb.setMailboxID(r.getLong("mailboxid"));
6      ...
7      return mb;
8    }
9  };
```

Listing 6. SQLException in AnsiMailboxDao$1/mapRow

*a) A Failed Failure-oblivious Experiment:* Listing 6 shows a perturbation point found by TRIPLEAGENT in Class `AnsiMailboxDao`, line 3. First, TRIPLEAGENT detects that when one `SQLException` is thrown from this location, the application fails to receive and send the test email. Hence, the perturbation point is a fragile one. By analyzing the call stack, method `mapRow` is considered as a candidate failure-oblivious method.

TRIPLEAGENT automatically wraps the method with a try-catch block. This specific failure-oblivious operation results

8

| 0 | Perturbation Point | Exception Type | Default Handling Method | Failure-oblivious Method | Category |
|---|---|---|---|---|---|
| 1 | AbstractDao/queryForLong@100 | DataAccessException | TransactionTemplate/execute | AnsiMessageDao/getHeaderNameID | fragile - sensitive |
| 2 | ImapServerHandler/handleUpstream@56 | Exception | DCPipeline/sendUpstream | ImapServerHandler/handleUpstream | fragile - sensitive |
| 3 | CountingInputStream/read@21 | IOException | BodySBuilder/build | BodySBuilder/simplePartDescriptor | fragile - immunized |
| 4 | CountingInputStream/read@21 | IOException | BodySBuilder/build | BodySBuilder/createDescriptor | fragile - immunized |
| 5 | MessageHeader/parse@118 | IOException | MessageHeader/<init> | MessageHeader/parse | fragile - immunized |
| 6 | FlagUtils/getFlags@57 | SQLException | JdbcT/doInPreparedStatement | FlagUtils/getFlags | fragile - immunized |
| 7 | PartContentBuilder/build@63 | IOException | FetchRespBuilder/bodyContent | PartContentBuilder/build | fragile - immunized |
| 8 | MailMessage/save@88 | IOException | ToRepository/service | MailMessage/save | sensitive - immunized |
| 9 | MailMessage/save@88 | IOException | ToRepository/service | ToRepository/saveMessage | sensitive - immunized |
| 10 | AliasingForwarding/service@70 | MessagingException | LocalDelivery/service | AliasingForwarding/service | alternative resilient method |
| 11 | ToRepository/deliver@118 | IOException | ToRepository/service | ToRepository/deliver | alternative resilient method |

TABLE II
SAMPLE OF PERTURBATION POINTS AND THE CORRESPONDING FAILURE-OBLIVIOUS METHODS IN HEDWIG

in inserting an incorrect record into the database, which influences the upcoming experiments. Thanks to running the checkpoint procedure described above, TRIPLEAGENT detects this problem, restarts the server, definitely labels this method as non failure-oblivious and excludes this perturbation point for later experiments.

```
1  public int read() throws IOException {
2    // perturbation point here
3    int next = in.read();
4    ...
5    return next;
6  }
```
Listing 7. IOException in CountingInputStream/read

*b) A Perturbation Point with Multiple Failure-oblivious Methods:* In Listing 7, line 2 is a fragile perturbation point in Class `CountingInputStream`. If an `IOException` is thrown from this location, the user is not able to fetch any emails. By default the exception is handled by `BodyStructureBuilder/build(Date d, Long l)`, 5 methods upper in the stack. It means that methods before the default exception handler are all candidate failure-oblivious methods. TRIPLEAGENT evaluates them one by one and verifies 3 out of them. In the call stack, if the exception is silenced in `BodyStructureBuilder/simplePartDescriptor`, `BodyStructureBuilder/createDescriptor` or `BodyStructureBuilder/build(InputStream i)`, the server works properly no matter how many exceptions are thrown. This is a strong improvement to the resilience.

---

**Insights from the HedWig experiment**

Under a production-like email task, TRIPLEAGENT performs 2188 experiments to evaluate 372 perturbation points spread over 13.8 kLOC. 261 fragile points, 68 sensitive points and 43 immunized points are identified in the original code. TRIPLEAGENT assesses that 60 out of 722 methods can be transformed into failure-oblivious methods. This further confirms that TRIPLEAGENT can improve the resilience of a server application in an automated manner.

---

### C. Overhead of TRIPLEAGENT

The overhead of TRIPLEAGENT varies a lot among different perturbation points, failure-oblivious methods. Considering that the ultimate goal of TRIPLEAGENT is to automatically improve resilience, we manually evaluate the overhead

TABLE III
THE OVERHEAD OF AN EXPERIMENT ON TTORRENT

| Evaluation Aspects | Original Version | Instrumented Version | Variation |
|---|---|---|---|
| Downloading time | 20.4s | 21.1s | 3.5% |
| CPU time | 15.0s | 18.3 | 22.2% |
| Memory usage | 47M | 49M | 4.3% |
| Peak thread count | 30 | 32 | 6.7% |
| Relevant class files size | 16.7KB | 16.8KB | 0.6% |

of failure-oblivious experiments. The overhead caused by TRIPLEAGENT is evaluated in 3 aspects: 1) at the application level, the execution time is compared. In TTorrent this metric means the downloading time. In HedWig experiments this means the time spent on sending and receiving the email. 2) at the operating system level, the CPU and memory usage, peak thread count are taken into consideration. 3) at the binary code level, the code bloat due to instrumentation is evaluated. For statistical purposes, we conduct the same measurement 30 times and calculate the average [1].

For TTorrent experiments, Table III records the overhead of verifying failure-oblivious method `HTTPTrackerClient/sendAnnounce` in Table I, row 5. For HedWig experiments, failure-oblivious method `BodySBuilder/simplePartDescriptor` in Table II, row 3 is taken as an example. The overhead of execution time, CPU time, memory usage, peak thread count, relevant class files size are respectively 0%, 6.0%, 0%, 5.4%, 3.0%.

The reason why TRIPLEAGENT has such a low overhead is that the instrumentation is small. The perturbation agent and failure-oblivious agent only instrumented one or two class files. Meanwhile, the monitoring agent does not cause high overhead thanks to the JVMTI framework. By evaluating the overhead, developers are more confident about the resilience improvement suggested by TRIPLEAGENT.

## V. DISCUSSION

*Fault model.* Currently TRIPLEAGENT considers two fault models. In both models, exceptions are injected at a single location. But there also exists common mode failures which involve a series of different exceptions. An exception could also be mixed with data errors. Devising and implementing fault models that stimulate common mode failures or data errors is an interesting direction for future work.

*Workload impact.* A threat to the validity of our experiments comes from the workload. TRIPLEAGENT takes a production-like workload to exercise the application code.

When TRIPLEAGENT identifies a failure-oblivious method, it guarantees that it works under the tested workload. But it may break some behavior with a more comprehensive workload. Overall, the more diverse the workload is, the more confidence TRIPLEAGENT has in the found failure-oblivious methods.

*Scalability.* During our experiments, the deepest call stack considered was composed of 39 methods. TRIPLEAGENT has not been tested with larger applications with deeper stacks. As such, the full scalability of TRIPLEAGENT is not yet verified. We note that the number of candidate failure-oblivious methods to assess is linear in the depth of the stack, which means that, in theory, TRIPLEAGENT would be scalable.

## VI. RELATED WORK

Now we discuss the related work along three aspects.

*1) Fault injection:* Fault injection is a widely-researched topic in software dependability. In the 1990s, the research was mostly about hardware implemented fault injection tools. For example, Madeira et al. [21] invented RIFLE, a pin-level fault injector to generate processor errors. Next, more software-based fault injection tools were invented. Kanawati et al. [18] proposed FERRARI, a tool for the emulation of hardware faults and control flow errors. Han et al. [16] designed DOCTOR, a tool for injecting hardware failures and network communication failures. Wei et al. [35] built a software-based hardware faults injector called LLFI, and quantitatively compared the accuracy of fault-injection with assembly code level injector PINFI. Lee et al. [20] presented SFIDA, a tool to test the dependability of distributed applications on the Linux platform. Kao et al. [19] invented "FINE", a fault injection and monitoring tool to inject both hardware-induced software errors and software faults. Kouwe and Tanenbaum [34] presented HSFI, a fault injection tool that injects faults with context information from source code and applies fault injection decisions efficiently on the binary.

Fu et al. [12] presented an approach to measure the coverage of recovery code with respect to operating system and I/O hardware faults. The common idea with TRIPLEAGENT is to inject exceptions to trigger error handling code. Yet, our and their goal are notably different. Fu el al. use fault injection to increase recovery code coverage. TRIPLEAGENT combines fault injection with failure-oblivious computing to improve resilience.

The novelty of TRIPLEAGENT is that it is designed to inject application-level exceptions (and not hardware faults) in Java applications. TRIPLEAGENT gives developers concrete insights at the source code level about their exception-handling implementation.

*2) Self-healing software:* Self-healing software follows the idea that it is possible to automatically make software recover from failures. Different techniques have been applied to achieve this goal, such as automatic reboot, checkpoint-restart, and failure-oblivious transformation.

Reboot techniques [4], [17], [33] require the system to be able to restart, which may bring some down-time. Checkpoint-restart techniques significantly reduce the recovery time by saving and reloading runtime states saved at checkpoints. Qin et al. [27] invented Rx, which enables the program to rollback to a recent checkpoint upon a software failure, and then to re-execute in a modified environment. Sidiroglou et al. [31] proposed ASSURE, a system that introduces rescue points to recover from unknown faults.

Regarding failure-oblivious computing, Rinard et al. [30] invented a safe compiler for C to enable servers to execute despite memory errors. Perkins et al. [26] proposed ClearView, a system for automatically patching errors in deployed software. It observes values of registers and memory locations and tries to detect violations of invariants at this level. Rigger et al. [28] presented an approach that allows C programmers to perform explicit sanity checks and to react according to invalid arguments or states. They also designed a C dialect called Lenient C [29] that checks undefined behaviours in the C standard including memory management, pointer operations and arithmetic operations.

None of these tools combine fault injection and failure-oblivious computing together as we do in TRIPLEAGENT. They do not actively inject failures into the system, nor do they conduct application-level analysis to detect valuable failure-oblivious positions.

*3) Exception analysis:* Byeong-Mo et al. [5] gave a comprehensive review on exception analysis. Magiel Bruntink et al. [3] proposed a characterization and evaluation method to statically discover faults in exception handling. Fu and Ryder [13] described a static analysis method for exception chains in Java. Martins et al. [22] presented VerifyEx to test Java exceptions by inserting exceptions at the beginning of try blocks. Zhang and Elbaum [37] presented an approach that amplifies tests to validate exception handling. Cornu et al. [7] proposed a classification of try-catch blocks at testing time.

Those tools rely on test suites to analyze resilience with respect to error-handling. On the contrary, TRIPLEAGENT analyzes the system behaviour based on user-level traffic and usages.

## VII. CONCLUSION

In this paper, we have presented TRIPLEAGENT, a system which combines automated monitoring, automated perturbation injection and automated resilience improvement. By evaluating TRIPLEAGENT on two real-world Java applications, we have shown that it is able to detect weaknesses in exception-handling of Java code and to improve resilience. In the future, we will further explore the design space of perturbation and failure-obliviousness strategies. For instance, we would like to inject timeout on requests and interactions in asynchronous software. Our long-term goal is to use TRIPLEAGENT in production, and consequently, we will also keep reducing the overhead of TRIPLEAGENT at runtime.

REFERENCES

[1] Andrea Arcuri and Lionel C. Briand. A practical guide for using statistical tests to assess randomized algorithms in software engineering. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu , HI, USA, May 21-28, 2011*, pages 1–10, 2011.

[2] A. Basiri, N. Behnam, R. de Rooij, L. Hochstein, L. Kosewski, J. Reynolds, and C. Rosenthal. Chaos engineering. *IEEE Software*, 33(3):35–41, May 2016.

[3] Magiel Bruntink, Arie van Deursen, and Tom Tourwé. Discovering faults in idiom-based exception handling. In *ICSE*, pages 242–251. ACM, 2006.

[4] George Candea and Armando Fox. Crash-only software. In *Proceedings of HotOS'03: 9th Workshop on Hot Topics in Operating Systems, May 18-21, 2003, Lihue (Kauai), Hawaii, USA*, pages 67–72, 2003.

[5] Byeong-Mo Chang and Kwanghoon Choi. A review on exception analysis. *Information & Software Technology*, 77:1–16, 2016.

[6] Michael Alan Chang, Bredan Tschaen, Theophilus Benson, and Laurent Vanbever. Chaos monkey: Increasing sdn reliability through systematic network destruction. *SIGCOMM Comput. Commun. Rev.*, 45(4):371–372, August 2015.

[7] Benoit Cornu, Lionel Seinturier, and Martin Monperrus. Exception Handling Analysis and Transformation Using Fault Injection: Study of Resilience Against Unanticipated Exceptions. *Information and Software Technology*, 57:66–76, January 2015.

[8] Benjamin Danglot, Philippe Preux, Benoit Baudry, and Martin Monperrus. Correctness attraction: a study of stability of software behavior under runtime perturbation. *Empirical Software Engineering*, 23(4):2086–2119, 2018.

[9] Thomas Durieux, Youssef Hamadi, Zhongxing Yu, Benoit Baudry, and Martin Monperrus. Exhaustive exploration of the failure-oblivious computing search space. In *11th IEEE International Conference on Software Testing, Verification and Validation*, pages 139–149, 2018.

[10] Felipe Ebert, Fernando Castor, and Alexander Serebrenik. An exploratory study on exception handling bugs in java programs. *Journal of Systems and Software*, 106:82–101, 2015.

[11] Bent Flyvbjerg. Five misunderstandings about case-study research. *Qualitative inquiry*, 12(2):219–245, 2006.

[12] C. Fu, R. P. Martin, K. Nagaraja, D. Wonnacott, T. D. Nguyen, and B. G. Ryder. Compiler-directed program-fault coverage for highly available java internet services. In *2003 International Conference on Dependable Systems and Networks, 2003. Proceedings.*, pages 595–604, June 2003.

[13] Chen Fu and Barbara G. Ryder. Exception-chain analysis: Revealing exception handling architecture in java server applications. In *Proceedings of the 29th International Conference on Software Engineering*, ICSE '07, pages 230–239, Washington, DC, USA, 2007. IEEE Computer Society.

[14] Sudipto Ghosh and John L. Kelly. Bytecode fault injection for java software. *J. Syst. Softw.*, 81(11):2034–2043, November 2008.

[15] James Gosling, Bill Joy, Guy L. Steele, Gilad Bracha, and Alex Buckley. *The Java Language Specification, Java SE 8 Edition*. Addison-Wesley Professional, 1st edition, 2014.

[16] Seungjae Han, K. G. Shin, and H. A. Rosenberg. Doctor: an integrated software fault injection environment for distributed real-time systems. In *Proceedings of 1995 IEEE International Computer Performance and Dependability Symposium*, pages 204–213, April 1995.

[17] Yennun Huang, Chandra M. R. Kintala, Nick Kolettis, and N. Dudley Fulton. Software rejuvenation: Analysis, module and applications. In *Digest of Papers: FTCS-25, The Twenty-Fifth International Symposium on Fault-Tolerant Computing, Pasadena, California, USA, June 27-30, 1995*, pages 381–390, 1995.

[18] Ghani A. Kanawati, Nasser A. Kanawati, and Jacob A. Abraham. FERRARI: A tool for the validation of system dependability properties. In *FTCS*, pages 336–344. IEEE Computer Society, 1992.

[19] W. I. Kao, R. K. Iyer, and D. Tang. Fine: A fault injection and monitoring environment for tracing the unix system behavior under faults. *IEEE Transactions on Software Engineering*, 19(11):1105–1118, Nov 1993.

[20] Hyosoon Lee, Youngshik Song, and Heonshik Shin. Sfida: a software implemented fault injection tool for distributed dependable applications. In *Proceedings Fourth International Conference/Exhibition on High Performance Computing in the Asia-Pacific Region*, volume 1, pages 410–415 vol.1, May 2000.

[21] Henrique Madeira, Mário Zenha Rela, Francisco Moreira, and João Gabriel Silva. RIFLE: A general purpose pin-level fault injector. In *Dependable Computing - EDCC-1, First European Dependable Computing Conference, Berlin, Germany, October 4-6, 1994, Proceedings*, pages 199–216, 1994.

[22] Alexandre Locci Martins, Simone Hanazumi, and Ana Cristina Vieira de Melo. Testing java exceptions: An instrumentation technique. In *COMPSAC Workshops*, pages 626–631. IEEE Computer Society, 2014.

[23] Martin Monperrus. Automatic software repair: A bibliography. *ACM Comput. Surv.*, 51(1):17:1–17:24, 2018.

[24] Roberto Natella, Domenico Cotroneo, and Henrique S. Madeira. Assessing dependability with software fault injection: A survey. *ACM Comput. Surv.*, 48(3):44:1–44:55, February 2016.

[25] Konstantinos Parasyris, Georgios Tziantzoulis, Christos D. Antonopoulos, and Nikolaos Bellas. Gemfi: A fault injection tool for studying the behavior of applications on unreliable substrates. In *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2014, Atlanta, GA, USA, June 23-26, 2014*, pages 622–629, 2014.

[26] Jeff H. Perkins, Sunghun Kim, Samuel Larsen, Saman P. Amarasinghe, Jonathan Bachrach, Michael Carbin, Carlos Pacheco, Frank Sherwood, Stelios Sidiroglou, Greg Sullivan, Weng-Fai Wong, Yoav Zibin, Michael D. Ernst, and Martin C. Rinard. Automatically patching errors in deployed software. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles 2009, SOSP 2009, Big Sky, Montana, USA, October 11-14, 2009*, pages 87–102, 2009.

[27] Feng Qin, Joseph Tucek, Yuanyuan Zhou, and Jagadeesan Sundaresan. Rx: Treating bugs as allergies - a safe method to survive software failures. *ACM Trans. Comput. Syst.*, 25(3):7, 2007.

[28] Manuel Rigger, Rene Mayrhofer, Roland Schatz, Matthias Grimmer, and Hanspeter Mössenböck. Introspection for C and its applications to library robustness. *Programming Journal*, 2(2):4, 2018.

[29] Manuel Rigger, Roland Schatz, Matthias Grimmer, and Hanspeter Mössenböck. Lenient execution of C on a java virtual machine: or: How I learned to stop worrying and run the code. In *Proceedings of the 14th International Conference on Managed Languages and Runtimes, ManLang 2017, Prague, Czech Republic, September 27 - 29, 2017*, pages 35–47, 2017.

[30] M. Rinard, C. Cadar, D. Dumitran, D.M. Roy, T. Leu, and W.S. Beebee Jr. Enhancing server availability and security through failure-oblivious computing. In *Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation*, pages 21–21. USENIX Association, 2004.

[31] Stelios Sidiroglou, Oren Laadan, Carlos Perez, Nicolas Viennot, Jason Nieh, and Angelos D. Keromytis. ASSURE: automatic software self-healing using rescue points. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2009, Washington, DC, USA, March 7-11, 2009*, pages 37–48, 2009.

[32] Nuno Silva, Ricardo Barbosa, João Carlos Cunha, and Marco Vieira. A view on the past and future of fault injection. In *2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Budapest, Hungary, June 24-27, 2013*, pages 1–2, 2013.

[33] Mark Sullivan and Ram Chillarege. Software defects and their impact on system availability: A study of field failures in operating systems. In *Proceedings of the 1991 International Symposium on Fault-Tolerant Computing, Montreal, Canada*, pages 2–9, 1991.

[34] Erik van der Kouwe and Andrew S. Tanenbaum. HSFI: accurate fault injection scalable to large code bases. In *46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2016, Toulouse, France, June 28 - July 1, 2016*, pages 144–155, 2016.

[35] Jiesheng Wei, Anna Thomas, Guanpeng Li, and Karthik Pattabiraman. Quantifying the accuracy of high-level fault injection techniques for hardware faults. In *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2014, Atlanta, GA, USA, June 23-26, 2014*, pages 375–382, 2014.

[36] Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, Yongle Zhang, Pranay Jain, and Michael Stumm. Simple testing can prevent most critical failures: An analysis of production failures in distributed data-intensive systems. In *11th USENIX Symposium on Operating Systems Design and Implementation*, pages 249–265, 2014.

[37] Pingyu Zhang and Sebastian Elbaum. Amplifying tests to validate excep-

tion handling code: An extended study in the mobile application domain. *ACM Trans. Softw. Eng. Methodol.*, 23(4):32:1–32:28, September 2014.

[38] Haissam Ziade, Rafic A. Ayoubi, and Raoul Velazco. A survey on fault injection techniques. *Int. Arab J. Inf. Technol.*, 1(2):171–186, 2004.