

# Automatic Software Repair: a Bibliography

Martin Monperrus  
University of Lille & Inria  
martin.monperrus@univ-lille1.fr

Technical report #hal-01206501

Version of Saturday 21<sup>st</sup> November, 2015. Critical feedback appreciated!  
Thanks, –Martin

Latest version: <http://www.monperrus.net/martin/survey-automatic-repair.pdf>

## Abstract

This article presents an annotated bibliography on automatic software repair. Automatic software repair consists of automatically finding a solution to software bugs, without human intervention. The uniqueness of this article is that it spans the research communities that contribute to this body of knowledge: software engineering, dependability, operating systems, programming languages and security. Furthermore, it provides a novel and structured overview of the diversity of bug oracles and repair operators used in the literature.

## 1 Introduction

This paper presents an annotated bibliography on automatic software repair. Automatic software repair consists of automatically finding a solution to software bugs, without human intervention. This idea of automatically repairing software bugs is both important and challenging. It is important because software has eaten the world<sup>1</sup>, but unfortunately each bite comes with bugs. The software we daily use sometimes crashes, sometimes gives erroneous results, and sometimes even kills people [81]. We do have millions of bugs in the wild, and many of them are being created every day in the new software products and releases we ship in production. To sum up, if automatic software repair could only repair a fraction of those bugs, it would bring value to society and humanity<sup>2</sup>.

Automatic software repair is challenging because fixing bugs is a difficult task. Of course there are stupid bugs – “blunder” as Knuth puts it [77] –

---

<sup>1</sup>paraphrasing Silicon Valley’s entrepreneur Marc Andreessen

<sup>2</sup>including economical value, but that’s less important

that can be trivially fixed. However, any programmer, whether professional or hobbyist, remembers a bug that took her hours, if not days and weeks to be understood and fixed, these are the “hairiest bugs” [39]. For those bugs, automatic repair is a challenging human-competitive task.

The goal of this paper is to draw the big picture of automatic software repair. In particular, I aim at presenting two families of automatic repair techniques: behavioral repair and state repair. The former is about automatically modifying the program code; the latter is about automatically modifying the execution state at runtime. The primary intended audience is researchers in computer science, with a focus on the research communities that contribute to this body of knowledge: software engineering, dependability, operating systems, programming languages and software security. Each section also provides an introductory explanation of the key concepts behind automatic repair, which could be of high interest for practitioners and curious students. The margin notes are here to help their reading, and to help the usage of this document as course note.

As any bibliographic survey, there is a will of completeness: I aim at covering all important works of the field published at the time of writing those lines (Fall 2015). The inclusion criterion is my own appreciation of novelty and significance. If several papers contain the same idea, I select only one, the best one, esp. if they come from the same authors. But the literature is vast, and I have surely forgotten important works<sup>3</sup>. In particular, the main pitfall I encountered for this bibliography was the terminology: the same concept “repair” has several names in the literature: patch, fix, heal, recover, etc. Table 1 lists the main ones, as well as example notable references that use the term. In my work and this paper, I use the term “repair”. I prefer “repair” because a program has something mechanical in nature, which fits well the daily usage of the word “repair”. Also, it is the name used by most excellent papers in the field.

To my knowledge, there is no comparable bibliography in the literature. The look-back paper by Le Goues et al. [85] is close but only covers a fraction of the papers, and only on behavioral repair. On the contrary, Rinard [133] only focused on runtime repair. Yet, there are surveys in related fields, for instance for fault-tolerance [164], fault localization [138, 173], algorithmic debugging [147] to only name a few.

The remainder of this paper reads as follows. Section 2 briefly presents the core concepts of automatic repair. Section 3 discusses the main approaches of behavioral repair and Section 4 is about state repair. Section 5 is dedicated to the empirical works that aim at understanding the foundations of automatic repair. Section 6 is an account on papers that are not directly about automatic repair yet have a close connection.

---

<sup>3</sup>I apologize in advance for the missing references, and I would be grateful if you point them to me

Expression	Example Ref.
automatic repair (program repair, self-repair)	[84, 94, 122, 112]
automatic fixing (bug fixing, program fixing)	[91, 170]
automatic patching	[171, 142, 90, 76]
healing (self-healing)	[50, 140, 143]
automatic correction (self-correcting)	[83, 74]
automatic recovery (self-recovering)	[154, 18]
resilience	[140, 26]
automatic workaround	[22]
survive (survival, survivability)	[109, 127, 133]
rejuvenation	[59]
biological metaphors: allergies, immunity, vaccination	[127, 144, 93, 67]

Table 1: The diverse terminology of automatic software repair

## 2 Core Concepts for Automatic Repair

Automatic repair is about bugs. The literature is full of synonyms for “bug”: defect, fault, error, failure, mistake, etc. There are rather accepted definitions between faults, errors and failures [9]: a *failure* is an observed unacceptable behavior; an *error* is a propagating incorrect state prior to the failure (without yet having being noticed); a *fault* is the root cause of the error (in particular incorrect code). Although the relative clarity of those three concepts, one can hardly say that the literature, incl. the most recent papers, sticks to those definitions. Furthermore, if we only consider the repair literature, there is absolutely no emerging separation between “automatic repair of failures”, “automatic repair of errors” and “automatic repair of faults”. Since, we need a common concept for all words, I like the term “bug” as an umbrella word because of its intuitiveness and wide usage. For this paper, I use the following definition. *A bug is a deviation between the expected behavior of a program execution and what it actually happened.*<sup>4</sup>

This definition of bug involves the notions of “behavior”, “execution”, “program” but has an implicit third subject: the observer, or reference point, that deems the behavior unexpected. This “observer” can obviously be a human user saying “this output is not correct”. It is also classically a *specification*, in its most general meaning: a specification is a set of expected behaviors. Specifications are polymorphic: they can be natural language documents, formal logic formulas, test suites, etc. They can even be implicit: for instance the specification “the program shall not crash on any input” holds for many programs while not often being explicitly written. To some extent, the user saying “this output is not correct”, is stating the specification on the fly. Consequently, automatic repair always refers to a specification and yields the following def-

<sup>4</sup>Note that some authors use “intended” instead of “expected”, I prefer the latter because it’s really the viewpoint of the user or client that matters, not the viewpoint of the engineer who designed and developed the software.

inition of automatic repair. *Automatic repair is the transformation of an unacceptable behavior of a program execution into an acceptable one according to a specification.*

oracle      A concept that is close to the one of specification is the one of *oracle*. Simply put, an oracle determines whether the result of executing a program is correct [155]. To this extent, specification and oracle refer to the same thing: expectation, acceptability, correctness. However, there is a major difference between both. An oracle is only a part of specifications, it is the part related to the expected output (when one such exists). In addition, a specification contains information about the input ranges, about non-functional properties, etc. For instance, a test suite is a specification, it contains test cases, which themselves contain assertions, the latter being the oracles.

bug oracle  
regression oracle      With respect to repair, the oracles can be split in two: the *bug oracle* refers to the oracle that detects the unexpected behaviors; the *regression oracle* refers to the oracles that check that no new bugs have been introduced during repair. The reason is that the program upon repair already satisfies all regression oracles, but a repair transformation may accidentally introduce a regression. There are more formal definitions of specification and oracle in the literature [155, 12] but they do not bring much in the context of this paper.

bug class      Finally, a repair technique often targets a *bug class*<sup>5</sup>. A bug class is an abstract concept referring to a family of bugs that have something in common: the same symptoms, the same root cause, the same solution [111]. For instance, well known bug classes include off-by-one errors, memory leaks, etc. However, there are many bug classes for which there are no clear definition and scope in the literature, and some of them even miss a name. While some initial taxonomies exist [166, 37], building a comprehensive taxonomy of bug classes will require years of research.

### 3 Behavioral Repair

Behavioral repair consists of changing the behavior of the program under repair, i.e. changing its code. The modification can be done on source code, but also on binary code (e.g. Java bytecode or x86 native code). Behavioral repair can be done offline or online at runtime.

When done offline, behavioral repair may happen in the development environment (IDE) of maintenance developers or in a continuous integration server. When done at runtime, it can be at testing time or in production. Technically, behavioral repair at runtime involves a kind of dynamic software update (DSU), which is a research topic per se.

---

<sup>5</sup>or fault class, or error class, etc.

Operator	Example Ref.
add/remove/replace code	[171, 6]
add a precondition	[31, 90]
replace a condition	[31, 115]
replace assignment RHS	[52, 115, 69]
addition or removal of method calls	[28]
adding a modulo for array read, truncating data for array write	[90]

Table 2: Examples of repair operators for behavioral repair

repair operator Behavioral repair involves a *repair operator*<sup>6</sup> which is a kind of modification on the program code. For instance, one repair operator is the addition of a precondition, as shown below.

```
+ if (age >= 18)
    serve_adult_content()
```

repair template The literature defines many different repair operators, that will be presented below and that are summarized in Table 2. Sometimes the repair operator involves a *repair template*<sup>7</sup>, which is a parameterized snippet of code which targets the repair a specific bug class. A *repair model* [103] is a set of repair operators.

repair model For instance, when considering a test suite as specification 3.1.1, a problem statement of behavioral repair is *given a program and its test suite with at least one failing test case, create a patch that makes the whole test suite passing*. This problem statement can be called test-suite based repair [111], and has been famously explored by Genprog, presented in Section 3.1.1

## 3.1 Repair & Oracles

As presented in Section 2, automatic repair is with respect to an oracle. I choose to order the related work by the kind of oracle they consider.

### 3.1.1 Test Suites

A test suite is an input-output based specification. In modern object-oriented software, the input can be as complex as a set of interrelated objects built with a rich sequence of method calls, and the output can also be a sequence of method calls that observe the execution state and behavior in various ways. In test-suite based repair, the failing test case acts as a bug oracle, the remaining passing test cases act as a regression oracle.

Genprog Genprog is a seminal and archetypal test-suite based repair system developed at the University of Virginia [171, 172, 42]. Genprog uses three repair

<sup>6</sup>or “repair action”

<sup>7</sup>or “repair strategy” or “fix schema” [170]

operators that are mutations over the abstract syntax trees (AST): deletion of AST nodes addition of AST nodes, and replacement of existing nodes. For addition and replacement, the nodes are taken from elsewhere in the code base. This is called the redundancy assumption [104, 11]. Genprog is able to handle real-world large scale C code. The largest evaluation of Genprog [84] claims that 55 out of 105 bugs can be fixed by Genprog. Those results have been later questioned, as discussed in 5. The Genprog thread of ideas yielded other papers in the original team [137, 53] and other laboratories [123, 125].

EXP Much before Genprog, in the mid 90ies, Stumptner and Wotawa [156] have proposed automatic repair in a simple toy language called EXP. The specification is a set of test cases (i.e. a test suite). To my knowledge, it is the first occurrence of test-suite based repair in the literature.

Jaff Arcuri [7, 5, 6] defines 7 repair operators based on abstract syntax tree modification. For instance, for “promote mutation”, a node is replaced by one of its child. The operators are stacked in a random way. The prototype implementation, called Jaff, handles a subset of Java and is evaluated on toy programs.

Mutation Debroy and Wong [30, 117] propose to use standard mutations from the mutation testing literature to fix programs. Consequently, their repair models are: replacement of an arithmetic, relational, logical, increment/decrement, or assignment operator by another operator from the same class; decision negation in an if or while statement. Conventionally, they locate fault statements with a spectrum based fault localization technique, Tarantula [66]. Nica et al. [117] also use mutations for repair. Compared to Debroy and Wong, they comprehensively explore the space of all mutations.

Metaprogram The key idea of Kern and Esparza [74] is to generate a meta-program that integrates all possible mutations according to a mutation operator. The mutations that are actually executed are driven by meta-variables. A repair is a set of values for those meta-variables. The meta-variables are valued using symbolic execution.

Semfix NGuyen et al. [115] proposed an approach for repair based on symbolic execution and code synthesis. The location of the repair is found with angelic debugging [23], then the repaired expression is synthesized with input-output component synthesis [63]. The repaired locations are right hand side (RHS) of assignments and boolean conditionals, the synthesized expressions mix arithmetic and first-order logics. The implementation works on C code.

PAR The PAR system [76] is an approach for automatically fixing bugs of Java code. PAR is based on repair templates: each of PAR’s ten repair templates represents a common way to fix a common kind of bug. For instance, a common bug is the access to a null pointer, and a common fix of this bug is to add a nullness check just before the undesired access: this is template “Null Pointer Checker”. Some templates are parameterized by variables, for instance the “Null Pointer Checker” template takes a variable name as parameter. The templates are applied and tested in a random search manner.

Nopol

Nopol [31] targets a specific fault class: conditional bugs. It repairs programs by either modifying an existing if-condition or adding a precondition (aka. a guard) to any statement or block in the code. The modified or inserted condition is synthesized via input-output based code synthesis with SMT [63] and predicate switching [175]. The Nopol system has been extended for also repairing infinite loops [102].

Relifix Tan and Roychoudhury proposed Relifix, a repair system dedicated to fixing regression bugs [161]. The approach consists of 8 repair templates, some being transformation operators, the other being parameterized repair templates. The key idea of Relifix is that the templates application are driven by the past changes, for instance, template “add statement” only add statements that were involved in the previous commits related to the regression.

DirectFix Mechtaev et al. also perform test-suite based repair [105], with the noble goal of synthesizing simple patches. In order to do so, they assume a very specific kind of programs: those that can be expressed as trace formulas (related to boolean programs of [54]). Under this assumption, they can state the repair problem as a Maximum Satisfiability (MaxSAT) problem, where the smallest patch is the one that satisfies the most constraints.

SPR SPR [98] defines a set of staged repair operators so as to early discard many candidate repairs that cannot pass the supplied test suite. This allows for exhaustively exploring a small and valuable search space.

CodePhage The idea of Codehage [145] is to transfer a check from one application to another application to avoid crashes. The system assumes an error-triggering input that crashes one application but not the other one. The considered errors are out of bounds access, integer overflow, and divide by zero errors. The missing check is inferred from a symbolic expression over the input fields and validated by a regression test suite.

SearchRepair Ke and colleagues proposes SearchRepair [73], a system inspired from code search. SearchRepair first indexes code fragments as SMT constraints, then at repair time, a fragment is retrieved by combining the desired input-output pairs and the fragments in a single constraint problem. The system is evaluated on small C programs written by students in an online course.

Prophet Prophet [97] is a repair system that uses past commits to drive the repair. What is learned on past commits from version control systems is a probability distribution over a set of features of the patch. This probability distribution is then used to both speed up the repair and increase the likelihood to find correct patches. The evaluation is done on 69 real world defects from the Genprog benchmark, and shows that 15 correct repairs are found.

### 3.1.2 Pre- and Post Conditions

Some works use classical pre- and post-conditions *à la* design-by-contract [110] as oracle for repair.

He and Gupta [56] use pre- and post-conditions to compute “hypothesized program states” (from the post condition) and “actual program states” (from

the failing input). The repair operators consist of changing the LHS or RHS of assignments, or changing a boolean condition with simple modifications (change variable, change relational operator) so that the hypothesized program state becomes compatible with the actual program state. A classical test suite is used for detecting regressions.

AutoFix-E      AutoFix-E is an approach by Wei et al. [170, 174], it generates fixes for Eiffel programs, relying on contracts (pre-conditions, post-conditions, invariants). AutoFix-E uses four repair templates that consist of a snippet and an empty conditional expression to be synthesized. The key intuition behind AutoFix-E is that both the snippet code and the conditional expression are taken from the existing contracts.

Alloy for repair      Gopinath et al. [52] uses pre- and post-conditions written in the Alloy specification language. The function body is also translated to Alloy formulas. Then, the bounded verification mechanism of Alloy is used both to detect bugs (similar to [60]) and to identify the repair. The repair operators are changing the RHS of assignments and modifying existing if-conditions.

Könighofer and Bloem [78] considers assertions as specifications in programs that can be translated to SMT. The approach is static and the repair is shown to not violate the assertion for the considered input domain. The approach is based on repair templates, such as changing the RHS of assignments or changing an arithmetic expression by a linear combination. The templates holes are filled by the SMT solver.

### 3.1.3 Abstract Behavioral Models

An abstract behavioral model, such as a state machine encoding the object state and the corresponding allowed method calls can be used to drive the repair.

In 2006, before Genprog, Weimer [] proposed a first patch generation technique. It requires as input a safety policy (i.e. a typestate property or an API usage rule) and the control-flow graph of a method. The whole approach is static: the bug is detected as a static violation of the safety property, and the correctness condition of the patch is only to pass the safety check. Interestingly, the word does not mention the term “repair”, it was not in the Zeitgeist at this time.

Pachika      Dallmeier et al. [28] presented Pachika, an approach for repairing Java programs. The idea of Pachika is to first infer an object usage model from executions, and then to generate a fix for failing runs in order to match the inferred expectedly correct behavior. The evaluation consists of fixing 18 bugs of ASPECTJ (75KLOC) and 8 of RHINO (38KLOC). The two repair operators of Pachika are addition and removal of method calls. The main difference with the previous approach is that the behavioral model is mined, and not given.

## 3.2 Static Analysis

Static analysis tools outputs errors and warnings. It is possible to automatically repair them. In this case, the correctness oracle is the static analysis itself.

.NET Logozzo and Ball [94] proposes a repair approach on top of their static analysis toolchain for .Net code. For a set of fault class identified statically (e.g. off-by-one errors), they propose a corresponding repair operations. The repair operators are specific to each fault class, for instance, it is adding a precondition, changing the size of an array allocation, etc. The static analysis is run again to verify the correctness of the repair.

Arithmetic overflow Logozzo and Martel [95] targets a specific fault class in integer arithmetic (linear combinations). The arithmetic overflow is detected statically, and the suggested fix is a re-ordering of the arithmetic operations. The fix ensures that that the overflow cannot happen anymore. On arithmetic overflows, there is also the work by Cocker et al. [25].

Buffer overflow Muntean etl al. [113] statically detects buffer overflows. Then they have templates parameterized by a variable. The correct variable to be used in the template is found using SMT.

## 3.3 Crashing inputs

Behavioral repair can happen as a response to a field failure (e.g. a crashing exception or a SegFault caused by a buffer overflow). The repair process happens once the crashing input has been identified and minimized if possible. The failing test case of a test suite can also be seen as a crashing input. However, the main difference or crashing inputs and test suites from the viewpoint of oracle for repair is the following. A test suite also contains passing test cases (the regression oracle), and that failing test case contains assertions on the expected value, while crashing inputs, as their name suggests, only refer to a violation of the non-functional contract “the program shall not crash”.

Overflows Sidoroglou and Keromytis [142] detect buffer overflow vulnerabilities at runtime in production, then they obtain the source of the vulnerability through the use of ProPolice [41]; finally, they use code transformation rules written in the transformation language TXL to modify source code. Regressions are caught by manually provided test suites.

AutoPag Lin et al. [90] tries to generate a source code patch from a working exploit that triggers an array overflow in C code. Its repair operators consist of fixing out-of-bound reads by adding a modulo in the read expression and out-of-bound writes by truncating data to be written (similarly to failure-oblivious computing).

AFix Jin et al. [64] present an approach to fix concurrency bugs. The bug oracle is a dynamic bug detector that triggers concurrency bugs such as deadlocks by cleverly injecting delays. The repair model of AFix consists of putting instructions into critical regions. This work on automatic repair of concurrency bugs has been further extended [92].

- Android Azim et al. [10] detect field failures on Android smartphone applications. The considered faults are unhandled exceptions, the repair operator consists of adding try/catch blocks with binary rewriting.
- SoupInt Wang et al. [168] target automatic repair of integer overflows. They have three repair operators. The first one is to force taking an error branch before the overflow happens, the second one is to force taking an error branch after the overflow has happened, and the last one is a program stop (exit). The generated conditions are path conditions obtained from dynamic symbolic execution.

### 3.4 Other Oracles

- Other specific oracles have been used in an automatic repair setting.
- malformed HTML Samimi et al. [136] have presented an approach for repairing PHP code which generates HTML tags. The oracle that is used is whether the output HTML string is malformed, i.e. that it does not contain a inconsistent sequence of opening and closing tags (e.g. “<a></i></a>”). They encode the repair as a constraint problem on strings.
- R2Fix Liu et al. [91] uses as oracle a manually written bug report. They have parameterized repair templates and extract the actual value of the template parameter from the bug report. For instance, for a not-null checker template, they extract the name of the variable to be checked from the bug report.
- Proof Dennis et al. [35] uses proof-based program verification on ML programs using Isabel as oracle. When the proof fails, the counter-example of the proof drives a repair approach based on repair templates (replacing one method call by another, adding some code).
- Reference implementation It is possible to use a reference implementation as specification for repair. In this case, the reference implementation both acts as the bug oracle (when the behavior of the reference implementation and of the buggy program do not correspond) and as a regression oracle. This has been little explored in the context of repair. The approach by Könighofer and Bloem [79] uses SMT-based templates. The approach by Singh et al. [149] is conceptually similar but is realized differently and the evaluation is much larger. The reference implementation and the program to be repaired are written in Python. The system translates them to a programming environment called Sketch, which is responsible for exploring the space of candidate fixes. The evaluation is made on thousands of buggy programs submitted for an online course.

### 3.5 Domain Specific Repair

The concept of automatic repair can be applied on many computational artifacts. Indeed, there are many works doing automatic repair in contexts that are specific to an application domain.

Lazaar et al. [83] repair constraint programs. With a domains-specific fault localization strategy, the repair consists of removing or adding new constraints.

Gopinath et al. [51] repair database selection statements in a specific data-oriented language called Abap. Kalyanpur et al. [71] state an automatic repair problem in the context of OWL ontologies. Griesmayer et al. [54] repair a specific class of programs called boolean programs: those that only contain boolean variables. Son et al. [153] repairs access-control policies in web applications, using a static analysis and transformations tailored to this domain.

architectural  
repair      Nentwich et al. [114] detect inconsistencies and propose repair actions on XML documents. Their approach is applicable to all structured documents with explicit static inconsistency rules. Along the same line, da Silva [148] use Prolog to propose a repair plan that fixes inconsistencies in UML models. Tran et al. [165] uses repair in the sense of forcing a match between source code dependencies and a dependency model that specifies the acceptable dependencies; this can be called “architectural repair”

test repair      The approach of Daniel et al. [29] does not repair programs but the test cases that are broken in the presence of refactoring. Memon [106] and Gao et al. [46] repair GUI test scripts. For instance, the approaches change the identifiers that are used for driving the GUI manipulation.

### 3.6 Fault Classes and Repair

Some fault classes are well-enough understood so that one can write a code transformation that suppresses all instances of the fault class at once. For instance, one can transform 64-bit integers to unlimited precision arithmetic objects (such as `BigInteger` in Java) to avoid all arithmetic overflows. In the related work, most repair transformations for fault classes are semantic-preserving, but not necessarily.

failure-oblivious  
computing      For instance, a seminal work on semantic modifying transformations is failure-oblivious computing [130]. Considering erroneous reads out of the bounds of an array, failure-oblivious computing transforms the code so that the read returns either the first non-null element, or the element modulo the length of the allocated array. Along the same line, Rinard et al. [134] proposes that out-of-bounds writes are stored in a hashtable and that subsequent reads to the out-of-bound index return the object previously stored in the hashtable. This line of research is based on the philosophical foundation that acceptable results is more important than correct results, this is called “acceptability-oriented computing” [132].

missing checks      Let us now review the approaches that target hardware fault classes. Kadav et al. [68] reason about the missing software checks where the software should not trust the hardware. The tool, called Carburizer, can also insert missing failure-reporting code which can be seen as a kind of automatic repair: a clean error message is better than an incorrect output or a dirty crash.

Carburizer      The same idea of adding condition checks after method and function calls has been explored by Khmelevsky et al. [75].

Rebaudengo et al. [129] describes a set of hardening transformation to

provide detection and repair against bit-flips. The core idea is to duplicate variables and calls and to compare the outputs of duplicated code.

SQL injection

Thomas and Williams [163] propose an approach to automatically transform PHP code to secure SQL statements. The transformations modify the abstract syntax trees in order to inject secured “prepared statements”.

Findbugs

At Google, they develop and use a tool called “error-prone” [2], it does automatic repair of Findbugs like errors [58]. Lawall et al. [82] also defined an approach for declaratively specifying bug patterns and the corresponding patches in a tool called Coccinelle. The same idea has been developed by Kalval and Warburton [70] where the repair strategy is written using a formal transformation language called Trans.

Shaw et al. [139] describe two transformations to fix C buffer overflows: replacement of unsafe calls by alternative safe libraries and replacement of unsafe types by safer ones. They show that the transformations scale to large programs, do not break the existing tests and do not slow down the programs. Coker and Hafiz employ a similar approach for another fault class: integer arithmetic bugs [25]. They propose three program transformations dedicated to integers, and show that the approach scales to real programs.

input filter  
generation

Long et al. [99] uses a static analysis specific to integer arithmetics that detects integer overflow. For all detected potential overflows, the system infers a filter that simply discards the input. To this extent, the repair action is denying the input, a technique also done at runtime and discussed in Section 4.5.

unhandled  
exceptions

Cornu et al. [26] targets unhandled exceptions in Java. They analyze test suite executions to identify the good catch blocks that have resilience capabilities. Then, they transform the caught exception type into a more generic one (i.e. a superclass exception) so as to catch exceptions that would not be caught otherwise. The code transformation, called “catch stretching” is a kind of proactive repair against unexpected exceptions.

## 4 State Repair

State repair consists in changing the state of the program under repair. The state is meant in its largest acceptance: it can be changing the input, the heap, the stack, the environment. For instance, automatic breaking a cycle in a linked list is one kind of state repair. As opposed to behavioral repair, state repair is necessarily done at runtime.

State repair can be rooted in classical fault tolerance [9]. In this large research field, much research has targeted “recovery”, which Avizienis et al. defines as transforming “*a system state that contains one or more errors and (possibly) faults into a state without detected errors*” [9]. In this paper, I use “state repair” instead of “recovery”. This terminological move allows to have

an umbrella term, “repair” above intrinsically related concepts (recovery, resilience, etc), and above behavioral and state repair. I refer to Figure 5.1 of [9] for a bird’s eye presentation of classical recovery, error handling and fault handling.

contract State repair requires an oracle of the bug, an oracle of incorrectness. As opposed to behavioral repair, those oracles have to be available in production, at runtime. This rules out certain oracles discussed in Section 3, such as test suites, and oracles based on static analysis. For state repair, there are three main families of bug oracles. First, state repair often considers violations of non-functional contracts. For instance, crashing with a Segfault or a null pointer exception violates the non-functional contract “the program shall never crash” Second, state repair can also consider functional contracts that are verifiable in production such as pre- and post-conditions. This will be much discussed in Section 4.7.1. Third, there are state repair approaches that reason on “inferred contracts”, obtained by observing the regularities of program states at runtime. In this case, a bug is defined as a program state or behavior that violates those inferred contracts and repair is a follow-up of anomaly detection on program states and executions.

In the following, we order the approaches by repair operators. This more fits to the history of the field than the ordering by kind of oracles, as what was done for behavioral repair.

## 4.1 Reinitialization & Restart

Software Rejuvenation Microreboot Restarting (aka rebooting) a software application is the simplest repair action . It has been much explored under the term “software rejuvenation” [59], but rather with a theoretical stance rather than a practical one.

Candea and colleagues [17, 18, 19] explored in depth the concept of microreboot. Microreboot consists of having a hierarchical structure of fine-grain rebootable components, and, in the presence of failures, to try to restart the application from the smallest component (an EJB) to the biggest one (the physical machine) (in a way that is similar to progressive retry in distributed computing [169]). Their experiments show that this can significantly improve the availability of systems.

## 4.2 Checkpoint & Rollback

Checkpoint A checkpoint and rollback mechanism takes regular snapshots of the execution state and is capable of restoring them later on. The challenges of checkpoint and rollback are first the size and boundaries of the captured state and second the point in time of checkpointing [80, 72]. When a system is equipped with a checkpoint and rollback mechanism, the rollback is the repair. Despite being an old technique, it is valuable in a number of contexts.

Dira [150] is a system that instruments code to detect and recover from control-hijacking attacks through malicious payloads. The repair consists of

Operator	Example Ref.
restart	[17, 169]
try an alternative implementation	[8, 128, 21]
modify the input	[3, 96, 89]
simulate a known error (aka error virtualization)	[143, 93]
change the execution environment	[127, 116, 48]

Table 3: Examples of state change operators for runtime repair

finding the least common ancestor of the function in which the attack is detected and the one in which the payload was read in. Then, the execution is resumed to this frame and all state changes are undone. Similarly, Assure [143] is a technique also based on checkpointing to provide self-healing capabilities. Recent papers also do checkpoint and rollback as part of the repair, such as [21].

### 4.3 Alternatives

recovery block

Another classical concept of fault-tolerance is n-version programming. Either with voting [8] or retrying with recovery blocks [128], it consists of relying on alternative implementations to recover from errors. This concept is now explored using natural sets of alternatives (as opposed to being engineered) or with automatically created sets of variants [15]. For instance, Carzaniga et al. [22] repair web applications at runtime with a repair strategy is based on a set of API-specific alternative rules: for instance calling `bar()` instead of `foo()`. They later applied the same idea for recovering from runtime exceptions in Java [21]. Hosek and Cadar [57] use a different kind of natural diversity: upon failures, they switch from past or newer versions of the same application. The key idea is that bugginess is not monotonic: some bugs disappear while others appear over time.

automatic  
workaround

### 4.4 Reconfiguration

Reconfiguring an application is one kind of recovery [9], thus one kind of state repair. Indeed, it has much been explored when “self-healing” was a hype term. For instance, Cheng et al. [24] use the three core runtime reconfiguration operators (add component, move component, delete component) to optimize quality-of-service values. The same line of repair can be found in [47, 141], which are relatively cited papers. In the context of web service orchestration, the repair actions of Friedrich et al. [121, 43] consist of substituting it a web service by another (which is a reconfiguration) and retrying a service call.

## 4.5 Input Modification

If the system fails on some input, one state repair action consists of modifying the input. Denying the input is also a possible option, which can be considered as an extreme case of input modification.

data diversity

Ammann and Knight’s “data diversity” [3] aims at enabling the computation of a program in the presence of failures. The idea of data diversity is that, when a failure occurs, the input data is changed so that the new input resulting from the change does not result in a failure. The assumption is that the output based on this artificial input, through an inverse transformation, remains acceptable in the domain under consideration.

input  
rectification

Long et al [96] present the idea of automated input rectification: instead of refusing anomalous inputs, they change it so that it fits into the space of typical and acceptable inputs, this is called “input rectification”.

Liand and Sekar [89] repair buffer overflows by learning common profiles between the characteristics of crashing inputs. Once a valid profile is identified, crashing inputs are denied. While the paper is about security, it can be seen as a runtime technique to repair memory errors of the form of buffer overflows. The fact that the buffer overflow is accidental (due to a bug) or maliciously triggered is irrelevant from a repair perspective. Along the same line of input denying, Vigilante [27] is an integrated approach for mitigating malicious attacks. The counter-measure to worm attacks is filtering: once invalid or malicious inputs are detected they are filtered out and the current request or task is aborted.

## 4.6 Environment Perturbation

If the system fails under certain conditions, one can get the next requests to succeed by changing the runtime environment (e.g. the memory, the scheduling) or the configuration.

Rx  
DieHard  
Exterminator

Qin et al. [127] shows that memory errors can be avoided by padding allocated memory blocks with extra space. Berger and Zorn [13] do the same thing and add replication. However, the difference with Rw is that their system allows for probabilistic reasoning on the resulting memory safety. Novark et al. [118] explores the same idea. Differently, Nguyen and Rinard [116] enforces a bounded memory size by cyclic memory allocation in a way that is similar to failure-oblivious computing (already presented in Section 3.6). Garvin et al. [48] address configuration bugs and propose “reconfiguration workarounds” that change the configuration causing a failure.

Jula et al. [67] presents a system to defend against deadlocks at runtime. The system first detects synchronization patterns of deadlocks, and when the pattern is detected, the system avoids re-occurrences of the deadlock with additional locks.

Tallam et al. [160] names this family of technique “execution perturbations”. For concurrency and memory bugs, they show that removing thread

interruptions, padding memory allocations, and performing denial of requests is a way to avoid failures.

## 4.7 Rollforward

Rollforward (or forward recovery) means transforming the current system state into a correct one. There are several techniques of forward recovery: invariant restoration, error virtualization, etc.

### 4.7.1 Invariant Restoration

In some cases, state correctness can be expressed as an invariant. Consequently, repair means restoring the invariant, if possible with a minimum of changes from the current erroneous state.

Demsky and Rinard [32] uses a specification language to express correctness properties on data structures. This specification is then used at runtime to automatically repair broken data structure (concrete instances at runtime, not the abstract data type). Elkarablieh et al. [40] also automatically repair data structures at runtime, the difference with Demsky and Rinard is that they rely on an invariant written in regular Java code (a “repOK” boolean method).

ClearView

Perkins et al. [120] presented ClearView a system for automatically repairing errors in production. The system works on low level x86 binaries and consists of monitoring the system execution to learn invariants. Those invariants are then monitored, and a violation is followed by a forced restoration. The repairs are at the level of CPU registers and memory location changes.

Lewis and Whitehead [88] have a generic repair approach for event-based system by defining a runtime fault-monitor, but the core idea is that same: when an invariant is violated, the repair system automatically restores it. The example in a video-game domain is fun: if Mario is hanged in the sky due to specific sequence of actions and interactions, it is forcefully put back on the ground. Beyond data structures and video-games, in real systems, many strange and undesired system states can happen from complex chains of events and interactions, but it is often possible to state simple invariants to guide runtime repair.

### 4.7.2 Error virtualization

Error virtualization consists of handling an unknown and unrecoverable error with error-handling code that is already present in the system yet designed for handling other errors.

This idea has been much explored at Columbia University. For instance, Sidiroglou et al. [144] do error virtualization in system that imitates biological immunity. They combine error virtualization with selective transactional emulation, a technique consisting of emulating the execution of native code with an an interpreter in a transactional manner. When a failure occurs in an emulated section, all state changes are undone (a kind of micro rollback at the

Assure level of functions). In Assure [143], the idea of error virtualization is associated with fuzzing to discover and test in advance valuable error virtualization points, called rescue points.

### 4.7.3 Other Forward Recovery

loop perforation

Carbin et al. [20] introduced a system that monitors programs in order to detect infinite loops and escaping them. The system works with binary code instrumentation and breaks the loops with no memory state changes detected during their execution. Along the same line is the concept of “loop perforation” [146]. Sidiroglou et al. have shown [146] that it is possible to skip the execution of loop iterations in certain application domains. For instance, in a video decoding algorithm (codec), skipping some loop iterations only has an effect on some pixels or contours but does not completely degrade or crash the software application. On the other hand, skipping loop iterations is key with respect to performance. In other words, there is a trade-off between the performance and accuracy. This trade-off can be set offline (e.g. by arbitrarily skipping one every two loops) or dynamically based on the current load of the machine.

Dobilyi and Weimer [36] target repair of null pointer exceptions. Using code transformation, they introduce hooks to a recovery framework. This framework is responsible for forward recovery of the form of creating a default object of an appropriate type to replace the null value or of skipping instructions.

RCV

Long et al. [100] introduces the idea of “recovery shepherding”. Upon certain errors (null dereferences and divide by zero), recovery shepherding consists in returning a manufactured value, as for failure oblivious computing. However, the key idea of recovery shepherding is to track the manufactured value so as to see 1) whether they are passed to system calls or files and 2) whether they disappear. In the former case, system calls and file writes are disabled if they involve a fake manufactured value, in order to limit error propagation. When a manufactured value is no longer used and referenced, it means that the error has somehow evaporated, and the experiments of the paper show that this is often the case.

## 4.8 Collaborative Repair

A cross-cutting concern of repair at runtime is to share the repairs that work across all instances of the same application. This has been explored under the name of “application community”. Locasto et al. [93] uses application communities to find and distribute repairs of the form of stack manipulation. Rinard et al. [131] also reports on experiments on the centralization of monitoring information and the distribution of repairs across a community of applications.

## 5 Empirical Knowledge on Repair

Beyond proposing new repair techniques, there is a thread of research on empirically investigating the foundations, impact and applicability of automatic repair, whether behavioral or state repair.

There is wealth of information in software repositories that can be used for repair. In particular, one can mine bug reports and commits for knowledge that is valuable for automatic repair. Martinez & Monperrus [103] studied 89,993 of commits to mine repair actions from manually-written patches. By repair actions, they mean kinds of changes on the abstract syntax trees of programs such as modifying an if condition. They later investigated [104] the *redundancy assumption* in automatic repair (whether you can fix bugs by rearranging existing code), and found that it holds in practice: many bug fix commits only rearrange existing code, a result confirmed by Barr et al. [11]. Zhong & Su [177] conducted a case study on over 9,000 real-world patches and found important facts for automatic repair: for instance, their analysis outlines that some bugs are repaired with changing the configuration files.

On the goodness of synthesized patches, Fry et al. [44] conducted a study of machine-generated patches based on 150 participants and 32 real-world defects. Their work shows that machine-generated patches are slightly less maintainable than human-written ones. Tao et al. [162] performed a similar study to study whether machine-generated patches assist human debugging. Monperrus [111] further discussed the *patch acceptability* criteria of synthesized patches and emphasized that assessing patch acceptability may require a high level of expertise, a result confirmed by [38]. Qi et al. [126] are the first to thoroughly analyze the patches generated by Genprog, and found that most of them are incorrect. It is an open question whether this holds for test-suite based repair in general or not [38]. When they are incorrect, it is because they exploit specificities and weaknesses of the test suite, which can be seen as a kind of *overfitting*. A repair technique is said to overfit when the synthesized patch only works on the failing inputs and fails to generalize. Smith et al. [151] also studied the problem of overfitting in automatic repair; on a dataset of student programs, they show that Genprog and related techniques do suffer from overfitting.

Finally, for the knowledge on repair to consolidate, there is a need for accepted, well-defined and publicly available benchmarks [111]. Le Goues et al. [86] have set up such a benchmarks for bugs in C programs, it totals 1183 bugs, collected in open-source projects and student code.

## 6 Related Techniques

We now present works that are related to automatic repair, yet not being “automatic repair” per se, according to the definitions we gave in Section 3 and 4. In particular, they either miss the full automation or the actual repair of real programs.

## 6.1 Forward Engineering For Repair

Many authors have tried to list the important principles to have robust, resilient if not self-repairable applications. These principles can be implemented and enforced as first-class concepts in frameworks and libraries. This is what I call “forward engineering for repair”.

Somayaji et al. describe principles to build immune computer systems [152]: distributability, multi-layering, diversity, disposability, autonomy, adaptability, behavioral sense-of-self, anomaly detection. Candea and Fox [16] define a set of characteristics for programs to recover quickly: with those characteristics an application becomes “*crash-only software*”. The two key characteristics are that all interactions between components have a timeout and all resources are leased. Sussmann [159] as well as Gabriel and Goldmann [45] also provide insightful perspectives on how to build resilient and self-repairable software.

There are also frameworks for supporting repair. Flora [154] is a framework to support local restart in applications. It is principally composed of a communication manager for dropping or queuing messages between components. Denaro et al. [34] proposes an architecture to fix interoperability bugs in service oriented systems. Adaptors between service variants are manually written and are selected at runtime to enable correct communication. Levinson [87] defines an embedded DSL to support runtime searches in a space of program variations. Zhou et al. [178] defines annotations for operating system C code in order to recover from driver errors in Linux. The annotations are checked by a type system and drives invariant restoration. Demsky and Dash [33] proposes Bristlecone, a language with built-in robustness capabilities. Bristlecone is based on tasks and dependences between tasks, as well as transactional state changes. Error-handling is thus fully automated.

A known characteristic of bugs is that the same kind of bug can affect many different locations in the same code base. In this case, it is desirable to write a unique patch that is then applied to all those locations. The generic patch can be inferred from a concrete instance at a given location or written in an abstract way. This has been called “systematic editing” by Meng et al. [108]. Similarly, Sun et al. [157, 158] propose tool support for patch applications. The Coccinelle tool [119] also provides this functionality. The abstract patches can be automatically inferred from concrete instances [4, 107].

## 6.2 Repair Suggestions

There are some systems which give “repair suggestions” to the developer. While it is not fully automated, if the suggestion is correct, such as system can be seen as providing partial automatic repair, where the repair system and the developer work in tandem.

Hartmann et al. [55] designed a system called “HelpMeout” that proposes suggestions to fix error messages. The system targets compiler error messages and runtime exceptions. It first collects error messages and the associated changes that occur on developer’s machines that are monitored. Then, when

the same error message is encountered by another developer, the system compares the erroneous source file with the closest fixed version that is in the database. It uses a tailored distance metric to increase the relevance of suggestions.

BugFix Jeffrey et al. [61] presented a fix suggestion approach based on association rules. The rules suggest a bug fix action for suspicious statements represented by a number of features (in the machine learning meaning). The features (called “descriptors” in the paper) are abstraction over the tokens of the statements. The prediction also uses “interesting value mapping pairs” (IVMP) which are concrete values that enable test cases to pass (aka value replacement [62] and angelic values [23, 31]). The bug fix recommendations are typical comparison operator change, constant change, add or increase numerical values.

MintHint Kaleeswaran et al. [69] have proposed a repair suggestion approach based on correlations variable values and expected output. The expected output is obtained through concolic executions, and the repair hints consist of changing the RHS of a single assignment statement.

Excel Abraham and Erwig [1] suggest change in Excel formulas. Malik et al. [101] transform runtime data structure repair (see 4.7.1) as fix suggestions. Brodie et al. [14] design a distance metric across call stacks (stack trace) to match issue reports and known fixes.

### 6.3 Theoretical Software Repair

Some authors explore automatic repair with strong assumptions under which there exists no program in practice. We mention here the most notable works in theoretical software repair.

For instance, Jobstmann et al. [65] repair programs that are expressed in linear temporal logics. George [49] describes a simple and theoretical programming model that supports automatic recovery via a kind of homeostasis that maintains invariants. Fisher et al. [124] also perform repair on a toy formal language. Wang and Cheng [167] state program repair as edit sequences on state machines. Zhang and Ding [176] repair computation tree logic models.

## 7 Conclusion

In this paper, I have presented an annotated bibliography on automatic software repair. This research field is both old and new. It is old because we can find techniques related to automatic repair in fault-tolerance papers from the 70es and 80es, for instance a 1973 paper is entitled “STAREX self-repair routines: software recovery in the JPL-STAR computer” [135]. It is new, because the idea of automatically changing the code, i.e. behavioral repair, has started to be explored only since the end of 2000. Whether old or new, the techniques have to scale to today’s size and complexity or software stacks, and we are not there yet. This means that this is only the beginning, and in

the upcoming years, we are going to have much fun, surprise and admiration in the field of automatic software repair.

## References

- [1] R. Abraham and M. Erwig. “Goal-Directed Debugging of Spreadsheets”. In: *Proceedings of the 2005 IEEE Symposium on Visual Languages and Human-Centric Computing*. 2005, pp. 37–44.
- [2] E. Aftandilian, R. Sauciuc, S. Priya, and S. Krishnan. “Building Useful Program Analysis Tools Using An Extensible Java Compiler”. In: *International Working Conference on Source Code Analysis and Manipulation (SCAM)*. 2012, pp. 14–23.
- [3] P. E. Ammann and J. C. Knight. “Data Diversity: An Approach to Software Fault Tolerance”. In: *IEEE Transactions on Computers* 37.4 (1988), pp. 418–425.
- [4] J. Andersen and J. L. Lawall. “Generic Patch Inference”. In: *Automated Software Engineering* 17.2 (2010), pp. 119–148.
- [5] A. Arcuri. “Automatic Software Generation and Improvement Through Search Based Techniques”. PhD thesis. The University of Birmingham, 2009.
- [6] A. Arcuri. “Evolutionary Repair of Faulty Software”. In: *Applied Soft Computing* 11.4 (2011), pp. 3494–3514.
- [7] A. Arcuri and X. Yao. “A Novel Co-evolutionary Approach to Automatic Software Bug Fixing”. In: *Proceedings of the IEEE Congress on Evolutionary Computation*. 2008, pp. 162–168.
- [8] A. Avizienis. “The N-version Approach to Fault-tolerant Software”. In: *IEEE Transactions on Software Engineering* 11.12 (1985), pp. 1491–1501.
- [9] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. “Basic Concepts and Taxonomy of Dependable and Secure Computing”. In: *IEEE Transactions on Dependable and Secure Computing* 1.1 (2004), pp. 11–33.
- [10] T. Azim, I. Neamtiu, and L. Marvel. “Towards Self-healing Smartphone Software via Automated Patching”. In: *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering*. 2014, pp. 623–628.
- [11] E. T. Barr, Y. Brun, P. T. Devanbu, M. Harman, and F. Sarro. “The Plastic Surgery Hypothesis”. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 2014, pp. 306–317.

- [12] E. Barr, M. Harman, P. McMinn, M. Shahbaz, and S. Yoo. “The Oracle Problem in Software Testing: a Survey”. In: *IEEE Transactions on Software Engineering* 41.5 (2015), pp. 507–525.
- [13] E. D. Berger and B. G. Zorn. “DieHard: Probabilistic Memory Safety for Unsafe Languages”. In: *ACM SIGPLAN Notices* 41.6 (2006), pp. 158–168.
- [14] M. Brodie, S. Ma, G. Lohman, L. Mignet, M. Wilding, J. Champlin, and P. Sohn. “Quickly Finding Known Software Problems via Automated Symptom Matching”. In: *Proceedings of the International Conference on Autonomic Computing*. 2005, pp. 101–110.
- [15] Y. Brun, E. Barr, M. Xiao, C. L. Goues, and P. Devanbu. *Evolution Vs. Intelligent Design in Program Patching*. Tech. rep. UC Davis, 2013.
- [16] G. Candea and A. Fox. “Crash-only Software”. In: *Proceedings of the 9th Conference on Hot Topics in Operating Systems*. 2003, pp. 12–12.
- [17] G. Candea and A. Fox. “Recursive Restartability: Turning the Reboot Sledgehammer Into a Scalpel”. In: *Proceedings of the 8th Workshop on Hot Topics in Operating Systems*. 2001, pp. 125–130.
- [18] G. Candea, E. Kiciman, S. Zhang, P. Keyani, and A. Fox. “JAGR: An Autonomous Self-recovering Application Server”. In: *Proceedings of the Workshop on Active Middleware Services*. 2003, pp. 168–177.
- [19] G. Candea, S. Kawamoto, Y. Fujiki, G. Friedman, and A. Fox. “Microreboot: a Technique for Cheap Recovery”. In: *Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation*. 2004, pp. 3–3.
- [20] M. Carbin, S. Misailovic, M. Kling, and M. C. Rinard. “Detecting and Escaping Infinite Loops with Jolt”. In: *Proceedings of ECOOP*. 2011, pp. 609–633.
- [21] A. Carzaniga, A. Gorla, A. Mattavelli, N. Perino, and M. Pezzè. “Automatic Recovery From Runtime Failures”. In: *Proceedings of the International Conference on Software Engineering*. 2013.
- [22] A. Carzaniga, A. Gorla, N. Perino, and M. Pezzè. “Automatic Workarounds for Web Applications”. In: *Proceedings of the Foundations of Software Engineering Conference*. 2010, pp. 237–246.
- [23] S. Chandra, E. Torlak, S. Barman, and R. Bodik. “Angelic Debugging”. In: *Proceeding of the International Conference on Software Engineering*. 2011, pp. 121–130.
- [24] S.-W. Cheng, D. Garlan, B. R. Schmerl, J. P. Sousa, B. Spitnagel, and P. Steenkiste. “Using Architectural Style As a Basis for System Self-repair”. In: *Proceedings of the IFIP 17th World Computer Congress / 3rd IEEE/IFIP Conference on Software Architecture: System Design, Development and Maintenance*. 2002, pp. 45–59.

- [25] Z. Coker and M. Hafiz. “Program Transformations to Fix C Integers”. In: *Proceedings of the International Conference on Software Engineering*. 2013, pp. 792–801.
- [26] B. Cornu, L. Seinturier, and M. Monperrus. “Exception Handling Analysis and Transformation Using Fault Injection: Study of Resilience Against Unanticipated Exceptions”. In: *Information and Software Technology* (2014).
- [27] M. Costa, J. Crowcroft, M. Castro, A. Rowstron, L. Zhou, L. Zhang, and P. Barham. “Vigilante: End-to-end Containment of Internet Worms”. In: *ACM SIGOPS Operating Systems Review*. Vol. 39. 5. 2005, pp. 133–147.
- [28] V. Dallmeier, A. Zeller, and B. Meyer. “Generating Fixes From Object Behavior Anomalies”. In: *Proceedings of the International Conference on Automated Software Engineering*. 2009.
- [29] B. Daniel, V. Jagannath, D. Dig, and D. Marinov. “ReAssert: Suggesting Repairs for Broken Unit Tests”. In: *Proceedings of the 24th IEEE/ACM International Conference on Automated Software Engineering*. 2009, pp. 433–444.
- [30] V. Debroy and W. Wong. “Using Mutation to Automatically Suggest Fixes for Faulty Programs”. In: *Proceedings of the International Conference on Software Testing, Verification and Validation*. 2010, pp. 65–74.
- [31] F. DeMarco, J. Xuan, D. Le Berre, and M. Monperrus. “Automatic Repair of Buggy If Conditions and Missing Preconditions with SMT”. In: *Proceedings of the 6th International Workshop on Constraints in Software Testing, Verification, and Analysis (CSTVA 2014)*. 2014.
- [32] B. Demsky and M. Rinard. “Automatic Detection and Repair of Errors in Data Structures”. In: *ACM SIGPLAN Notices* 38.11 (2003), pp. 78–95.
- [33] B. Demsky and A. Dash. “Bristlecone: a Language for Robust Software Systems”. In: *Proceedings of ECOOP*. 2008, pp. 490–515.
- [34] G. Denaro, M. Pezzè, and D. Tosi. “Ensuring Interoperable Service-oriented Systems Through Engineered Self-healing”. In: *Proceedings of ESEC/FSE*. 2009, pp. 253–262.
- [35] L. A. Dennis, R. Monroy, and P. Nogueira. “Proof-directed Debugging and Repair”. In: *Seventh Symposium on Trends in Functional Programming*. 2006, pp. 131–140.
- [36] K. Dobolyi and W. Weimer. “Changing Java’s Semantics for Handling Null Pointer Exceptions”. In: *19th International Symposium on Software Reliability Engineering*. 2008, pp. 47–56.
- [37] J. Duraes and H. Madeira. “Emulation of Software Faults: a Field Data Study and a Practical Approach”. In: *IEEE Transactions on Software Engineering* 32.11 (2006), pp. 849–867.

- [38] T. Durieux, M. Martinez, M. Monperrus, R. Sommerard, and J. Xuan. *Automatic Repair of Real Bugs: An Experience Report on the Defects4J Dataset*. Tech. rep. 1505.07002. Arxiv, 2015.
- [39] M. Eisenstadt. “My Hairiest Bug War Stories”. In: *Communications of the ACM* 40.4 (1997), pp. 30–37.
- [40] B. Elkarablieh, I. Garcia, Y. Suen, and S. Khurshid. “Assertion-based Repair of Complex Data Structures”. In: *Proceedings of the Twenty-second IEEE/ACM International Conference on Automated Software Engineering*. 2007, pp. 64–73.
- [41] H. ETO and K. Yoda. “Propolice: Improved Stacksmashing Attack Detection”. In: *IPSJ SIG Notes* 75 (2001), pp. 181–188.
- [42] S. Forrest, T. Nguyen, W. Weimer, and C. Le Goues. “A Genetic Programming Approach to Automated Software Repair”. In: *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation*. 2009, pp. 947–954.
- [43] G. Friedrich, M. Fugini, E. Mussi, B. Pernici, and G. Tagni. “Exception Handling for Repair in Service-Based Processes”. In: *IEEE Transactions on Software Engineering* 36.2 (2010), pp. 198–215.
- [44] Z. P. Fry, B. Landau, and W. Weimer. “A Human Study of Patch Maintainability”. In: *Proceedings of the International Symposium on Software Testing and Analysis*. 2012, pp. 177–187.
- [45] R. P. Gabriel and R. Goldman. “Conscientious Software”. In: *Acm Sigplan Notices*. Vol. 41. 10. 2006, pp. 433–450.
- [46] Z. Gao, Z. Chen, Y. Zou, and A. Memon. “SITAR: GUI Test Script Repair”. In: *IEEE Transactions on Software Engineering* (2015).
- [47] D. Garlan, S.-W. Cheng, and B. Schmerl. “Increasing System Dependability Through Architecture-based Self-repair”. In: *Architecting Dependable Systems*. 2003, pp. 61–89.
- [48] B. J. Garvin, M. B. Cohen, and M. B. Dwyer. “Using Feature Locality: Can We Leverage History to Avoid Failures During Reconfiguration?” In: *Proceedings of the 8th Workshop on Assurances for Self-adaptive Systems*. 2011, pp. 24–33.
- [49] S. George, D. Evans, and S. Marchette. “A Biological Programming Model for Self-Healing”. In: *Proceedings of SSRS*. 2013.
- [50] D. Ghosh, R. Sharman, H. Raghav Rao, and S. Upadhyaya. “Self-healing Systems—survey and Synthesis”. In: *Decision Support Systems* 42.4 (2007), pp. 2164–2185.
- [51] D. Gopinath, S. Khurshid, D. Saha, and S. Chandra. “Data-guided Repair of Selection Statements”. In: *Proceedings of the 36th International Conference on Software Engineering*. 2014, pp. 243–253.

- [52] D. Gopinath, M. Z. Malik, and S. Khurshid. “Specification-based Program Repair Using SAT”. In: *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 2011.
- [53] C. L. Goues, T. Nguyen, S. Forrest, and W. Weimer. “GenProg: a Generic Method for Automatic Software Repair”. In: *IEEE Transactions on Software Engineering* 38 (2012), pp. 54–72.
- [54] A. Griesmayer, R. Bloem, and B. Cook. “Repair of Boolean Programs with An Application to C”. In: *Computer Aided Verification*. 2006, pp. 358–371.
- [55] B. Hartmann, D. MacDougall, J. Brandt, and S. R. Klemmer. “What Would Other Programmers Do: Suggesting Solutions to Error Messages”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2010, pp. 1019–1028.
- [56] H. He and N. Gupta. “Automated Debugging Using Path-Based Weakest Preconditions”. In: *FASE*. 2004, pp. 267–280.
- [57] P. Hosek and C. Cadar. “Safe Software Updates via Multi-version Execution”. In: *Proceedings of the International Conference on Software Engineering*. 2013, pp. 612–621.
- [58] D. Hovemeyer and W. Pugh. “Finding Bugs Is Easy”. In: *ACM SIGPLAN Notices* 39.12 (2004).
- [59] Y. Huang, C. Kintala, N. Kolettis, and N. Fulton. “Software Rejuvenation: Analysis, Module and Applications”. In: *Proceedings of the International Symposium on Fault-Tolerant Computing*. 1995, pp. 381–390.
- [60] D. Jackson and M. Vaziri. “Finding Bugs with a Constraint Solver”. In: *Proceedings of the 2000 ACM SIGSOFT International Symposium on Software Testing and Analysis*. 2000, pp. 14–25.
- [61] D. Jeffrey, M. Feng, N. Gupta, and R. Gupta. “BugFix: a Learning-based Tool to Assist Developers in Fixing Bugs”. In: *ICPC*. 2009, pp. 70–79.
- [62] D. Jeffrey, N. Gupta, and R. Gupta. “Fault Localization Using Value Replacement”. In: *Proceedings of the International Symposium on Software Testing and Analysis*. 2008, pp. 167–178.
- [63] S. Jha, S. Gulwani, S. A. Seshia, and A. Tiwari. “Oracle-guided Component-based Program Synthesis”. In: *Proceedings of the International Conference on Software Engineering*. Vol. 1. 2010, pp. 215–224.
- [64] G. Jin, L. Song, W. Zhang, S. Lu, and B. Liblit. “Automated Atomicity-violation Fixing”. In: *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*. 2011, pp. 389–400.

- [65] B. Jobstmann, A. Griesmayer, and R. Bloem. “Program Repair As a Game”. In: *Computer Aided Verification*. 2005, pp. 226–238.
- [66] J. Jones and M. Harrold. “Empirical Evaluation of the Tarantula Automatic Fault-localization Technique”. In: *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering*. 2005, pp. 273–282.
- [67] H. Jula, D. Tralamazza, C. Zamfir, and G. Candea. “Deadlock Immunity: Enabling Systems to Defend Against Deadlocks”. In: *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation*. 2008, pp. 295–308.
- [68] A. Kadav, M. J. Renzelmann, and M. M. Swift. “Tolerating Hardware Device Failures in Software”. In: *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*. 2009, pp. 59–72.
- [69] S. Kaleeswaran, V. Tulsian, A. Kanade, and A. Orso. “Minthint: Automated Synthesis of Repair Hints”. In: *Proceedings of the International Conference on Software Engineering*. 2014, pp. 266–276.
- [70] S. Kalvala and R. Warburton. “A Formal Approach to Fixing Bugs”. In: *Formal Methods, Foundations and Applications*. 2011, pp. 172–187.
- [71] A. Kalyanpur, B. Parsia, E. Sirin, and B. Cuenca-Grau. “Repairing Unsatisfiable Concepts in OWL Ontologies”. In: *The Semantic Web: Research and Applications*. Vol. 4011. 2006, pp. 170–184.
- [72] M. Kasbekar, C. Narayanan, and C. Das. “Selective Checkpointing and Rollbacks in Multi-threaded Object-oriented Environment”. In: *IEEE Transactions on Reliability* 48.4 (1999), pp. 325–337.
- [73] Y. Ke, K. T. Stolee, C. Le Goues, and Y. Brun. “Repairing Programs with Semantic Code Search”. In: *Proceedings of the International Conference on Automated Software Engineering*. 2015.
- [74] C. Kern and J. Esparza. “Automatic Error Correction of Java Programs”. In: *Formal Methods for Industrial Critical Systems*. 2010, pp. 67–81.
- [75] Y. Khmelevsky, M. C. Rinard, and S. Sidiroglou-Douskos. “A Source-to-source Transformation Tool for Error Fixing.” In: *Proceedings of CASCON*. 2013, pp. 147–160.
- [76] D. Kim, J. Nam, J. Song, and S. Kim. “Automatic Patch Generation Learned From Human-Written Patches”. In: *Proceedings of ICSE’2013*. 2013.
- [77] D. E. Knuth. “The Errors of TEX”. In: *Softw., Pract. Exper.* 19.7 (1989), pp. 607–685.
- [78] R. Könighofer and R. Bloem. “Automated Error Localization and Correction for Imperative Programs”. In: *Formal Methods in Computer-Aided Design (FMCAD), 2011*. 2011, pp. 91–100.

- [79] R. Könighofer and R. Bloem. “Repair with On-the-fly Program Analysis”. In: *Hardware and Software: Verification and Testing*. 2013, pp. 56–71.
- [80] R. Koo and S. Toueg. “Checkpointing and Rollback-recovery for Distributed Systems”. In: *IEEE Transactions on Software Engineering* 1 (1987), pp. 23–31.
- [81] M. Lake. *Epic Failures: 11 Infamous Software Bugs*. <http://www.computerworld.com/article/2515483/enterprise-applications/epic-failures--11-infamous-software-bugs.html>. 2010.
- [82] J. L. Lawall, J. Brunel, N. Palix, R. R. Hansen, H. Stuart, and G. Muller. “WYSIWIB: a Declarative Approach to Finding API Protocols and Bugs in Linux Code”. In: *International Conference on Dependable Systems & Networks*. 2009, pp. 43–52.
- [83] N. Lazaar, A. Gotlieb, and Y. Lebbah. “A Framework for the Automatic Correction of Constraint Programs”. In: *Proceedings of the International Conference on Software Testing, Verification and Validation*. 2011, pp. 319–326.
- [84] C. Le Goues, M. Dewey-Vogt, S. Forrest, and W. Weimer. “A Systematic Study of Automated Program Repair: Fixing 55 Out of 105 Bugs for \$8 Each”. In: *Proceedings of the International Conference on Software Engineering*. 2012, pp. 3–13.
- [85] C. Le Goues, S. Forrest, and W. Weimer. “Current Challenges in Automatic Software Repair”. In: *Software Quality Journal* 21.3 (2013), pp. 421–443.
- [86] C. Le Goues, N. Holtschulte, E. K. Smith, Y. Brun, P. Devanbu, S. Forrest, and W. Weimer. “The ManyBugs and IntroClass Benchmarks for Automated Repair of C Programs”. In: *IEEE Transactions on Software Engineering (TSE)*, in press (2015).
- [87] R. Levinson. “Unified Planning and Execution for Autonomous Software Repair”. In: *Workshop on Plan Execution: a Reality Check*. 2005.
- [88] C. Lewis and J. Whitehead. “Runtime Repair of Software Faults Using Event-driven Monitoring”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - ICSE '10 2* (2010), p. 275.
- [89] Z. Liang and R. Sekar. “Fast and Automated Generation of Attack Signatures: a Basis for Building Self-protecting Servers”. In: *Proceedings of the 12th ACM Conference on Computer and Communications Security*. 2005, pp. 213–222.

- [90] Z. Lin, X. Jiang, D. Xu, B. Mao, and L. Xie. “AutoPaG: Towards Automated Software Patch Generation with Source Code Root Cause Identification and Repair”. In: *Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security*. 2007, pp. 329–340.
- [91] C. Liu, J. Yang, L. Tan, and M. Hafiz. “R2Fix: Automatically Generating Bug Fixes From Bug Reports”. In: *Proceedings of the International Conference on Software Testing, Verification and Validation (ICST)*. 2013, pp. 282–291.
- [92] P. Liu and C. Zhang. “Axis: Automatically Fixing Atomicity Violations Through Solving Control Constraints”. In: *Proceedings of the 2012 International Conference on Software Engineering*. 2012, pp. 299–309.
- [93] M. E. Locasto, S. Sidiroglou, and A. D. Keromytis. “Software Self-Healing Using Collaborative Application Communities”. In: *Proceedings of the Network and Distributed System Security Symposium*. 2006.
- [94] F. Logozzo and T. Ball. “Modular and Verified Automatic Program Repair”. In: *Proceedings of the 27th ACM International Conference on Object Oriented Programming Systems Languages and Applications*. 2012.
- [95] F. Logozzo and M. Martel. “Automatic Repair of Overflowing Expressions with Abstract Interpretation”. In: *Semantics, Abstract Interpretation, and Reasoning About Programs: Essays Dedicated to David A. Schmidt on the Occasion of His Sixtieth Birthday*. Vol. 129. 2013, pp. 341–357.
- [96] F. Long, V. Ganesh, M. Carbin, S. Sidiroglou, and M. Rinard. “Automatic Input Rectification”. In: *Proceedings of ICSE*. 2012.
- [97] F. Long and M. C. Rinard. “Prophet: Automatic Patch Generation via Learning From Successful Patches”. In: *Proceedings of the Symposium on Principles of Programming Languages*. 2016.
- [98] F. Long and M. C. Rinard. “Staged Program Repair with Condition Synthesis”. In: *Proceedings of ESEC/FSE*. 2015.
- [99] F. Long, S. Sidiroglou-Douskos, D. Kim, and M. Rinard. “Sound Input Filter Generation for Integer Overflow Errors”. In: *ACM SIGPLAN Notices* 49.1 (2014), pp. 439–452.
- [100] F. Long, S. Sidiroglou-Douskos, and M. C. Rinard. “Automatic Runtime Error Repair and Containment via Recovery Shepherding”. In: *ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2014.
- [101] M. Malik, J. Siddiqi, and S. Khurshid. “Constraint-Based Program Debugging Using Data Structure Repair”. In: *International Conference on Software Testing, Verification and Validation (ICST)*. 2011, pp. 190–199.

- [102] S. L. Marcote and M. Monperrus. *Automatic Repair of Infinite Loops*. Tech. rep. 1504.05078. Arxiv, 2015.
- [103] M. Martinez and M. Monperrus. “Mining Software Repair Models for Reasoning on the Search Space of Automated Program Fixing”. In: *Empirical Software Engineering* 20.1 (Sept. 2015), pp. 176–205.
- [104] M. Martinez, W. Weimer, and M. Monperrus. “Do the Fix Ingredients Already Exist? An Empirical Inquiry into the Redundancy Assumptions of Program Repair Approaches”. In: *Proceedings of the International Conference on Software Engineering*. 2014.
- [105] S. Mechtaev, J. Yi, and A. Roychoudhury. “DirectFix: Looking for Simple Program Repairs”. In: *Proceedings of the 37th International Conference on Software Engineering*. 2015.
- [106] A. M. Memon. “Automatically Repairing Event Sequence-based GUI Test Suites for Regression Testing”. In: *ACM Transactions on Software Engineering and Methodology* 18.2 (2008), p. 4.
- [107] N. Meng, M. Kim, and K. S. McKinley. “LASE: Locating and Applying Systematic Edits by Learning From Examples”. In: *Proceedings of the International Conference on Software Engineering*. 2013, pp. 502–511.
- [108] N. Meng, M. Kim, and K. S. McKinley. “Systematic Editing: Generating Program Transformations From An Example”. In: *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation*. 2011, pp. 329–342.
- [109] M. G. Merideth. “Enhancing Survivability with Proactive Fault-containment”. In: *The 2003 International Conference on Dependable Systems and Networks*. 2003.
- [110] B. Meyer. “Applying ’design by Contract’”. In: *Computer* 25.10 (1992), pp. 40–51.
- [111] M. Monperrus. “A Critical Review of "Automatic Patch Generation Learned from Human-Written Patches": Essay on the Problem Statement and the Evaluation of Automatic Software Repair”. In: *Proceedings of the International Conference on Software Engineering*. 2014, pp. 234–242.
- [112] M. Monperrus. *Principles of Antifragile Software*. Arxiv 1404.3056. 2014.
- [113] P. Muntean, V. K. Kommanapalli, A. Ibing, and C. Eckert. “Automated Generation of Buffer Overflows Quick Fixes Using Symbolic Execution and SMT”. In: *International Conference on Computer Safety, Reliability & Security (SAFECOMP’15)*. 2015.
- [114] C. Nentwich, W. Emmerich, and A. Finkelstein. “Consistency Management with Repair Actions”. In: *Proceedings of the 25th International Conference on Software Engineering*. 2003, pp. 455–464.

- [115] H. D. T. Nguyen, D. Qi, A. Roychoudhury, and S. Chandra. “SemFix: Program Repair via Semantic Analysis”. In: *Proceedings of the International Conference on Software Engineering*. 2013.
- [116] H. H. Nguyen and M. Rinard. “Detecting and Eliminating Memory Leaks Using Cyclic Memory Allocation”. In: *Proceedings of the 6th International Symposium on Memory Management*. 2007, pp. 15–30.
- [117] M. Nica, S. Nica, and F. Wotawa. “On the Use of Mutations and Testing for Debugging”. In: *Software: Practice and Experience* 43.9 (2013), pp. 1121–1142.
- [118] G. Novark, E. Berger, and B. Zorn. “Exterminator: Automatically Correcting Memory Errors with High Probability”. In: *ACM SIGPLAN Notices* 42.6 (2007), pp. 1–11.
- [119] Y. Padioleau, J. Lawall, R. Hansen, and G. Muller. “Documenting and Automating Collateral Evolutions in Linux Device Drivers”. In: *ACM SIGOPS Operating Systems Review* 42.4 (2008), pp. 247–260.
- [120] J. H. Perkins, G. Sullivan, W.-F. Wong, Y. Zibin, M. D. Ernst, M. Rinard, S. Kim, S. Larsen, S. Amarasinghe, J. Bachrach, M. Carbin, C. Pacheco, F. Sherwood, and S. Sidiroglou. “Automatically Patching Errors in Deployed Software”. In: *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles - SOSP '09* (2009), p. 87.
- [121] B. Pernici and A. M. Rosati. “Automatic Learning of Repair Strategies for Web Services”. In: *Proceedings of the Fifth European Conference on Web Services*. 2007, pp. 119–128.
- [122] M. Pezzè, M. C. Rinard, W. Weimer, and A. Zeller. *Self-Repairing Programs, Report From Dagstuhl Seminar*. Tech. rep. 11062. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2011.
- [123] Y. Qi, X. Mao, and Y. Lei. “Efficient Automated Program Repair Through Fault-Recorded Testing Prioritization”. In: *Proceedings of ICSM*. 2013.
- [124] Y. Qi, X. Mao, and Y. Lei. “Program Repair As Sound Optimization of Broken Programs”. In: *International Symposium on Theoretical Aspects of Software Engineering*. 2009.
- [125] Y. Qi, X. Mao, Y. Lei, Z. Dai, and C. Wang. “The Strength of Random Search on Automated Program Repair”. In: *Proceedings of the 36th International Conference on Software Engineering*. 2014, pp. 254–265.
- [126] Z. Qi, F. Long, S. Achour, and M. Rinard. “An Analysis of Patch Plausibility and Correctness for Generate-And-Validate Patch Generation Systems”. In: *Proceedings of ISSTA*. 2015.
- [127] F. Qin, J. Tucek, J. Sundaresan, and Y. Zhou. “Rx: Treating Bugs As Allergies—a Safe Method to Survive Software Failures”. In: *ACM SIGOPS Operating Systems Review* 39.5 (2005), pp. 235–248.

- [128] B. Randell. “System Structure for Software Fault Tolerance”. In: *IEEE Transactions on Software Engineering* 2 (1975), pp. 220–232.
- [129] M. Rebaudengo, M. S. Reorda, M. Violante, and M. Torchiano. “A Source-to-source Compiler for Generating Dependable Software”. In: *International Workshop on Source Code Analysis and Manipulation*. 2001, pp. 33–42.
- [130] M. Rinard, C. Cadar, D. Dumitran, D. Roy, T. Leu, and W. Beebee Jr. “Enhancing Server Availability and Security Through Failure-oblivious Computing”. In: *Proceedings of the 6th Conference on Symposium on Operating Systems, Design & Implementation*. 2004, pp. 21–21.
- [131] M. Rinard, M. Ernst, and J. Perkins. *Collaborative Learning for Security and Repair in Application Communities*. Tech. rep. Massachusetts Institute of Technology, 2011.
- [132] M. Rinard. “Acceptability-oriented Computing”. In: *ACM SIGPLAN Notices* 38.12 (2003), pp. 57–75.
- [133] M. C. Rinard. *Survival Techniques for Computer Programs*. Tech. rep. MIT, 2006.
- [134] M. Rinard, C. Cadar, D. Dumitran, D. M. Roy, and T. Leu. “A Dynamic Technique for Eliminating Buffer Overflow Vulnerabilities (and Other Memory Errors)”. In: *Proceedings of the 20th Annual Computer Security Applications Conference*. 2004, pp. 82–90.
- [135] J. Rohr. “STAREX Self-repair Routines: Software Recovery in the JPL-STAR Computer”. In: *Proceedings of FTCS*. 1973.
- [136] H. Samimi, M. Schäfer, S. Artzi, T. D. Millstein, F. Tip, and L. J. Hendren. “Automated Repair of HTML Generation Errors in PHP Applications Using String Constraint Solving”. In: *Proceedings of ICSE*. 2012, pp. 277–287.
- [137] E. Schulte, S. Forrest, and W. Weimer. “Automated Program Repair Through the Evolution of Assembly Code”. In: *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering*. 2010, pp. 313–316.
- [138] A. Sethi et al. “A Survey of Fault Localization Techniques in Computer Networks”. In: *Science of Computer Programming* 53.2 (2004), pp. 165–194.
- [139] A. Shaw, D. Doggett, and M. Hafiz. “Automatically Fixing C Buffer Overflows Using Program Transformations”. In: *International Conference on Dependable Systems and Networks*. 2014, pp. 124–135.
- [140] M. Shaw. “Self-healing: Softening Precision to Avoid Brittleness”. In: *Proceedings of the First Workshop on Self-healing Systems*. 2002, pp. 111–114.

- [141] S. Sicard, F. Boyer, and N. De Palma. “Using Components for Architecture-based Management: the Self-repair Case”. In: *Proceedings of the 30th International Conference on Software Engineering*. 2008, pp. 101–110.
- [142] S. Sidiroglou and A. Keromytis. “Countering Network Worms Through Automatic Patch Generation”. In: *Security & Privacy* 3.6 (2005), pp. 41–49.
- [143] S. Sidiroglou, O. Laadan, C. Perez, N. Viennot, J. Nieh, and A. Keromytis. “Assure: Automatic Software Self-healing Using Rescue Points”. In: *ACM Sigplan Notices* 44.3 (2009), pp. 37–48.
- [144] S. Sidiroglou, M. Locasto, S. Boyd, and A. Keromytis. “Building a Reactive Immune System for Software Services”. In: *Proceedings of the USENIX Annual Technical Conference*. Vol. 161. 2005.
- [145] S. Sidiroglou-Douskos, E. Lahtinen, F. Long, and M. Rinard. “Automatic Error Elimination by Horizontal Code Transfer Across Multiple Applications”. In: *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 2015, pp. 43–54.
- [146] S. Sidiroglou-Douskos, S. Misailovic, H. Hoffmann, and M. Rinard. “Managing Performance Vs. Accuracy Trade-offs with Loop Perforation”. In: *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering*. 2011, pp. 124–134.
- [147] J. Silva. “A Survey on Algorithmic Debugging Strategies”. In: *Advances in Engineering Software* 42.11 (2011), pp. 976–991.
- [148] M. A. A. da Silva, A. Mougénot, X. Blanc, and R. Bendraou. “Towards Automated Inconsistency Handling in Design Models”. In: *Proceedings of the 22nd International Conference on Advanced Information Systems Engineering*. 2010, pp. 348–362.
- [149] R. Singh, S. Gulwani, and A. Solar-Lezama. “Automated Feedback Generation for Introductory Programming Assignments”. In: *ACM SIGPLAN Notices*. Vol. 48. 6. 2013, pp. 15–26.
- [150] A. Smirnov and T. Chiueh. “DIRA: Automatic Detection, Identification, and Repair of Control-hijacking Attacks”. In: *The 12th Annual Network and Distributed System Security Symposium*. 2005.
- [151] E. K. Smith, E. Barr, C. Le Goues, and Y. Brun. “Is the Cure Worse Than the Disease? Overfitting in Automated Program Repair”. In: *Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 2015.
- [152] A. Somayaji, S. Hofmeyr, and S. Forrest. “Principles of a Computer Immune System”. In: *Proceedings of the 1997 Workshop on New Security Paradigms*. 1998, pp. 75–82.

- [153] S. Son, K. S. McKinley, and V. Shmatikov. “Fix Me Up: Repairing Access-Control Bugs in Web Applications.” In: *Proceedings of the Network and Distributed System Security Symposium*. 2013.
- [154] H. Sözer, B. Tekinerdoğan, and M. Akşit. “FLORA: a Framework for Decomposing Software Architecture to Introduce Local Recovery”. In: *Software: Practice and Experience* 39.10 (2009), pp. 869–889.
- [155] M. Staats, M. W. Whalen, and M. P. E. Heimdahl. “Programs, Tests, and Oracles: the Foundations of Testing Revisited”. In: *Proceedings of the International Conference on Software Engineering*. 2011, pp. 391–400.
- [156] M. Stumptner and F. Wotawa. “A Model-based Approach to Software Debugging”. In: *Proceedings on the Seventh International Workshop on Principles of Diagnosis*. 1996.
- [157] B. Sun, R. Chang, X. Chen, and A. Podgurski. “Automated Support for Propagating Bug Fixes”. In: *Proceedings of the International Symposium on Software Reliability Engineering*. 2008, pp. 187–196.
- [158] B. Sun, G. Shu, A. Podgurski, S. Li, S. Zhang, and J. Yang. “Propagating Bug Fixes with Fast Subgraph Matching”. In: *Proceedings of the International Symposium on Software Reliability Engineering*. 2010, pp. 21–30.
- [159] G. J. Sussman. *Building Robust Systems An Essay*. 2007.
- [160] S. Tallam, C. Tian, R. Gupta, and X. Zhang. “Avoiding Program Failures Through Safe Execution Perturbations”. In: *International Conference on Computer Software and Applications*. 2008.
- [161] S. H. Tan and A. Roychoudhury. “Relifix: Automated Repair of Software Regressions”. In: *Proceedings of ICSE*. 2015.
- [162] Y. Tao, J. Kim, S. Kim, and C. Xu. “Automatically Generated Patches As Debugging Aids: a Human Study”. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 2014, pp. 64–74.
- [163] S. Thomas and L. Williams. “Using Automated Fix Generation to Secure SQL Statements”. In: *Proceedings of the Third International Workshop on Software Engineering for Secure Systems*. 2007, p. 9.
- [164] W. Torres-Pomales et al. *Software Fault Tolerance: a Tutorial*. Tech. rep. NASA-2000-tm210616. NASA, 2000.
- [165] J. Tran, M. Godfrey, E. Lee, and R. Holt. “Architectural Repair of Open Source Software”. In: *Proceedings of the International Workshop on Program Comprehension*. 2000, pp. 48–59.
- [166] K. Tsipenyuk, B. Chess, and G. McGraw. “Seven Pernicious Kingdoms: a Taxonomy of Software Security Errors”. In: *Security & Privacy* 3.6 (2005), pp. 81–84.

- [167] F. Wang and C.-H. Cheng. “Program Repair Suggestions From Graphical State-Transition Specifications”. In: *Proceedings of FORTE 2008*. 2008.
- [168] T. Wang, C. Song, and W. Lee. “Diagnosis and Emergency Patch Generation for Integer Overflow Exploits”. In: *Detection of Intrusions and Malware, and Vulnerability Assessment*. 2014, pp. 255–275.
- [169] Y. Wang, Y. Huang, and C. Kintala. “Progressive Retry for Software Failure Recovery in Message-passing Applications”. In: *IEEE Transactions on Computers* 46.10 (1997), pp. 1137–1141.
- [170] Y. Wei, Y. Pei, C. A. Furia, L. S. Silva, S. Buchholz, B. Meyer, and A. Zeller. “Automated Fixing of Programs with Contracts”. In: *Proceedings of the International Symposium on Software Testing and Analysis*. 2010.
- [171] W. Weimer, T. Nguyen, C. L. Goues, and S. Forrest. “Automatically Finding Patches Using Genetic Programming”. In: *Proceedings of the International Conference on Software Engineering*. 2009.
- [172] W. Weimer, S. Forrest, C. Le Goues, and T. Nguyen. “Automatic Program Repair with Evolutionary Computation”. In: *Communications of the ACM* 53.5 (2010), p. 109.
- [173] W. Wong and V. Debroy. “A Survey on Software Fault Localization”. In: *University of Texas at Dallas, Tech. Rep. UTDCS-45-09* (2009).
- [174] A. Zeller, Y. Wei, B. Meyer, M. Nordio, C. A. Furia, and Y. Pei. “Automated Fixing of Programs with Contracts”. In: *IEEE Transactions on Software Engineering* 40.5 (2014), pp. 427–449.
- [175] X. Zhang, N. Gupta, and R. Gupta. “Locating Faults Through Automated Predicate Switching”. In: *Proceedings of the 28th International Conference on Software Engineering*. 2006, pp. 272–281.
- [176] Y. Zhang and Y. Ding. “CTL Model Update for System Modifications”. In: *Journal of Artificial Intelligence Research* 31.1 (2008), pp. 113–155.
- [177] H. Zhong and Z. Su. “An Empirical Study on Real Bug Fixes”. In: *Proceedings of the 37th International Conference on Software Engineering*. 2015.
- [178] F. Zhou, J. Condit, Z. Anderson, I. Bagrak, R. Ennals, M. Harren, G. Necula, and E. Brewer. “SafeDrive: Safe and Recoverable Extensions Using Language-based Techniques”. In: *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. 2006, pp. 45–60.