

Towards Extending the Range of Bugs That Automated Program Repair Can Handle

Omar I. Al-Bataineh
Simula Research Laboratory
Oslo, Norway
omar@simula.no

Leon Moonen
Simula Research Laboratory
Oslo, Norway
leon.moonen@computer.org

Abstract—Modern *automated program repair* (APR) is well-tuned to finding and repairing bugs that introduce *observable* erroneous behavior to a program. However, a significant class of bugs does not lead to such observable behavior (e.g., liveness/termination bugs, non-functional bugs, and information flow bugs). Such bugs can generally not be handled with current APR approaches, so, as a community, we need to develop complementary techniques.

To stimulate the systematic study of alternative APR approaches and hybrid APR combinations, we devise a novel bug classification system that enables methodical analysis of their bug detection power and bug repair capabilities. To demonstrate the benefits, we analyze the repair of termination bugs in sequential and concurrent programs. The study shows that integrating dynamic APR with formal analysis techniques, such as termination provers and software model checkers, reduces complexity and improves the overall reliability of these repairs.

Index Terms—automated program repair, bug classification, non-observable and liveness bugs, hybrid techniques.

I. INTRODUCTION

Corrective maintenance, i.e., finding and repairing software defects, is one of the main categories of software maintenance, and responsible for a large part of the overall costs of software development [1]. *Automated program repair* (APR) promises to increase developer productivity and drastically reduce the costs of corrective maintenance [2, 3]. Despite the advances of APR for real-world programs [4], these approaches can only handle certain types of bugs because they generally rely on dynamic analysis for functional verification, where a test suite is used to simulate the input and monitor the output to check correct behavior. However, this is only viable if the effects of a bug can be observed when executing the program.

Detecting and repairing *non-observable bugs* and *liveness bugs* (i.e., bugs that do not lead to incorrect results or crashes) pose a far greater challenge. For example, identifying a liveness bug requires finding an infinite execution that will never satisfy the desired liveness property [5]. It is not known how long one would need to run the program to reveal an existing liveness bug, making it impractical to find such cases using dynamic analysis. Another critical challenge that makes detection and repair of liveness bugs notoriously hard is that the effects that a liveness bug is triggered are generally unobservable (i.e., they typically produce little debugging information). One option for finding this class of bugs is applying formal program

analysis techniques that use correctness specifications to detect liveness bugs. Such a rigorous analysis can both help to detect the presence of liveness bugs as well as assure the absence of these bugs in automatically generated patches.

The question which (combinations of) techniques will be most effective at handling certain bugs is an open research question that forms the foundation of this paper. To stimulate the systematic study of alternative APR approaches and hybrid APR combinations, we devise a novel bug classification system that enables methodical analysis of their bug detection power and bug repair capabilities. In earlier work, various bug classification schemes were developed to understand when and why specific bugs arise, and how they are fixed. These classifications use a number of criteria, such as cause-impact [6, 7], severity-priority [8], and bug complexity [9]. However, since they were designed for different goals, they do not capture the properties required to determine if a bug is amenable to a particular technique. To that end, we introduce a bug classification system explicitly aimed at comparing different techniques and evaluating the feasibility of their integration.

Contributions: (1) We propose a *novel bug classification system* based on three fundamental properties: *bug observability*, *bug reproducibility*, and *bug tractability*. This classification provides the APR community with a tool to methodologically explore and compare alternative and hybrid APR approaches by (i) analyzing the detection power of different bug detection techniques, (ii) distinguishing APR approaches based on their bug repair capabilities, and (iii) providing a common terminology that helps identify gaps in current APR research.

(2) We discuss four *APR approaches* that can handle different classes of bugs: *dynamic APR*, *static APR*, *dynamic-static APR*, and *formal APR*. Moreover, we identify the conditions under which each approach can be effectively applied.

(3) To demonstrate the benefits of our method, we study *termination bugs* in sequential and concurrent programs, and sketch novel *hybrid APR algorithms* for repairing such bugs. The study shows that termination bugs in *sequential* programs can be effectively addressed using *dynamic-static APR*, by first generating plausible patches using test cases and then using termination provers [10–12] to check their correctness. The non-deterministic nature of termination bugs in *concurrent* programs makes them challenging for dynamic analysis, and they are best addressed with *formal APR* that combines termination provers with software model checkers [13–19].

II. BUG CLASSIFICATION SCHEMES

There are many different ways to expose bugs in programs, including manual inspection, dynamic analysis (testing), static analysis, model checking, or a combination of these techniques. Effective bug classification schemes can help understand why bugs arise and how to fix them. Classification can also help identify the most appropriate analysis technique for handling each class of bugs. Next, we discuss three existing bug classification systems, analyze their limitations, and introduce a new classification system that addresses them:

- 1) *Cause-impact criteria* [6, 7]: Bugs are classified based on their cause: algorithmic, concurrency, memory, generic programming, and unknown, as well as based on their impact: security, performance, failure, and unknown.
- 2) *Severity and priority criteria*: This classification is used in many bug tracking systems [8]. Severity indicates the impact of the bug on the program’s functionality and can be categorized as critical, major, moderate, minor, etc. Priority indicates how soon the bug should be fixed and is categorized into levels such as low, medium, and high.
- 3) *Bug complexity criteria* [9]: These criteria distinguish four main categories: (i) easy to detect, easy to repair bugs, (ii) easy to detect, difficult to repair bugs, (iii) difficult to detect, easy to repair bugs, and (iv) difficult to detect, difficult to repair bugs.

These three existing bug classification systems were not designed for comparing the capabilities and limitations of different bug detection or program repair techniques. As a result, their criteria do not capture the specific properties needed to determine whether a bug b is amenable to a particular technique T . To address this gap, we propose a new bug classification system that is based on three key properties of bugs, namely *bug observability*, *bug reproducibility*, and *bug tractability*. In Sections III and IV, these properties are then used to analyze the power of different bug detection techniques and APR approaches.

Before proceeding further, let us first define a program bug. We base ourselves on a specification of *expected behavior*: the expected responses (output) of the program to a given input.

DEFINITION 1: Program bug. Let P be a program, I a set of inputs, and φ_{beh} be a specification of expected behavior of P . We say that P suffers from a bug iff there exists at least one input $i \in I$ that leads to an execution trace under which program P violates φ_{beh} , formalised as $(P, i) \not\models \varphi_{beh}$.

Since a complete specification of a program’s expected behavior is often not available, test cases are generally used to model the expected behavior of a program P . We assume there exists (i, o_{exp}) , where o_{exp} is the expected output for input $i \in I$. When the observed output $o_{obs} = (P, i)$ does not match the expected output o_{exp} , we say that P contains a bug.

DEFINITION 2: Observable bug. Let P be a program containing bug b . We say that b is an observable bug iff there exists an execution of P where, in a finite number of execution steps, the erroneous behavior of b can be seen by an observer O .

DEFINITION 3: Classifying bugs by observability. We classify bugs based on the notion of observability in three types:

- 1) observable bugs whose erroneous behavior is fully observable in finite execution steps (e.g., arithmetic bugs),
- 2) partially observable bugs whose erroneous behavior is only partially observable at runtime because the faulty trace is infinite, so not all output of the program can be observed (e.g., termination bugs),
- 3) non-observable bugs whose erroneous behavior is fully unobservable at runtime (e.g., non-functional bugs).

Note that observability is a relative notion that depends on the observation power of O . In the simplest case, the observer can witness the output produced by P and the corresponding execution time. If we increase the observation power of O (i.e., the amount of information O can gather about the program’s execution), some non-observable bugs may become observable. For example, bugs that adversely affect the memory or energy consumption can easily go unnoticed during the execution of the program. Such bugs can be exposed by using monitoring at the virtual machine or operating system level [20–23]. Alternatively, the program can be augmented with additional variables and checks that help to keep track of these non-functional aspects at runtime [24]. However, increasing observation power also increases the chance of affecting the program’s execution [25].

We distinguish five common types of observable erroneous behavior: $EB = \{crash, exception, incorrectResult, softHang, hardHang\}$. While most types in EB are easy to understand, we will define the notions of *soft* and *hard* hang bugs. Hang bugs are a particular type of bugs that concern (temporary or permanent) lack of progress in observable behavior [26, 27]. Hang bugs can have various causes, such as iteration errors or communication deadlocks. To define hang bugs, we use a temporal specification [28, 29] that checks if any of the locations where the program might terminate can be reached.

DEFINITION 4: Halting Statements. We refer to a statement s in a program P as a halting statement iff the expected behavior of P is that execution terminates after executing statement s .

Examples of halting statements include special termination statements such as `exit`, or simply the final statement in a program. Observe that programs whose expected behavior is to never terminate have no halting statements (e.g., a webserver).

DEFINITION 5: Hang bugs. Let P be a program with a set of inputs I and H be the set of halting statements of P . Let φ_{temp} be a temporal property that puts an upper bound on the execution time of P , and φ_{reach} be a temporal property that checks whether P reaches a halting statement. Let also $EB' = EB \setminus \{softHang, hardHang\}$. We distinguish:

- 1) **Soft hang bugs** occur when there exists an input $i \in I$ that makes P unresponsive for a finite amount of time before execution is resumed and a halting statement is reached:

$$S : (P, i) \not\models \varphi_{temp} \wedge (P, i) \models \varphi_{reach} \wedge \text{output}(P, i) \notin EB'$$
- 2) **Hard hang bugs**, also known as **termination bugs**, occur when there exists an input $i \in I$ that makes P unresponsive

for an unbounded amount of time, never resuming to normal execution or reaching a halting statement:

$$H : (P, i) \not\models \varphi_{temp} \wedge (P, i) \not\models \varphi_{reach} \wedge \text{output}(P, i) \notin EB'$$

Hang bugs are also referred to as *liveness violations* in model checking and formal program analysis literature [30–33], and we will use these terms interchangeably in the remainder.

We now turn to discuss the property of bug reproducibility.

DEFINITION 6: Bug reproducibility. Let P be a program containing a bug b and t_b be a test case that exposes b . We say that b is a reproducible or deterministic bug iff every time program P is executed under test t_b , bug b is exposed and the same erroneous behavior is observed. On the other hand, we say that b is a hard-to-reproduce or non-deterministic bug iff bug b is exposed in rare circumstances when repeating the execution of P under test t_b (i.e., the result of program P depends not only on the code of P but also on the timing of the execution).

Reproducible bugs are easy to detect, provided that the bug is observable (see Definition 2). Not surprisingly, hard-to-reproduce bugs are also hard to detect. Arithmetic bugs are examples of *easy-to-reproduce* bugs, while concurrency bugs are examples of *hard-to-reproduce* bugs. The last property we study is *bug tractability*, which depends on the *depth* of the bug and the size of the faulty trace it produces.

DEFINITION 7: Bug tractability. Let P be a program containing bug b and L be the set of reachable locations of P and $\ell_b \in L$ be the buggy location to the bug b . We say that the trace of b is a tractable trace iff for each execution of P that is buggy to b , the number of execution steps that are required to reach ℓ_b is bounded and that ℓ_b is not part of a loop that can be executed infinitely often. On the other hand, we say that the trace of b is intractable iff ℓ_b is visited infinitely often during the execution of P (i.e., ℓ_b is part of an infinite loop).

The presence of loops plays a crucial role in determining the tractability of a bug. The size of faulty traces for non-loop programs is typically shorter than those in loop programs. Based on the size of the faulty trace, we can further distinguish the class of tractable bugs: (i) *shallow bugs* are tractable bugs with finite short faulty traces, (ii) *deep bugs* are tractable bugs with finite but long faulty traces. For example, a bug in a loop program P that does not occur until a vast number of iterations are executed can be viewed as an example of a deep bug.

Bug Classification System: Table I summarises the three properties in our classification system with their distinguishing attributes and impact on bug detection. Our classification system associates each bug with a three-tuple of concrete attributes for $\{\text{observability}, \text{reproducibility}, \text{tractability}\}$. For example, an arithmetic bug has the properties: $\{\text{observable}, \text{easy-to-reproduce}, \text{shallow}\}$.

III. BUG DETECTION TECHNIQUES

Various bug detection techniques can be used to expose bugs in programs. In this work, we are interested in studying three well-known bug detection techniques: dynamic analysis, static analysis, and model checking. While dynamic analysis detects

bugs in programs by executing them, static analysis and model checking use different techniques in bug detection that perform bug checking statically, without running the program. We start by discussing the requirements needed to expose bugs in each technique, the advantages and disadvantages of the techniques, and the theoretical foundation and detection power of each (i.e., the classes of bugs that each technique can handle).

A. Dynamic Analysis

Dynamic analysis is a technique to identify bugs and vulnerabilities in programs by exercising various runs through the program based on valid inputs. Dynamic analysis can be performed using a test suite, which can be developed manually or via test case generation, or through fuzzing, which systematically explores a large amount of automatically generated tests. Fuzzing is one of the most common methods used to find vulnerabilities in programs [34, 35]. Early fuzz testing was based on sending random inputs to a program to check if it could be made to crash. The techniques have evolved to systematically explore the input space using knowledge from the source code or input formats to discover bugs that are hidden deep in the code.

Dynamic analysis has several advantages over the other program analysis techniques: (i) the program behavior can be monitored, and bugs can be exposed while the program is running; (ii) it allows for analysis of programs for which we do not have access to the actual code; (iii) it can be conducted against any program; and last but not least, (iv) it can identify bugs that are hard to find using static analysis. We now discuss the conditions under which bugs may be discovered in dynamic analysis.

DEFINITION 8: Bug detection in dynamic analysis. Let P be a program containing bug b , and D be a dynamic program analysis checker (i.e., an automated testing tool such as a fuzzer). We say that b is detectable in D iff

- 1) P is given in an executable form,
- 2) bug b is observable in some executions of P ,
- 3) there exists a test suite T containing at least one failing test t by which bug b can be exposed in P .

However, many bugs whose detection requires the analysis of infinite traces or the satisfaction of complex composite properties cannot be found using an approach solely relying on test cases. Examples of such classes of bugs include (i) bugs in non-executable programs, (ii) liveness bugs such as termination and starvation bugs, and (iii) non-observable bugs such as non-functional and information flow bugs [36, 37].

OBSERVATION 1: *Dynamic analysis techniques can handle observable classes of bugs with finite execution traces, provided that a testing mechanism is implemented by which the bug can be exposed, and provided that the program is executable.*

To address these limitations, we need to pair dynamic analysis with complementary bug detection techniques to improve the detection power of the approach. The remainder of this section discusses static analysis and model checking, which instead of using the program itself as in dynamic analysis,

TABLE I
A SUMMARY OF THE KEY PROPERTIES FOR BUG CLASSIFICATION WITH THEIR ATTRIBUTES AND IMPACT

bug property	property attributes	impact on bug detection techniques
observability	{observable, partially-observable, non-observable}	affects detection power
reproducibility	{easy-to-reproduce, hard-to-reproduce}	affects efficiency and scalability
tractability	{shallow, deep, liveness}	affects detection power and efficiency

analyzes abstractions of the program, to improve observability and tractability and enable the detection of deeper bugs [38].

B. Static Program Analysis

Static program analysis is an approach for analyzing a computer program without actually executing it. The most significant advantage of static analysis is the ability to quickly and automatically examine the complete code of the program to find flaws that might be missed by dynamic analysis. The literature on static program analysis for bug detection is rich and mature [39–41]. Many of these techniques build on automatically evaluated analysis rules and bug detection patterns that capture the general conditions under which specific bugs can occur, providing a systematic way for their detection.

To capture the notion of bug detectability in static analysis, one needs to ensure the availability of the source code of the buggy program at which the bug b occurs and the availability of some valid solid theory for the detection of bug b .

DEFINITION 9: Bug detection in static analysis. Let P be a program containing bug b and S be a static program analyzer that can be used to expose bugs of type b . We say that b is a bug detectable by analyzer S if all of the following conditions hold:

- 1) the source code of P is available, and
- 2) S contains sound and complete detection method for b , and
- 3) P is written in a language that is accepted by S .

The increasing complexity of (loop) programs and the large variety of vulnerabilities make it difficult for static code analyzers to detect and identify vulnerabilities in a precise manner. One of the most significant disadvantages of the static code analysis methodology is the presence of false-positive warnings: the tool may signal possible bugs where there are none. However, reducing the number of false positives in static analysis tools is still an open problem.

OBSERVATION 2: *Static checkers can handle observable, non-observable, reproducible, and non-reproducible bugs, provided that the checker is built based on some solid mathematical foundation, and provided that the source code of the program is available in a programming language accepted by the checker.*

Dynamic and static analysis techniques have different bug detection powers as they rely on distinct assumptions and use orthogonal detection methods. Overall, these techniques have complementary strengths and weaknesses that are worth combining to improve the reliability of APR systems.

C. Model Checking

Model checking [42–44] is an automated formal method for checking whether a finite-state model of a system meets a given specification. The technique has been used successfully

to debug complex computer hardware, concurrent systems, and real-world safety-critical systems.

Model checking has several advantages. It can detect errors that would be very difficult to notice with other methods, such as in concurrent programs. The properties that can be verified are more expressive than with traditional testing, depending on the formalism used to express them. For example, properties that require something to happen infinitely often, or properties that require that some alternative is always available. In addition, because every possible behavior of the model is checked, the result is inevitable, provided that the model checking tool itself has no serious errors.

Model checking can be an expensive procedure in a repair process because of its exhaustive nature. Expressing both the model of the system and the properties formally requires great care and expertise. Moreover, one of the most significant problems with model checking in practice is the so-called “state explosion problem”: When the number of state variables in the system increases, the size of the system state space grows exponentially. However, abstractions can be applied to bring the verification within feasible bounds of model checking technology.

Software model checking tools [13–19] verify the correctness of software models in a rigorous and automated fashion. Most tools construct a (symbolic) reachability graph for the program-under-analysis (i.e., a graph that contains reachable run-time states of the program) *without* running the program. This graph is then used to check if a property of interest holds. They typically implement sophisticated data structures that enable clever search algorithms and optimizations.

The answer returned by a model checker is either a notion of a successful verification (i.e., the specification holds), or a counterexample – an execution path that violates a given property. However, if the program being verified has an infinite state space, certain types of abstractions are needed, or the analysis may simply not terminate.

DEFINITION 10: Bug detection in model checking. Let P be a program containing bug b and SMC be a software model checker that can be used to expose bugs of type b . We say that b is a detectable bug in the model checker SMC iff:

- 1) a formal property φ_b is available that is written in the input specification language of SMC, which captures the conditions under which b can occur, and
- 2) the source code of p is available, and
- 3) P has finite states or an equivalent finite abstract program P_{abs} can be constructed for the properties of interest, and
- 4) the program P or its reduced equivalent program P_{abs} is written in a modeling language that is acceptable by SMC.

TABLE II
A SUMMARY OF KEY DIFFERENCES AND SIMILARITIES OF DYNAMIC ANALYSIS, STATIC ANALYSIS, AND MODEL CHECKING

criteria	dynamic analysis	static analysis	model checking
bug detection mechanism	test cases	pattern-based specifications	formal correctness specifications
accuracy of the analysis	accurate	inaccurate (suffers from false positives)	accurate relative to the accuracy of model
the need of code availability	not needed	needed to perform the analysis	needed to construct the model
the need of bug observability	needed to detect the bug	not needed (performs code analysis)	not needed (performs state analysis)
code coverage of the program	incomplete code coverage	complete code coverage	complete code coverage
the need of code executability	needed to detect the bug	the program will not be executed	the program will not be executed
language dependency	language-independent	language dependent	language dependent
automation of the analysis	can be automated (fuzz testing)	fully automated	a model needs to be manually written

OBSERVATION 3: *Model checking tools can handle observable, non-observable, reproducible, and non-reproducible bugs, provided that the size of the program is finite or a sound abstraction can be developed to bring the program within the feasibility bound of model checking, and provided that a specification is available for the bug of interest.*

Bug Detection Properties: Based on the characteristics of the three bug detection techniques discussed, one can make the following general observations. Model checking is more expensive than static analysis, requiring longer running times and more resources. Static analysis is not as accurate as model checking, and testing is not as complete as model checking. Testing suffers from coverage challenges: it is challenging to cover all possible executions of the program, in particular for programs with an infinite input space where this becomes prohibitively expensive. To ensure the correctness of the program for all inputs, a correctness specification must exist that can be formally analyzed.

Table II summarises the key differences and similarities between the three program analysis techniques using several criteria: code coverage, the need for code executability, accuracy of the analysis, the need for bug observability, the need for code availability, and automation of the technique. By code coverage, we mean the number of feasible execution paths of the program that the technique can cover during the analysis, and accuracy indicates whether the detected bug is a real bug.

IV. APR APPROACHES

This section describes four APR approaches that can handle different classes of bugs. The four approaches combine the detection power of dynamic analysis, static analysis, and model checking techniques to improve the reliability of existing APR techniques. We discuss the applicability of these approaches to three classes of bugs: arithmetic bugs (observable bugs), non-functional bugs (non-observable bugs), and liveness bugs (partially observable bugs).

An APR approach generally consists of four steps: fault identification, fault localization, patch generation, and patch validation. The most challenging step in the APR process is the patch validation step, in which the generated patch is extensively evaluated to ensure that the bug is resolved, and that the patch does not introduce any unwanted behavior. In dynamic APR, as the name implies, the patch validation step is primarily performed using test cases. Since these rarely capture the expected behavior in full detail, the technique suffers from

the so-called patch overfitting problem, where the patched program may pass the tests in the given test suite, while it is failing for valid inputs not covered by the test suite. It is therefore desirable to combine the power of different analysis techniques while taking into account the distinctive properties of each class of bugs. This leads to our examination of the following four APR approaches:

- 1) *dynamic APR*: in which fault identification and patch validation are performed using dynamic analysis techniques. GenProg [45] is an example of a dynamic APR tool;
- 2) *static APR*: in which fault identification and patch validation are performed using static analysis techniques;
- 3) *dynamic-static APR*: in which fault identification and patch validation are performed using a combination of static and dynamic analysis, i.e., using test cases and static analysis tools developed for the same class of bugs;
- 4) *formal APR*: in which fault identification and patch validation are performed using formal methods and verification techniques such as model checking.

A hybrid APR approach aims to improve the overall quality of the generated repairs and alleviate the patch overfitting challenge of dynamic APR systems.

A. Arithmetic Bugs

Arithmetic calculations affect a wide variety of applications, including safety-critical systems such as control systems for vehicles, medical equipment, and industrial plants. The key properties of arithmetic bugs can be summarised as follows:

- 1) Arithmetic bugs are observable classes of bugs or can be easily made observable to external observers: the root causes of arithmetic bugs are limited and easy to identify.
- 2) Arithmetic bugs are tractable bugs with finite traces.
- 3) Arithmetic bugs introduce various erroneous behavior to the program in which they occur: they may cause the program to crash or may produce incorrect outputs.

These properties make them directly amenable to dynamic (i.e., test-based) APR. Note that many of the available dynamic APR systems rely explicitly or implicitly on observability strategies to expose this class of bugs (e.g., integer overflow, division by zero, etc.). There are also several static analysis tools that can handle arithmetic bugs. Thus, arithmetic bugs can be repaired using dynamic APR, static APR, or dynamic-static APR. These repair approaches differ mainly in the correctness specification used to validate generated patches for the detected arithmetic bug. In dynamic APR, the spec-

ification is captured by the test cases, while in static APR, the specification is captured by the formal bug detection rules.

EXAMPLE 1: **Integer overflow (IO)** is a type of arithmetic bug that occurs when the computation of an arithmetic operation, such as multiplication or addition, exceeds the maximum size of the integer type used to store it. IO bugs are an observable class of bugs, and thus they are amenable to dynamic APR. IO bugs are also amenable to static APR, and there are several reliable static analysis tools available that can address IO bugs [46, 47]. Thus, devising a hybrid static-dynamic APR system for IO bugs is feasible and will help increase confidence about the soundness of the generated repairs.

OBSERVATION 4: Arithmetic bugs are observable, tractable, and reproducible classes of bugs with finite execution traces. They are amenable to both dynamic and static APR since the root causes of arithmetic bugs are easy to identify.

B. Non-functional Bugs

Non-functional bugs [24, 48, 49] are a class of bugs that affect the way a program operates, rather than the functional behavior of the program. Inefficiently written loops in programs and synchronization issues in concurrent programs (i.e., using a large unnecessary number of locks) can be viewed as examples of non-functional bugs. Non-functional bugs are as important as functional bugs. For example, energy consumption saving is getting more urgent, particularly for applications running on embedded systems and IoT in Smart Cities.

Fixing non-functional bugs is generally more complex than fixing functional bugs, since non-functional bugs can hide themselves well in the code. While most functional bugs can be detected through observing the erroneous behavior of bugs, a large percentage of non-functional bugs are detected through manual code review [50–52]. Non-functional bugs usually do not generate incorrect results or crashes. Therefore, they cannot be observed by checking the program output. The key properties of non-functional bugs can be described as follows.

- 1) Non-functional bugs are generally non-observable classes of bugs: they do not introduce direct observable erroneous behavior to the program in which they occur.
- 2) Non-functional bugs increase the anticipated running cost of a program (execution time, memory and energy consumption, etc.) due to the inefficient use of resources.

Depending on what quality attributes are considered, programs may suffer from many different types of non-functional bugs. For example, consider the class of non-functional bugs that adversely affect the run-time costs of executing the program, such as execution time, memory consumption, and energy consumption. To expose such types of bugs, the program may need to be augmented with additional variables or online monitors that can be used to observe aspects at runtime [24].

Therefore, there is a need to develop effective bug detection tools that can be used to expose non-functional bugs at the early stages of the software development life cycle. Specifically, this requires efficient profiling techniques and oracles that help decide whether the program’s non-functional requirements are met under a particular workload. Unfortunately, the

lack of effective test oracles for non-functional bugs is a well-known problem that will need to be addressed in the future.

OBSERVATION 5: Non-functional bugs are an example of non-observable classes of bugs (i.e., a program that suffers from a non-functional bug executes normally and terminates normally). Thus, they are not directly amenable to traditional dynamic bug detection techniques that rely on test cases and observing a program’s outputs.

C. Liveness Bugs

In this section, we discuss a class of bugs that has received little attention from the APR community, namely liveness bugs. A *liveness property* asserts that “something good will eventually occur when executing a program” [30, 31]. Freedom of starvation and program termination are examples of liveness properties. A program that violates a liveness property cannot make progress and thus suffers from a *liveness bug*.

Two fundamental properties make detecting and repairing liveness bugs far more challenging than other classes of bugs. First, the behavioral effects of triggering a liveness bug are generally unobservable. Second, identifying a liveness bug requires finding an infinite execution that will never satisfy the desired liveness property [5], making it impractical to find such bugs using dynamic analysis. Therefore, detecting and repairing liveness bugs generally require more sophisticated repair algorithms since they must be able to generate a finite representation of infinite counterexamples.

To better understand the complexity of repairing liveness bugs, we study a subset of liveness bugs known as termination bugs. There are two advantages to this choice: On the one hand, a termination bug is a specific type of liveness bug whose repair is essential for ensuring software reliability. On the other hand, by examining techniques for handling termination bugs, we gain knowledge that can help address other liveness bugs. Termination bugs have the following specific properties:

- 1) Termination bugs are *partially observable*: an observer monitoring the behavior of a non-terminating loop program will not witness any erroneous behavior but rather experience unexpectedly long execution times. The only observable behavior of termination bugs is that the program becomes *non-responsive at runtime*.
- 2) Termination bugs have *infinite faulty traces*: a counterexample to a termination property violation is infinite.

OBSERVATION 6: Termination bugs are an example of partially observable bugs with infinite execution traces. Termination bugs are amenable to dynamic-static APR and static APR.

V. HYBRID APR FOR TERMINATION BUGS

This section describes a hybrid program repair approach for termination bugs that combines the strengths of termination provers with those of software model checkers. Such a combination has two key advantages. First, it considerably reduces the overall computational complexity of the problem by avoiding the exhaustive exploration of the program’s input space. Second, it helps avoid the known overfitting problem by generating verified repairs for termination bugs.

The presence of loops in programs can complicate the detection of certain classes of bugs. Recall the two types of hang bugs from Definition 5; It is not clear how one can distinguish between the following two types of loops using a dynamic analysis technique: (i) inefficiently written loops that introduce a soft hang bug, and (ii) incorrect infinite loops that introduce a hard hang or termination bug.

Earlier work [53] that evaluated the effectiveness of different APR tools on the ManyBugs and IntroClass datasets, used a simple timeout mechanism to handle termination bugs in these two datasets: when the execution time of a program exceeds some pre-specified period, they consider the program to be *likely* non-terminating due to an infinite loop. Marcote and Monperrus instrument loops with iteration counters that are monitored to detect infinite loops [54]. Both options *can* lead to false conclusions about the program under analysis (i.e., even when the watchdog triggers, the program may not have a termination bug but suffer from a soft hang bug). Moreover, hang bugs can have complicated causes: programs that become unresponsive may contain deadlocks, infinite loops, or other bugs that lead to non-termination but are *not* infinite loops.

An effective solution for addressing termination bugs is to apply *termination provers*: tools that can check combinations of many complex termination criteria. They take a program as input and return one of three answers: *terminating* (*TR*), *non-terminating* (*NT*), or *unknown* (*UN*). In general, when the prover returns definite answer for a given program (i.e., $answer \in \{TR, NT\}$), the answer is valid with high confidence. Termination provers have been successfully used to analyse termination of a wide variety of loop programs [55–65].

DEFINITION 11: Valid Termination Bug Repair. Let P be a buggy non-terminating program with a set of inputs I , φ_{beh} a specification of expected behavior of P , and φ_{reach} a specification that checks the reachability of some halting statement of P . We say that P' is a valid repair of P iff for every input $i \in I$ we have $(P', i) \models \varphi_{beh}$ and $(P', i) \models \varphi_{reach}$.

In other words, the patched version P' should preserve the expected behavior of P , and it should terminate.

A. Termination Bugs in Sequential Programs

To repair termination bugs in sequential programs, we first generate plausible patches, and then use termination provers to check the correctness of these patches. AProVE [10] and 2LS [11] are among the most reliable candidates to analyze the termination of sequential programs.

DEFINITION 12: Validity of Patches for Termination Bugs in Sequential Programs. Let P_s be a non-terminating sequential program and $T = (T_p \cup T_f)$ be a test-suite consisting of passing test cases T_p and failing test cases T_f . Let P'_s be a candidate patch of P_s and TP be a termination prover that returns one of the verification answers $\{TR, NT, UN\}$. We say that P'_s is a valid patch of P_s iff all of the following conditions hold:

- 1) all failing test cases from T_f pass on program P'_s ,
- 2) none of the passing test cases from T_p fail on program P'_s ,

- 3) termination prover TP returns “*TR*” when analyzing termination of P'_s .

The soundness of generated patches for termination bugs in sequential programs can be captured formally as:

$$\varphi_{seq} = (\forall t \in T (P'_s \vdash t) \wedge TP(P'_s) = TR) \quad (1)$$

where T is the set of available test cases, and $P'_s \vdash t$ indicates that patch P'_s runs successfully against test t .

Figure 1 sketches a hybrid repair procedure for termination bugs in sequential programs. The algorithm takes five inputs: the buggy sequential program P_s , a specification of expected behavior φ_{beh} , a test suite T_s , a termination prover TP , and the allocated time budget $TimeBudget$. It uses two functions: (i) $faultLocalizer(P_s, CE, T_s)$ computes the set of suspicious statements $SuspStats$ whose mutation may lead to generate a valid patch. It finds these suspicious statements by combining the counterexamples CE generated by termination prover TP with the results of spectrum-based fault localization [66–69] on the test suite (executed with a timeout mechanism to avoid getting stuck in infinite loops); (ii) the function $mutate(P_s, SuspStats)$ is used to construct the patch space (i.e., patch generation) by mutating the computed set $SuspStats$ that may affect the truth value of the termination condition of the detected buggy non-terminating loop in P_s .

B. Termination Bugs in Concurrent Programs

In the automated repair of concurrent programs, the goal is to generate a patch that ensures that a concurrent program is correct under all interleavings. It is difficult, if not impossible, to examine all possible executions of a concurrent program using dynamic analysis techniques. Therefore, a concurrent program cannot be debugged and repaired in the same manner as sequential programs. In the case of concurrency, it usually refers to action interleaving. That is, if the processes p_i and p_j are in parallel composition ($p_i \parallel p_j$) then the actions of these will be interleaved. Each process executes a sequence of actions (sub-program), then the set of possible interleavings of several processes consists of all possible sequences of

```

1: Inputs:  $P_s, \varphi_{beh}, T_s, TP, TimeBudget$ 
2: Output:  $P_{repaired}$ 
3:  $PatchSpace := \emptyset, P_{repaired} := NoPatchFound, found := false$ 
4:  $(*, CE) := TP(P_s)$  # get CE, verdict known to be NT
5:  $SuspStats := faultLocalizer(P_s, CE, T_s)$ 
6:  $PatchSpace := mutate(P_s, SuspStats)$  # patch generation
7: while  $PatchSpace \neq \emptyset \wedge TimeBudget > 0 \wedge \neg found$  do
8:   select  $patch$  from  $PatchSpace$ 
9:   if  $patch \models \varphi_{beh}$  then # 2-step patch validation
10:    if  $TP(patch) = (TR, *)$  then
11:       $P_{repaired} := patch$ 
12:       $found := true$ 
13:    end if
14:   end if
15: end while
16: return  $P_{repaired}$ 

```

Fig. 1. Repair algorithm for sequential programs

actions. Before proceeding further, let us introduce the notion of successful termination in concurrent programs.

DEFINITION 13: Termination of Concurrent Programs. Let $P_c = (p_1 \parallel p_2 \parallel \dots \parallel p_n)$ be a concurrent program that consists of a collection of sub-programs, where each process p_i executes a sub-program. Let H_i be the set of halting statements at the sub-program executed by process p_i . We say that P_c is terminating iff every sub-program is eventually terminating. That is, the sub-program executed by process p_i eventually executes some halting statement $s \in H_i$ and terminates.

DEFINITION 14: Termination Failures in Concurrent Programs. Let $P_c = (p_1 \parallel p_2 \parallel \dots \parallel p_n)$ be a concurrent program that consists of a collection of sub-programs executed by processes p_1, \dots, p_n . The program P_c fails to terminate iff:

- 1) there exists a logical bug that leads to an infinite loop, or
- 2) there exists a concurrency bug, such as deadlock or livelock, that prevents the program from making any further progress to reach a halting statement and terminate.

Thus, if deadlocks and livelocks are formally proven to never occur in the program-under-analysis, and all loops in the sub-programs are proven to be terminating, then one can conclude that the concurrent program is terminating. The distinction between the two causes of termination failures in concurrent programs (logical bug or concurrency bug) helps to select a strategy for fixing the detected termination bug.

Combining Model Checking and Termination Provers:

Repairing termination bugs in concurrent programs can be a computationally complex task. This is mainly because termination bugs in concurrent programs can be caused by either a logical or concurrency bug. Furthermore, the vast number of possible interleavings of parallel processes of a given concurrent program can increase the complexity of the repair problem. Therefore, it is necessary to employ both termination provers and model checking to reduce the computational complexity of the problem. Fortunately, we know how to write specifications to check the absence of concurrency bugs in concurrent programs [70–72]. In Definition 15, we describe the conditions that are necessary to ensure the correctness of the generated patches for termination bugs in the buggy non-terminating concurrent program.

DEFINITION 15: Validity of Generated Patches for Termination Bugs in Concurrent Programs. Let $P_c = (p_1 \parallel p_2 \parallel \dots \parallel p_n)$ be a buggy non-terminating concurrent program, TP be a termination prover, and SMC be a software model checker. Let $\varphi_{\text{deadlock}}$ and $\varphi_{\text{livelock}}$ be specifications that check respectively the absence of deadlocks and livelocks in P_c and φ_{beh} be a specification that captures the expected behavior of P_c . We say that a candidate patch P'_c for the buggy program P_c is a valid patch iff it meets the following requirements:

- 1) the prover TP returns “terminating” when analysing termination of the sub-program executed by process p_i , and
- 2) the checker SMC returns “holds” when checking the specifications $\varphi_{\text{deadlock}}$ and $\varphi_{\text{livelock}}$ against P'_c , and
- 3) the checker SMC returns “holds” when checking the specification of expected behavior φ_{beh} against the patch P'_c .

Formally, we can capture the above described requirements in a 4-part correctness specification of the following form:

$$\varphi_{\text{con}} = \forall_{p_i \in P'_c} (TP(p_i) = TR) \wedge SMC(P'_c, \varphi_{\text{deadlock}}) \wedge SMC(P'_c, \varphi_{\text{livelock}}) \wedge SMC(P'_c, \varphi_{\text{beh}}) \quad (2)$$

A key challenge when dealing with termination bugs is to ensure that the generated repair guarantees not only the termination of the program for each possible input, but also the semantic preservation of the program. This requires the analysis of a composite correctness property that checks both termination and semantic preservation. By using formula (2), we entirely avoid the patch overfitting problem, which is one of the major problems of dynamic APR. Formula (2) uses sequential termination provers to check the absence of infinite loops in each individual process. This can be performed while abstracting away concurrency details that are irrelevant to the local computations of processes. Note that one cannot prove the termination of program P_c by simply applying a sequential termination prover: a sound proof of termination must consider all possible interactions among the sub-programs of P_c .

An alternative way to use model checking in detecting and repairing termination bugs of the concurrent program is to reduce the termination problem to the reachability analysis problem. That is, to check whether each process will eventually reach some halting location and terminate. However, the feasibility of the approach relies mainly on the size and number of processes of the buggy program under analysis (i.e., computing state-reachability is known to be PSPACE-complete when processes are finite state [73]). The reduction of the termination problem to the reachability problem in model checking leads to the following temporal formula

$$\varphi'_{\text{con}} = \forall_{p_i \in P'_c} (\mathbf{AF}(p_i \cdot \ell_h^{(j)} \mid \ell_h^{(j)} \in H_i)) \wedge SMC(P'_c, \varphi_{\text{beh}}) \quad (3)$$

where \mathbf{A} is a temporal path quantifier which should be read as “for all paths”, and \mathbf{F} is the “future” temporal operator [28]. Intuitively, formula (3) checks whether for each reachable execution path of process p_i , some halting location $\ell_h^{(j)} \in H_i$ will eventually be reached. However, the proper termination analysis of the concurrent program P_c using formula (3) requires the precise computation of the set of halting locations H_i for each process p_i of the patched concurrent program P'_c .

There are several software model checkers¹ that can be used to verify reachability properties and detect concurrency bugs, including VeriSoft [14], Java Pathfinder [16], CMC [17], DIVINE [19], and GMC [74]. DIVINE is a modern, explicit-state model checker that can verify programs written in multiple real-world programming languages, including C and C++. On the other hand, GMC is a model checker based on the generic Monte-Carlo model-checking algorithm. It takes as input a C program, the target program to be verified, and the linear temporal logic specification that needs to be checked.

There are also a few termination provers that can be used to analyze the termination of concurrent programs. For instance,

¹ Unlike traditional model checking, a software model checker does not require a user to manually construct an abstract model of the program to be checked, but instead, the tool works directly on the program’s source code.

Cook et al. [75] have extended the termination prover T2 [12] to support the analysis of concurrent programs, which can be used to validate generated patches for termination bugs in concurrent programs. Termination checker T2 supports nested loops, recursive functions, pointers, side-effects, and function-pointers, as well as concurrent programs. Of course, the prover cannot handle termination of all concurrent programs since the general problem is undecidable. The use of concurrent termination provers leads to the following specification

$$\varphi''_{con} = \forall p_i \in P'_c (TP(p_i) = TR) \wedge SMC(P'_c, \varphi_{beh}) \quad (4)$$

The termination provers AProVE, 2LS, and T2 can be viewed as complementary tools: it is possible that some tool fails to detect certain forms of termination bugs while other succeeds, depending on the implemented theory and the complexity of the program under analysis. Therefore, termination checkers can be run in parallel to expose termination bugs.

While checking the satisfaction of formula (2) may require higher computational complexity than formulas (3) and (4) (i.e., it employs both termination provers and software model checkers to check the absence of deadlocks, livelocks, and infinite loops), it has several advantages. First, it helps identify the root causes of non-termination in the program-under-repair. Second, the patch validation approach that uses formula (2) can benefit from the counterexamples generated by both termination provers and software model checkers. This helps to develop effective program synthesis for termination bugs.

Repairing Termination Bugs in Concurrent Programs: To fix termination bugs in concurrent programs, it is first essential to identify the root cause of the termination bug since both logical and concurrency bugs can cause non-termination. On the one hand, there are several mechanisms for handling deadlocks in concurrent programs [71, 72, 76]. On the other, repair algorithms based on genetic programming can help fix termination bugs that occur due to infinite loops [45, 77].

Figure 2 sketches a hybrid repair procedure for termination bugs in concurrent programs. It uses three helper functions: (i) *faultLocalizer*(P_c, CE) finds the statements that are suspected to be faulty in P_c using the counterexamples CE generated by termination prover TP and software model checker SMC , (ii) *mutateConcur*($P_c, SuspStatsD$) constructs the patch space for a concurrency bug by mutating the synchronization primitives of the program, and (iii) *mutateLogic*($P_c, SuspStatsL$) constructs the patch space for a logical bug by mutating expressions that affect the control of faulty loops.

Generating verified repairs of termination bugs in both sequential and concurrent programs is a challenging open problem that requires formal analysis techniques. Our preliminary investigation of using termination provers in sequential programs shows that this is a rather promising approach that can be effectively integrated with dynamic APR approaches to generate verified repairs for termination bugs. The application of the termination provers 2LS and AProVE on the programs in the two datasets, SNU real-time benchmark and the PowerStone benchmark suite [78], show that the tools are able to successfully prove termination of around 85% of the examined

```

1: Inputs:  $P_c, \varphi_{beh}, \varphi_{deadlock}, TP, SMC, TimeBudget$ 
2: Output:  $P_{repaired}$ 
3:  $PatchSpace := \emptyset, P_{repaired} := NoPatchFound, found := false$ 
4:  $(Outcome_1, CE_1) := SMC(P_c, \varphi_{deadlock})$ 
5: if  $(Outcome_1 = fails)$  then #  $P_c$  contains a deadlock
6:    $SuspStatsD := faultLocalizer(P_s, CE_1)$ 
7:    $PatchSpace := mutateConcur(P_c, SuspStatsD)$ 
8: end if
9:  $(Outcome_2, CE_2) := TP(P_c)$ 
10: if  $Outcome_2 = NT$  then #  $P_c$  contains an infinite loop
11:    $SuspStatsL := faultLocalizer(P_s, CE_2)$ 
12:    $PatchSpace := PatchSpace \cup mutateLogic(P_c, SuspStatsL)$ 
13: end if
14: while  $PatchSpace \neq \emptyset \wedge TimeBudget > 0 \wedge \neg found$  do
15:   select  $patch$  from  $PatchSpace$ 
16:   if  $patch \models \varphi_{beh}$  then # 2-step patch validation
17:     if  $TP(patch) = (TR, *) \wedge$ 
18:        $SMC(patch, \varphi_{deadlock}) = (holds, *)$  then
19:        $P_{repaired} := patch, found := true$ 
20:     end if
21:   end if
22: end while
23: return  $P_{repaired}$ 

```

Fig. 2. Repair algorithm for concurrent programs

programs using very little computational time (a few seconds). This demonstrates the feasibility of using termination provers to validate the generated patches of termination bugs.

Table III compares the complexity of termination bugs in both sequential and concurrent programs using several criteria: (i) root causes of termination bugs, (ii) feasible APR approach to be applied for termination bugs in both classes of programs, (iii) patch validation procedure to validate generated patches, (iv) patch validation tool that can be used to check termination.

VI. RELATED WORK

We discuss the related literature on APR, bug classification, use of formal specifications in APR, previous attempts to integrate different analysis techniques, and termination analysis. **Automated Program Repair:** Source-based, automated program repair approaches [3] can be separated into search-based and semantic-based approaches. *Search-based* approaches such as GenProg [45], Astor [79], and SCRepair [77] predominantly use failing test cases to identify bugs, and then mutate the source code until the program passes all failing test cases. They do not provide patch correctness guarantees beyond the fact that the provided test cases now pass. Recent work introduced property-based testing to strengthen the validation of candidate repairs and address overfitting [80]. Nevertheless, these approaches require executing the program-under-repair, first to find the bug, and then to generate and validate candidate repairs. *Semantic-based* approaches like SemFix [81], Nopol [82], DirectFix [83], SPR [84], Angelix [85], and JFIX [86] infer repair constraints for the buggy program via symbolic execution of the given tests. The completeness of inferred constraints relies on the available test suite. Similarly, Infinitel [54] uses an SMT solver to synthesise a loop termination condition and then uses test cases for patch validation.

TABLE III
A COMPARISON OF THE COMPLEXITY OF THE REPAIR PROBLEM OF TERMINATION BUGS IN SEQUENTIAL AND CONCURRENT PROGRAMS

program class	root causes of non-termination	feasible APR	patch validation procedure	patch validation tools
sequential programs	infinite loops	dynamic-static	test cases and termination provers	AProVE, 2LS, T2, and GMC
concurrent programs	infinite loops, deadlocks, or livelocks	static or formal	termination provers and SMC	T2 and GMC

Bug Classification Systems: Several works target bug classification using a wide variety of classification criteria and for different goals. The work of Tan et al. [6, 7] introduced a bug classification system based on the cause-impact criteria. They studied bug characteristics of around 2,060 real-world bugs in three representative open-source projects. They concluded that semantic bugs are the dominant root cause of bugs, and memory-related bugs have decreased due to the development of effective detection tools. Many bug tracking systems classify bugs using severity and priority criteria, where severity indicates the seriousness of the bug on the program functionality and priority indicates how soon the bug should be fixed [8]. Cotroneo et al. present a maintenance-oriented bug classification system in which the characteristics of the bug manifestation are studied [9]. The study identifies the set of failure-exposing conditions under which a bug may occur. Neither of these classification systems considers properties that can be used to analyze the detection power of different bug detection techniques and the conditions under which they can be integrated, which we add with the classification system proposed in this paper. Abbaspour et al. [87] do use bug observability to study the erroneous behavior of a wide variety of concurrency bugs. However, the authors restrict themselves to concurrency bugs and do not study the detection power of different bug detection techniques.

Integrating Bug Detection Techniques: Few attempts have been made to integrate different program analysis techniques to alleviate the impact of the patch overfitting problem. Al-Bataineh et al. used the static detection patterns/rules as the source for formulating formal specifications and discussed the possibility of integrating static and dynamic analysis techniques to improve the overall quality of generated patches [47]. There also exists a few examples of using information from debugging to aid APR: Facebook’s APR tool SapFix takes information generated during the bug detection process and applies various techniques, including templates specific to given bug types, to fix the program [4].

Termination Analysis of Programs: A huge body of work has been published on proving termination of programs based on a variety of techniques, such as abstract interpretation [55, 58, 59], bounds analysis [60, 61], ranking functions [62, 63], recurrence sets [64, 65], and transition invariants [56, 57]. Based on these techniques, a number of program termination checkers have been developed in the prior literature including AProVE [10], 2LS [11], T2 [12], and ARMC [88]. Unfortunately, termination provers have not yet been used to *validate* the generated candidate patches of termination bugs in the previous APR approaches. We strongly believe that integrating APR approaches with contemporary termination

checkers would help advance the current state-of-the-art of APR, not only for repairing termination bugs, but also for other classes of concurrency bugs. This is mainly because fixing termination bugs in concurrent programs would ensure the absence of certain classes of concurrency bugs, such as deadlock and livelock bugs. It would also help avoid the patch overfitting problem by generating verified repairs for termination bugs in both sequential and concurrent programs.

VII. CONCLUDING REMARKS

A significant class of bugs cannot be handled with current APR approaches, and there is a need to study complementary techniques. To stimulate this work, we propose a novel bug classification system based on three key properties: bug observability, bug tractability, and bug reproducibility. This provides a tool to methodologically explore and compare alternative and hybrid APR approaches by (i) analyzing the detection power of different bug detection techniques, (ii) distinguishing APR approaches based on their bug repair capabilities, and (iii) providing a common terminology that helps identify gaps in current APR research. Moreover, it allows analysis of how techniques can be combined to handle challenging classes of bugs. As a demonstrating example, we study termination bugs in sequential and concurrent programs, and present novel *hybrid* algorithms for their repair by integrating termination provers and software model checkers in the APR pipeline. Such integration reduces the complexity of the repair algorithms and improves the overall reliability.

We identify the following directions for future research: (1) we are in the process of empirically validating the ideas described in this work by combining a selection of APR tools with some of the termination provers and software model checkers mentioned earlier; (2) there is a need for efficient fault localization mechanisms for termination and liveness bugs, since these create infinite traces that cannot be handled with the current spectrum-based fault localization approaches; We currently circumvent this with a timeout mechanism, but more efficient techniques could find a more precise set of suspicious statements; (3) the combination of CounterExample-Guided Inductive Synthesis (CEGIS) [89] with termination provers and software model checkers could enable efficient patch space exploration without the user guidance that is normally needed for CEGIS; (4) aside from the termination bugs studied in Section V, various other bugs cannot be handled by APR approaches solely based on dynamic analysis. One particularly interesting class of non-observable bugs for future research are security-related vulnerabilities where sensitive information may be disclosed to unauthorized parties as a result of violations of information flow security [36, 37].

REFERENCES

- [1] E. B. Swanson. "The Dimensions of Maintenance." In: *Int'l Conference on Software Engineering (ICSE)*. IEEE, 1976, pp. 492–497.
- [2] C. Le Goues, M. Pradel, and A. Roychoudhury. "Automated Program Repair." In: *Communications of the ACM* 62.12 (2019), pp. 56–65. DOI: 10/gkgf29.
- [3] M. Monperrus. "Automatic Software Repair: A Bibliography." In: *ACM Computing Surveys* 51.1 (2018), pp. 1–24. DOI: 10/ggssbj.
- [4] A. Marginean, J. Bader, S. Chandra, M. Harman, Y. Jia, K. Mao, A. Mols, and A. Scott. "SapFix: Automated End-to-End Repair at Scale." In: *Int'l Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. 2019, pp. 269–278. DOI: 10/gkgf2c.
- [5] B. Alpern and F. B. Schneider. "Recognizing Safety and Liveness." In: *Distributed Computing* 2.3 (1987), pp. 117–126. DOI: 10/fmbptq.
- [6] Z. Li, L. Tan, X. Wang, S. Lu, Y. Zhou, and C. Zhai. "Have Things Changed Now? An Empirical Study of Bug Characteristics in Modern Open Source Software." In: *Workshop on Architectural and System Support for Improving Software Dependability*. ACM, 2006, pp. 25–33. DOI: 10/cqsw86.
- [7] L. Tan, C. Liu, Z. Li, X. Wang, Y. Zhou, and C. Zhai. "Bug Characteristics in Open Source Software." In: *Empirical Software Engineering* 19.6 (2014), pp. 1665–1705. DOI: 10/f6m38x.
- [8] N. Serrano and I. Ciordia. "Bugzilla, ITracker, and Other Bug Trackers." In: *IEEE Software* 22.2 (2005), pp. 11–13. DOI: 10/fbgs99.
- [9] D. Cotroneo, K. Trivedi, S. Russo, and R. Pietrantuono. "How Do Bugs Surface? A Comprehensive Study on the Characteristics of Software Bugs Manifestation." In: *Journal of Systems and Software* 113 (2016), pp. 27–43. DOI: 10/gpd4t5.
- [10] J. Giesl et al. "Proving Termination of Programs Automatically with AProVE." In: *Int'l Joint Conference on Automated Reasoning (IJCAR)*. Ed. by S. Demri, D. Kapur, and C. Weidenbach. Springer, 2014, pp. 184–191. DOI: 10/f3ssz2.
- [11] H.-Y. Chen, C. David, D. Kroening, P. Schrammel, and B. Wachter. "Synthesising Interprocedural Bit-Precise Termination Proofs (T)." In: *Int'l Conference on Automated Software Engineering (ASE)*. 2015, pp. 53–64. DOI: 10/gp6krb.
- [12] M. Brockschmidt, B. Cook, S. Ishtiaq, H. Khlaaf, and N. Piterman. "T2: Temporal Property Verification." In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. Ed. by M. Chechik and J.-F. Raskin. Springer, 2016, pp. 387–393. DOI: 10/gp6kq9.
- [13] R. Jhala and R. Majumdar. "Software Model Checking." In: *ACM Computing Surveys* 41.4 (2009), pp. 1–54. DOI: 10/fd3pxq.
- [14] P. Godefroid. "VeriSoft: A Tool for the Automatic Analysis of Concurrent Reactive Software." In: *Int'l Conference on Computer Aided Verification (CAV)*. Ed. by O. Grumberg. Springer, 1997, pp. 476–479. DOI: 10/b7ntq3.
- [15] G. Holzmann. "The Model Checker SPIN." In: *IEEE Transactions on Software Engineering* 23.5 (1997), pp. 279–295. DOI: 10/d7wqxt.
- [16] K. Havelund and T. Pressburger. "Model Checking JAVA Programs Using JAVA PathFinder." In: *International Journal on Software Tools for Technology Transfer* 2.4 (2000), pp. 366–381. DOI: 10/d2pmx6.
- [17] M. Musuvathi, D. Y. W. Park, A. Chou, D. R. Engler, and D. L. Dill. "CMC: A Pragmatic Approach to Model Checking Real Code." In: *Symposium on Operating Systems Design and Implementation (OSDI)*. Usenix, 2002.
- [18] S. J. Thompson, G. Brat, and A. Venet. "Software Model Checking of ARINC-653 Flight Code with MCP." In: *NASA Formal Methods Symposium*. NASA, 2010.
- [19] Z. Baranová, J. Barnat, K. Kejstová, T. Kučera, H. Lauko, J. Mrázek, P. Ročkal, and V. Štill. "Model Checking of C and C++ with DIVINE 4." In: *Automated Technology for Verification and Analysis*. Ed. by D. D'Souza and K. Narayan Kumar. Springer, 2017, pp. 201–207. DOI: 10/gp6kq7.
- [20] Y. Hebbal, S. Laniece, and J.-M. Menaud. "Virtual Machine Introspection: Techniques and Applications." In: *Int'l Conference on Availability, Reliability and Security*. 2015, pp. 676–685. DOI: 10/gp6kq4.
- [21] P. Dovgalyuk, N. Fursova, I. Vasiliev, and V. Makarov. "QEMU-based Framework for Non-Intrusive Virtual Machine Instrumentation and Introspection." In: *Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, 2017, pp. 944–948. DOI: 10/gp6kqz.
- [22] B. Gregg. *Systems Performance: Enterprise and the Cloud*. 2nd Ed. Addison-Wesley, 2020.
- [23] B. Gregg. *BPF Performance Tools: Linux System and Application Observability*. Addison-Wesley.
- [24] O. Al-Bataineh, D. J. X. Ng, and A. Easwaran. "Monitoring Cumulative Cost Properties." In: *Int'l Conference on Formal Methods in Software Engineering (FormalSE)*. 2021, pp. 19–30. DOI: 10/gp6kqw.
- [25] T. Mytkowicz, P. F. Sweeney, M. Hauswirth, and A. Diwan. *Observer Effect and Measurement Bias in Performance Analysis*. Technical Report CU-CS-1042-08. University of Colorado at Boulder, 2008.
- [26] X. Wang, Z. Guo, X. Liu, Z. Xu, H. Lin, X. Wang, and Z. Zhang. "Hang Analysis: Fighting Responsiveness Bugs." In: *ACM SIGOPS Operating Systems Review* 42.4 (2008), pp. 177–190. DOI: 10/b9chm8.
- [27] D. J. Dean, P. Wang, X. Gu, W. Enck, and G. Jin. "Automatic Server Hang Bug Diagnosis: Feasible Reality or Pipe Dream?" In: *IEEE Int'l Conference on Autonomic Computing*. 2015, pp. 127–132. DOI: 10/gp6kq5.
- [28] A. Pnueli. "The Temporal Logic of Programs." In: *Annual Symposium on Foundations of Computer Science (SFCS)*. 1977, pp. 46–57. DOI: 10/dn8cpn.
- [29] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1992.
- [30] L. Lamport. "Proving the Correctness of Multiprocess Programs." In: *IEEE Transactions on Software Engineering* SE-3.2 (1977), pp. 125–143. DOI: 10/d25dpw.
- [31] B. Alpern and F. B. Schneider. "Defining Liveness." In: *Information Processing Letters* 21.4 (1985), pp. 181–185. DOI: 10/d97bw4.
- [32] C. Killian, J. W. Anderson, R. Jhala, and A. Vahdat. "Life, Death, and the Critical Transition: Finding Liveness Bugs in Systems Code." In: *USENIX Symposium on Networked Systems Design & Implementation (NSDI 07)*. 2007.
- [33] P. Li and J. Regehr. "T-Check: Bug Finding for Sensor Networks." In: *Int'l Conference on Information Processing in Sensor Networks*. ACM, 2010, pp. 174–185. DOI: 10/djrwkg.
- [34] B. P. Miller, L. Fredriksen, and B. So. "An Empirical Study of the Reliability of UNIX Utilities." In: *Communications of the ACM* 33.12 (1990), pp. 32–44. DOI: 10/fqnt9s.
- [35] B. P. Miller, G. Cooksey, and F. Moore. "An Empirical Study of the Robustness of MacOS Applications Using Random Testing." In: *Int'l Workshop on Random Testing*. ACM, 2006, pp. 46–54. DOI: 10/bw7w9h.
- [36] A. Sabelfeld and A. Myers. "Language-Based Information-Flow Security." In: *IEEE Journal on Selected Areas in Communications* 21.1 (2003), pp. 5–19. DOI: 10/dnvjwg.
- [37] G. Smith. "Principles of Secure Information Flow Analysis." In: *Malware Detection*. Ed. by M. Christodorescu, S. Jha, D. Maughan, D. Song, and C. Wang. Springer, 2007, pp. 291–307. DOI: 10/fnfjff.
- [38] C. David, P. Kesseli, D. Kroening, and M. Lewis. "Danger Invariants." In: *Formal Methods*. Ed. by J. Fitzgerald, C. Heitmeyer, S. Gnesi, and A. Philippou. Springer, 2016, pp. 182–198. DOI: 10/gp6kqr.
- [39] V. D'Silva, D. Kroening, and G. Weissenbacher. "A Survey of Automated Techniques for Formal Software Verification." In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 27.7 (2008), pp. 1165–1178. DOI: 10/frt7ww.
- [40] A. Bessey et al. "A Few Billion Lines of Code Later: Using Static Analysis to Find Bugs in the Real World." In: *Communications of the ACM* 53.2 (2010), pp. 66–75. DOI: 10/bj8r36.
- [41] C. Sadowski, E. Aftandilian, A. Eagle, L. Miller-Cushon, and C. Jaspán. "Lessons from Building Static Analysis Tools at Google." In: *Communications of the ACM* 61.4 (2018), pp. 58–66. DOI: 10/ggsshq.
- [42] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. "Symbolic Model Checking: 10²⁰ States and Beyond." In: *Information and Computation* 98.2 (1992), pp. 142–170. DOI: 10/bvrsx5.
- [43] B. Bérard, M. Bidoit, A. Finkel, F. Laroussinie, A. Petit, L. Petrucci, P. Schnoebelen, and P. McKenzie. *Systems and Software Verification: Model-Checking Techniques and Tools*. Springer, 2001.
- [44] E. Clarke, O. Grumberg, D. Kroening, D. Peled, and H. Veith. *Model Checking*. 2nd Ed. MIT Press, 2018.
- [45] C. Le Goues, T. Nguyen, S. Forrest, and W. Weimer. "GenProg: A Generic Method for Automatic Software Repair." In: *IEEE Transactions on Software Engineering* 38.1 (2012), pp. 54–72. DOI: 10/cfzt3f.
- [46] P. Muntean, M. Monperrus, H. Sun, J. Grossklags, and C. Eckert. "IntRepair: Informed Repairing of Integer Overflows." In: *IEEE Transactions on Software Engineering* 47.10 (2021), pp. 2225–2241. DOI: 10/gh97rm.

- [47] O. I. Al-Bataineh, A. Grishina, and L. Moonen. "Towards More Reliable Automated Program Repair by Integrating Static Analysis Techniques." In: *IEEE Int'l Conference on Software Quality, Reliability and Security (QRS)*. 2021, pp. 654–663. DOI: 10/gp6kq6.
- [48] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu. "Understanding and Detecting Real-World Performance Bugs." In: *ACM SIGPLAN Notices* 47.6 (2012), pp. 77–88. DOI: 10/f372jr.
- [49] A. Radu and S. Nadi. "A Dataset of Non-Functional Bugs." In: *Int'l Conference on Mining Software Repositories (MSR)*. IEEE, 2019, pp. 399–403. DOI: 10/gp6kq2.
- [50] M. E. Fagan. "Design and Code Inspections to Reduce Errors in Program Development." In: *IBM Systems Journal* 15.3 (1976), pp. 182–211. DOI: 10/btfv3v.
- [51] T. Gilb, D. Graham, and S. Finzi. *Software Inspection*. Addison-Wesley, 1993.
- [52] A. Bacchelli and C. Bird. "Expectations, Outcomes, and Challenges of Modern Code Review." In: *Int'l Conference on Software Engineering (ICSE)*. 2013, pp. 712–721. DOI: 10/gf2h2r.
- [53] 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* 41.12 (2015), pp. 1236–1256. DOI: 10/gpd4jv.
- [54] S. R. L. Marcote and M. Monperrus. *Automatic Repair of Infinite Loops*. 2015. DOI: 10/jb2f. arXiv: 1504.05078.
- [55] J. Berdine, A. Chawdhary, B. Cook, D. Distefano, and P. O'Hearn. "Variance Analyses from Invariance Analyses." In: *SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 2007, pp. 211–224. DOI: 10/frk5km.
- [56] D. Kroening, N. Sharygina, A. Tsitovich, and C. M. Wintersteiger. "Termination Analysis with Compositional Transition Invariants." In: *Computer Aided Verification*. Ed. by T. Touili, B. Cook, and P. Jackson. Springer, 2010, pp. 89–103. DOI: 10/b54tf5.
- [57] A. Podelski and A. Rybalchenko. "Transition Invariants." In: *Symposium on Logic in Computer Science*. 2004, pp. 32–41. DOI: 10/fbsbdm.
- [58] A. Chawdhary, B. Cook, S. Gulwani, M. Sagiv, and H. Yang. "Ranking Abstractions." In: *European Symposium on Programming Languages and Systems (ESOP)*. Ed. by S. Drossopoulou. Springer, 2008, pp. 148–162. DOI: 10/d4wv5s.
- [59] A. Tsitovich, N. Sharygina, C. M. Wintersteiger, and D. Kroening. "Loop Summarization and Termination Analysis." In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. Ed. by P. A. Abdulla and K. R. M. Leino. Springer, 2011, pp. 81–95. DOI: 10/cd8sdb.
- [60] S. Gulwani, S. Jain, and E. Koskinen. "Control-Flow Refinement and Progress Invariants for Bound Analysis." In: *ACM SIGPLAN Notices* 44.6 (2009), pp. 375–385. DOI: 10/fxdmzw.
- [61] S. Gulwani, K. K. Mehra, and T. Chilimbi. "SPEED: Precise and Efficient Static Estimation of Program Computational Complexity." In: *ACM SIGPLAN Notices* 44.1 (2009), pp. 127–139. DOI: 10/fxgdsx.
- [62] A. R. Bradley, Z. Manna, and H. B. Sipma. "Linear Ranking with Reachability." In: *Int'l Conference on Computer Aided Verification (CAV)*. Ed. by K. Etessami and S. K. Rajamani. Springer, 2005, pp. 491–504. DOI: 10/fpv58.
- [63] P. Cousot. "Proving Program Invariance and Termination by Parametric Abstraction, Lagrangian Relaxation and Semidefinite Programming." In: *Verification, Model Checking, and Abstract Interpretation*. Ed. by R. Cousot. Springer, 2005, pp. 1–24. DOI: 10/dzfp66.
- [64] A. Gupta, T. A. Henzinger, R. Majumdar, A. Rybalchenko, and R.-G. Xu. "Proving Non-Termination." In: *SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 2008, pp. 147–158. DOI: 10/dd76q5.
- [65] W. R. Harris, A. Lal, A. V. Nori, and S. K. Rajamani. "Alternation for Termination." In: *Static Analysis Symposium (SAS)*. Ed. by R. Cousot and M. Martel. Springer, 2010, pp. 304–319. DOI: 10/djthds.
- [66] J. A. Jones, M. J. Harrold, and J. Stasko. "Visualization of Test Information to Assist Fault Localization." In: *Int'l Conference on Software Engineering (ICSE)*. 2002, p. 467. DOI: 10/bx264c.
- [67] L. Naish, H. J. Lee, and K. Ramamohanarao. "A Model for Spectra-Based Software Diagnosis." In: *ACM Transactions on Software Engineering and Methodology (TOSEM)* 20.3 (2011), pp. 1–32. DOI: 10/c5xnmnd.
- [68] S. Heiden, L. Grunske, T. Kehler, F. Keller, A. van Hoorn, A. Filieri, and D. Lo. "An Evaluation of Pure Spectrum-Based Fault Localization Techniques for Large-Scale Software Systems." In: *Software: Practice and Experience* 49.8 (2019), pp. 1197–1224. DOI: 10/gp6kq8.
- [69] X. Xie and B. Xu. *Essential Spectrum-based Fault Localization*. Springer, 2021. DOI: 10/jb2d.
- [70] A. Gupta, V. Kahlon, S. Qadeer, and T. Touili. "Model Checking Concurrent Programs." In: *Handbook of Model Checking*. Ed. by E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem. Springer, 2018, pp. 573–611. DOI: 10/jkpx.
- [71] Y. Lin and S. S. Kulkarni. "Automatic Repair for Multi-Threaded Programs with Deadlock/Livelock Using Maximum Satisfiability." In: *Int'l Symposium on Software Testing and Analysis (ISSTA)*. ACM, 2014, pp. 237–247. DOI: 10/gp6khr.
- [72] J. Zhou, S. Silvestro, H. Liu, Y. Cai, and T. Liu. "UNDEAD: Detecting and Preventing Deadlocks in Production Software." In: *Int'l Conference on Automated Software Engineering (ASE)*. 2017, pp. 729–740. DOI: 10/gp6kqx.
- [73] D. Kozen. "Lower Bounds for Natural Proof Systems." In: *Annual Symposium on Foundations of Computer Science (SFCS)*. 1977, pp. 254–266. DOI: 10/dbkc79.
- [74] S. Thompson and G. Brat. "Verification of C++ Flight Software with the MCP Model Checker." In: *IEEE Aerospace Conference*. IEEE, 2008, pp. 1–9. DOI: 10/ds9chc.
- [75] B. Cook, A. Podelski, and A. Rybalchenko. "Proving Thread Termination." In: *Conference on Programming Language Design and Implementation (PLDI)*. ACM, 2007, pp. 320–330. DOI: 10/dp236p.
- [76] Y. Cai and L. Cao. "Fixing Deadlocks via Lock Pre-Acquisitions." In: *Int'l Conference on Software Engineering (ICSE)*. ACM, 2016, pp. 1109–1120. DOI: 10/gp6kq3.
- [77] X. L. Yu, O. Al-Bataineh, D. Lo, and A. Roychoudhury. "Smart Contract Repair." In: *ACM Transactions on Software Engineering and Methodology* 29.4 (2020), pp. 1–32. DOI: 10/gpd4hr.
- [78] K. Ku, T. E. Hart, M. Chechik, and D. Lie. "A Buffer Overflow Benchmark for Software Model Checkers." In: *Int'l Conference on Automated Software Engineering (ASE)*. ACM, 2007, pp. 389–392. DOI: 10/fds27b.
- [79] M. Martinez and M. Monperrus. "ASTOR: A Program Repair Library for Java." In: *Int'l Symposium on Software Testing and Analysis (ISSTA)*. ACM, 2016, pp. 441–444. DOI: 10/gndn55.
- [80] M. P. Gissurarson, L. Applis, A. Panichella, A. van Deursen, and D. Sands. "PROPR: Property-Based Automatic Program Repair." In: *Int'l Conference on Software Engineering (ICSE)*. 2022, pp. 1768–1780. DOI: 10/gqhgs7.
- [81] H. D. T. Nguyen, D. Qi, A. Roychoudhury, and S. Chandra. "SemFix: Program Repair via Semantic Analysis." In: *Int'l Conference on Software Engineering (ICSE)*. 2013, pp. 772–781. DOI: 10/gg82z6.
- [82] J. Xuan, M. Martinez, F. DeMarco, M. Clément, S. L. Marcote, T. Durieux, D. Le Berre, and M. Monperrus. "Nopol: Automatic Repair of Conditional Statement Bugs in Java Programs." In: *IEEE Transactions on Software Engineering* 43.1 (2017), pp. 34–55. DOI: 10/gm5s3h.
- [83] S. Mechtaev, J. Yi, and A. Roychoudhury. "DirectFix: Looking for Simple Program Repairs." In: *Int'l Conference on Software Engineering (ICSE)*. Vol. 1. 2015, pp. 448–458. DOI: 10/gndpcc.
- [84] F. Long and M. Rinard. "Staged Program Repair with Condition Synthesis." In: *Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, 2015, pp. 166–178. DOI: 10/gfvmzm.
- [85] S. Mechtaev, J. Yi, and A. Roychoudhury. "Angelix: Scalable Multiline Program Patch Synthesis via Symbolic Analysis." In: *Int'l Conference on Software Engineering (ICSE)*. 2016, pp. 691–701. DOI: 10/ggsskp.
- [86] X.-B. D. Le, D.-H. Chu, D. Lo, C. Le Goues, and W. Visser. "JFIX: Semantics-Based Repair of Java Programs via Symbolic PathFinder." In: *Int'l Symposium on Software Testing and Analysis (ISSTA)*. ACM, 2017, pp. 376–379. DOI: 10/gp6krg.
- [87] S. A. Asadollah, H. Hansson, D. Sundmark, and S. Eldh. "Towards Classification of Concurrency Bugs Based on Observable Properties." In: *Int'l Workshop on Complex Faults and Failures in Large Software Systems (COUFLESS)*. 2015, pp. 41–47. DOI: 10/f3nh6c.
- [88] A. Podelski and A. Rybalchenko. "ARMC: The Logical Choice for Software Model Checking with Abstraction Refinement." In: *Practical Aspects of Declarative Languages*. Ed. by M. Hanus. Springer, 2007, pp. 245–259. DOI: 10/dmrfmf.
- [89] A. Solar-Lezama, L. Tancau, R. Bodik, S. Seshia, and V. Saraswat. "Combinatorial Sketching for Finite Programs." In: *Int'l Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2006, pp. 404–415. DOI: 10/fdsnnt.