

The Art, Science, and Engineering of Fuzzing: A Survey

Valentin J.M. Manès, HyungSeok Han, Choongwoo Han, Sang Kil Cha, Manuel Egele, Edward J. Schwartz, and Maverick Woo

Abstract—Among the many software vulnerability discovery techniques available today, *fuzzing* has remained highly popular due to its conceptual simplicity, its low barrier to deployment, and its vast amount of empirical evidence in discovering real-world software vulnerabilities. At a high level, fuzzing refers to a process of repeatedly running a program with generated inputs that may be syntactically or semantically malformed. While researchers and practitioners alike have invested a large and diverse effort towards improving fuzzing in recent years, this surge of work has also made it difficult to gain a comprehensive and coherent view of fuzzing. To help preserve and bring coherence to the vast literature of fuzzing, this paper presents a unified, general-purpose model of fuzzing together with a taxonomy of the current fuzzing literature. We methodically explore the design decisions at every stage of our model fuzzer by surveying the related literature and innovations in the art, science, and engineering that make modern-day fuzzers effective.

Index Terms—software security, automated software testing, fuzzing.

1 INTRODUCTION

Ever since its introduction in the early 1990s [152], *fuzzing* has remained one of the most widely-deployed techniques to discover software security vulnerabilities. At a high level, fuzzing refers to a process of repeatedly running a program with generated inputs that may be syntactically or semantically malformed. In practice, attackers routinely deploy fuzzing in scenarios such as exploit generation and penetration testing [21], [109]; several teams in the 2016 DARPA Cyber Grand Challenge (CGC) also employed fuzzing in their cyber reasoning systems [37], [200], [9], [93]. Fueled by these activities, defenders have started to use fuzzing in an attempt to discover vulnerabilities before attackers do. For example, prominent vendors such as Adobe [1], Cisco [2], Google [61], [5], [15], and Microsoft [38], [8] all employ fuzzing as part of their secure development practices. More recently, security auditors [237] and open-source developers [4] have also started to use fuzzing to gauge the security of commodity software packages and provide some suitable forms of assurance to end-users.

The fuzzing community is extremely vibrant. As of this writing, GitHub alone hosts over a thousand public repositories related to fuzzing [86]. And as we will demonstrate, the literature also contains a large number of fuzzers (see

Figure 1 on p. 5) and an increasing number of fuzzing studies appear at major security conferences (e.g. [225], [52], [37], [176], [83], [239]). In addition, the blogosphere is filled with many success stories of fuzzing, some of which also contain what we consider to be gems that warrant a permanent place in the literature.

Unfortunately, this surge of work in fuzzing by researchers and practitioners alike also bears a warning sign of impeded progress. For example, the description of some fuzzers do not go much beyond their source code and manual page. As such, it is easy to lose track of the design decisions and potentially important tweaks in these fuzzers over time. Furthermore, there has been an observable fragmentation in the terminology used by various fuzzers. For example, whereas AFL [231] uses the term “test case minimization” to refer to a technique that reduces the size of a crashing input, the same technique is called “test case reduction” in funfuzz [187]. At the same time, while BFF [49] includes a similar-sounding technique called “crash minimization”, this technique actually seeks to minimize the number of bits that differ between the crashing input and the original seed file and is not related to reducing input size. This makes it difficult, if not impossible, to compare fuzzers using the published evaluation results. We believe such fragmentation makes it difficult to discover and disseminate fuzzing knowledge and this may severely hinder the progress in fuzzing research in the long run.

Due to the above reasons, we believe it is prime time to consolidate and distill the large amount of progress in fuzzing, many of which happened after the three trade-books on the subject were published in 2007–2008 [79], [203], [205].

As we attempt to unify the field, we will start by using §2 to present our fuzzing terminology and a unified model of fuzzing. Staying true to the purpose of this paper, our terminology is chosen to closely reflect the current predominant usages, and our model fuzzer (Algorithm 1, p. 3) is designed to suit a large number of fuzzing tasks

- V. J. M. Manès is with KAIST Cyber Security Research Center, Korea
E-mail: valentin.manes@kaist.ac.kr.
- H. Han and S. K. Cha are with KAIST, Korea
E-mail: hyungseok.han@kaist.ac.kr and sangkilc@kaist.ac.kr.
- C. Han is with Naver Corp., Korea
E-mail: cwahan.tunz@navercorp.com.
- M. Egele is with Boston University
E-mail: megele@bu.edu.
- E. J. Schwartz is with SEI, Carnegie Mellon University
E-mail: edmcman@cmu.edu
- M. Woo is with Carnegie Mellon University
E-mail: pooh@cmu.edu.

Corresponding author: Sang Kil Cha.
Manuscript submitted on April 8, 2019.

as classified in a taxonomy of the current fuzzing literature (Figure 1, p. 5). With this setup, we will then explore every stage of our model fuzzer in §3–§7, and present a detailed overview of major fuzzers in Table 1 (p. 6). At each stage, we will survey the relevant literature to explain the design choices, discuss important trade-offs, and highlight many marvelous engineering efforts that help make modern-day fuzzers effective at their task.

2 SYSTEMIZATION, TAXONOMY, AND TEST PROGRAMS

The term “fuzz” was originally coined by Miller *et al.* in 1990 to refer to a program that “generates a stream of random characters to be consumed by a target program” [152, p. 4]. Since then, the concept of fuzz as well as its action—“fuzzing”—has appeared in a wide variety of contexts, including dynamic symbolic execution [90], [226], grammar-based test case generation [88], [105], [213], permission testing [24], [80], behavioral testing [122], [175], [224], complexity testing [135], [222], kernel testing [216], [196], [186], representation dependence testing [121], function detection [227], robustness evaluation [223], exploit development [111], GUI testing [197], signature generation [72], and penetration testing [81], [156]. To systematize the knowledge from the vast literature of fuzzing, let us first present a terminology of fuzzing extracted from modern uses.

2.1 Fuzzing & Fuzz Testing

Intuitively, fuzzing is the action of running a Program Under Test (PUT) with “fuzz inputs”. Honoring Miller *et al.*, we consider a fuzz input to be an input that the PUT *may not* be expecting, i.e., an input that the PUT may process incorrectly and trigger a behavior that was unintended by the PUT developer. To capture this idea, we define the term *fuzzing* as follows.

Definition 1 (Fuzzing). Fuzzing is the execution of the PUT using input(s) sampled from an input space (the “fuzz input space”) that *protrudes* the expected input space of the PUT.

Three remarks are in order. First, although it may be common to see the fuzz input space to contain the expected input space, this is *not* necessary—it suffices for the former to contain an input *not in* the latter. Second, in practice fuzzing almost surely runs for *many* iterations; thus writing “repeated executions” above would still be largely accurate. Third, the sampling process is *not* necessarily randomized, as we will see in §5.

Fuzz testing is a form of software testing technique that utilizes fuzzing. To differentiate it from others and to honor what we consider to be its most prominent purpose, we deem it to have a specific goal of finding security-related bugs, which include program crashes. In addition, we also define *fuzzer* and *fuzz campaign*, both of which are common terms in fuzz testing:

Definition 2 (Fuzz Testing). Fuzz testing is the use of fuzzing to test if a PUT violates a security policy.

Definition 3 (Fuzzer). A fuzzer is a program that performs fuzz testing on a PUT.

Definition 4 (Fuzz Campaign). A fuzz campaign is a specific execution of a fuzzer on a PUT with a specific security policy.

The goal of running a PUT through a fuzzing campaign is to find bugs [26] that violate the specified security policy. For example, a security policy employed by early fuzzers tested only whether a generated input—the *test case*—crashed the PUT. However, fuzz testing can actually be used to test any security policy observable from an execution, i.e., EM-enforceable [183]. The specific mechanism that decides whether an execution violates the security policy is called the *bug oracle*.

Definition 5 (Bug Oracle). A bug oracle is a program, perhaps as part of a fuzzer, that determines whether a given execution of the PUT violates a specific security policy.

We refer to the algorithm implemented by a fuzzer simply as its “fuzz algorithm”. Almost all fuzz algorithms depend on some parameters beyond (the path to) the PUT. Each concrete setting of the parameters is a *fuzz configuration*:

Definition 6 (Fuzz Configuration). A fuzz configuration of a fuzz algorithm comprises the parameter value(s) that control(s) the fuzz algorithm.

The definition of a fuzz configuration is intended to be broad. Note that the type of values in a fuzz configuration depend on the type of the fuzz algorithm. For example, a fuzz algorithm that sends streams of random bytes to the PUT [152] has a simple configuration space $\{(PUT)\}$. On the other hand, sophisticated fuzzers contain algorithms that accept a set of configurations and evolve the set over time—this includes adding and removing configurations. For example, CERT BFF [49] varies both the mutation ratio and the seed over the course of a campaign, and thus its configuration space is $\{(PUT, s_1, r_1), (PUT, s_2, r_2), \dots\}$. A seed is a (commonly well-structured) input to the PUT, used to generate test cases by modifying it. Fuzzers typically maintain a collection of seeds, and some fuzzers evolve the collection as the fuzz campaign progresses. This collection is called a *seed pool*. Finally, a fuzzer is able to store some data within each configuration. For instance, coverage-guided fuzzers may store the attained coverage in each configuration.

2.2 Paper Selection Criteria

To achieve a well-defined scope, we have chosen to include all publications on fuzzing in the last proceedings of 4 major security conferences and 3 major software engineering conferences from Jan 2008 to February 2019. Alphabetically, the former includes (i) ACM Conference on Computer and Communications Security (CCS), (ii) IEEE Symposium on Security and Privacy (S&P), (iii) Network and Distributed System Security Symposium (NDSS), and (iv) USENIX Security Symposium (USEC); and the latter includes (i) ACM International Symposium on the Foundations of Software Engineering (FSE), (ii) IEEE/ACM International Conference on Automated Software Engineering (ASE), and (iii) International Conference on Software Engineering (ICSE). For writings that appear in other venues or mediums, we include them based on our own judgment on their relevance.

ALGORITHM 1: Fuzz Testing

```

Input:  $\mathcal{C}, t_{\text{limit}}$ 
Output:  $\mathbb{B}$  // a finite set of bugs
1  $\mathbb{B} \leftarrow \emptyset$ 
2  $\mathcal{C} \leftarrow \text{PREPROCESS}(\mathcal{C})$ 
3 while  $t_{\text{elapsed}} < t_{\text{limit}} \wedge \text{CONTINUE}(\mathcal{C})$  do
4    $\text{conf} \leftarrow \text{SCHEDULE}(\mathcal{C}, t_{\text{elapsed}}, t_{\text{limit}})$ 
5    $\text{tcs} \leftarrow \text{INPUTGEN}(\text{conf})$ 
   //  $O_{\text{bug}}$  is embedded in a fuzzer
6    $\mathbb{B}', \text{execinfos} \leftarrow \text{INPUTEVAL}(\text{conf}, \text{tcs}, O_{\text{bug}})$ 
7    $\mathcal{C} \leftarrow \text{CONFUPDATE}(\mathcal{C}, \text{conf}, \text{execinfos})$ 
8    $\mathbb{B} \leftarrow \mathbb{B} \cup \mathbb{B}'$ 
9 return  $\mathbb{B}$ 

```

As mentioned in §2.1, *fuzz testing* only differentiates itself from software testing in that fuzz testing is security related. In theory, focusing on security bugs does not imply a difference in the testing process beyond the selection of a bug oracle. The techniques used often vary in practice, however. When designing a testing tool, access to source code and some knowledge about the PUT are often assumed. Such assumptions often drive the development of testing tools to have different characteristics compared to those of fuzzers, which are more likely to be employed by parties other than the PUT’s developer. Nevertheless, the two fields are still closely related to one another. Therefore, when we are unsure whether to classify a publication as relating to “fuzz testing” and include it in this survey, we follow a simple rule of thumb: we include the publication if the word *fuzz* appears in it.

2.3 Fuzz Testing Algorithm

We present a generic algorithm for fuzz testing, Algorithm 1, which we imagine to have been implemented in a *model fuzzer*. It is general enough to accommodate existing fuzzing techniques, including black-, grey-, and white-box fuzzing as defined in §2.4. Algorithm 1 takes a set of fuzz configurations \mathcal{C} and a timeout t_{limit} as input, and outputs a set of discovered bugs \mathbb{B} . It consists of two parts. The first part is the `PREPROCESS` function, which is executed at the beginning of a fuzz campaign. The second part is a series of five functions inside a loop: `SCHEDULE`, `INPUTGEN`, `INPUTEVAL`, `CONFUPDATE`, and `CONTINUE`. Each execution of this loop is called a *fuzz iteration* and each time `INPUTEVAL` executes the PUT on a test case is called a *fuzz run*. Note that some fuzzers do *not* implement all five functions. For example, to model Radamsa [102], which never updates the set of fuzz configurations, `CONFUPDATE` always returns the current set of configurations unchanged.

PREPROCESS (\mathcal{C}) $\rightarrow \mathcal{C}$

A user supplies `PREPROCESS` with a set of fuzz configurations as input, and it returns a potentially-modified set of fuzz configurations. Depending on the fuzz algorithm, `PREPROCESS` may perform a variety of actions such as inserting instrumentation code to PUTs, or measuring the execution speed of seed files. See §3.

SCHEDULE ($\mathcal{C}, t_{\text{elapsed}}, t_{\text{limit}}$) $\rightarrow \text{conf}$

`SCHEDULE` takes in the current set of fuzz configurations, the current time t_{elapsed} , and a timeout t_{limit} as input, and

selects a fuzz configuration to be used for the current fuzz iteration. See §4.

INPUTGEN (conf) $\rightarrow \text{tcs}$

`INPUTGEN` takes a fuzz configuration as input and returns a set of concrete test cases tcs as output. When generating test cases, `INPUTGEN` uses specific parameter(s) in conf . Some fuzzers use a seed in conf for generating test cases, while others use a model or grammar as a parameter. See §5.

INPUTEVAL ($\text{conf}, \text{tcs}, O_{\text{bug}}$) $\rightarrow \mathbb{B}', \text{execinfos}$

`INPUTEVAL` takes a fuzz configuration conf , a set of test cases tcs , and a bug oracle O_{bug} as input. It executes the PUT on tcs and checks if the executions violate the security policy using the bug oracle O_{bug} . It then outputs the set of bugs found \mathbb{B}' and information about each of the fuzz runs execinfos , which may be used to update the fuzz configurations. We assume O_{bug} is embedded in our model fuzzer. See §6.

CONFUPDATE ($\mathcal{C}, \text{conf}, \text{execinfos}$) $\rightarrow \mathcal{C}$

`CONFUPDATE` takes a set of fuzz configurations \mathcal{C} , the current configuration conf , and the information about each of the fuzz runs execinfos , as input. It may update the set of fuzz configurations \mathcal{C} . For example, many grey-box fuzzers reduce the number of fuzz configurations in \mathcal{C} based on execinfos . See §7.

CONTINUE (\mathcal{C}) $\rightarrow \{\text{True}, \text{False}\}$

`CONTINUE` takes a set of fuzz configurations \mathcal{C} as input and outputs a boolean indicating whether a new fuzz iteration should occur. This function is useful to model white-box fuzzers that can terminate when there are no more paths to discover.

2.4 Taxonomy of Fuzzers

For this paper, we have categorized fuzzers into three groups based on the granularity of semantics a fuzzer observes in each fuzz run. These three groups are called black-, grey-, and white-box fuzzers, which we define below. Note that this classification is different from traditional software testing, where there are only two major categories (black- and white-box testing) [158]. As we will discuss in §2.4.3, grey-box fuzzing is a variant of white-box fuzzing that can only obtain some partial information from each fuzz run.

2.4.1 Black-box Fuzzer

The term “black-box” is commonly used in software testing [158], [32] and fuzzing to denote techniques that do *not* see the internals of the PUT—these techniques can observe only the input/output behavior of the PUT, treating it as a black-box. In software testing, black-box testing is also called IO-driven or data-driven testing [158]. Most traditional fuzzers [13], [103], [49], [6], [50] are in this category. Some modern fuzzers, e.g., funfuzz [187] and Peach [76], also take the structural information about inputs into account to generate more meaningful test cases while maintaining the characteristic of not inspecting the PUT. A similar intuition is used in adaptive random testing [57].

2.4.2 White-box Fuzzer

At the other extreme of the spectrum, white-box fuzzing [90] generates test cases by analyzing the internals of the PUT

and the information gathered when executing the PUT. Thus, white-box fuzzers are able to explore the state space of the PUT systematically. The term *white-box fuzzing* was introduced by Godefroid [87] in 2007 and refers to dynamic symbolic execution (DSE), which is a variant of symbolic execution [39], [126], [108]. In DSE, symbolic and concrete execution operate concurrently, where concrete program states are used to simplify symbolic constraints, e.g., concretizing system calls. DSE is thus often referred to as *concolic testing* (concrete + symbolic) [191], [89]. In addition, white-box fuzzing has also been used to describe fuzzers that employ taint analysis [84]. The overhead of white-box fuzzing is typically much higher than that of black-box fuzzing. This is partly because DSE implementations [90], [46], [25] often employ dynamic instrumentation and SMT solving [155]. While DSE is an active research area [90], [88], [38], [172], [112], many DSEs are *not* white-box fuzzers because they do not aim to find security bugs. As such, this paper does not provide a comprehensive survey on DSEs and we refer the reader to recent survey papers [17], [185] for more information on DSEs for non-security applications.

2.4.3 Grey-box Fuzzer

Some fuzzers [78], [68], [205] take a middle ground approach which is dubbed *grey-box fuzzing*. In general, grey-box fuzzers can obtain *some* information internal to the PUT and/or its executions. Unlike white-box fuzzers, grey-box fuzzers do not reason about the full semantics of the PUT; instead, they may perform lightweight static analysis on the PUT and/or gather dynamic information about its executions, such as code coverage. Grey-box fuzzers rely on approximated, imperfect information in order to gain speed and be able to test more inputs. Although there usually is a consensus between security experts, the distinction between black-, grey- and white-box fuzzing is not always clear. Black-box fuzzers may collect some information about fuzz runs, and white-box fuzzers often use some approximations. When classifying the fuzzers in this survey, particularly in Table 1, we used our best judgement.

An early example of grey-box fuzzer is EFS [68], which uses code coverage gathered from each fuzz run to generate test cases with an evolutionary algorithm. Randoop [166] also used a similar approach, though it did not target security vulnerabilities. Modern fuzzers such as AFL [231] and VUzzer [176] are exemplars in this category.

2.5 Fuzzer Genealogy and Overview

Figure 1 (p. 5) presents our categorization of existing fuzzers in chronological order. Starting from the seminal work by Miller *et al.* [152], we manually chose popular fuzzers that either appeared in a major conference or obtained more than 100 GitHub stars, and showed their relationships as a graph. Black-box fuzzers are in the left half of the figure, and grey- and white-box fuzzers are in the right half. Furthermore, fuzzers are subdivided depending on the type of input the PUT uses: file, network, UI, web, kernel I/O, or threads (in the case of concurrency fuzzers).

Table 1 (p. 6) presents a detailed summary of the techniques used in the most notable fuzzers in Figure 1. We had to omit some of fuzzers in Figure 1 due to space

constraints. Each fuzzer is summarized based on its implementation of the five functions of our model fuzzer, and a miscellaneous section that provides other details on the fuzzer. We describe the properties described by each column below. The first column (feedback gathering granularity) indicates whether the fuzzer is black- (●), white- (○), or grey-box (◐). Two circles appear when a fuzzer has two phases which use different kinds of feedback gathering. For example, SymFuzz [52] runs a white-box analysis as a preprocessing step in order to optimize the performance of a subsequent black-box campaign (●+○), and hybrid fuzzers, such as Driller [200], alternate between white- and grey-box fuzzing (◐+○). The second column shows whether the source code of the fuzzer is publicly available. The third column denotes whether fuzzers need the source code of the PUT to operate. The fourth column points out whether fuzzers support in-memory fuzzing (see §3.1.2). The fifth column is about whether fuzzers can infer models (see §5.1.2). The sixth column shows whether fuzzers perform either static or dynamic analysis in `PREPROCESS`. The seventh column indicates if fuzzers support handling multiple seeds, and perform scheduling. The mutation column specifies if fuzzers perform input mutation to generate test cases. We use ◐ to indicate fuzzers that guide input mutation based on the execution feedback. The model-based column is about whether fuzzers generate test cases based on a model. The constraint-based column shows that fuzzers perform a symbolic analysis to generate test cases. The taint analysis column means that fuzzers leverage taint analysis to guide their test case generation process. The two columns in the `INPUTEVAL` section show whether fuzzers perform crash triage using either stack hashing or code coverage. The first column of the `CONFUPDATE` section indicates if fuzzers evolve the seed pool during `CONFUPDATE`, such as adding new seeds to the pool (see §7.1). The second column of the `CONFUPDATE` section is about whether fuzzers learn an input model in an online fashion. Finally, the third column of the `CONFUPDATE` section shows which fuzzers remove seeds from the seed pool (see §7.2).

3 PREPROCESS

Some fuzzers modify the initial set of fuzz configurations before the first fuzz iteration. Such preprocessing is commonly used to instrument the PUT, to weed out potentially-redundant configurations (i.e., “seed selection” [177]), to trim seeds, and to generate driver applications. As will be detailed in §5.1.1 (p. 9), `PREPROCESS` can also be used to prepare a model for future input generation (`INPUTGEN`).

3.1 Instrumentation

Unlike black-box fuzzers, both grey- and white-box fuzzers can instrument the PUT to gather execution feedback as `INPUTEVAL` performs fuzz runs (see §6), or to fuzz the memory contents at runtime. The amount of collected information defines the color of a fuzzer (i.e., black-, white-, or grey-box). Although there are other ways of acquiring information about the internals of the PUT (e.g., processor traces or system call usage [204], [92]), instrumentation is often the method that collects the most valuable feedback.

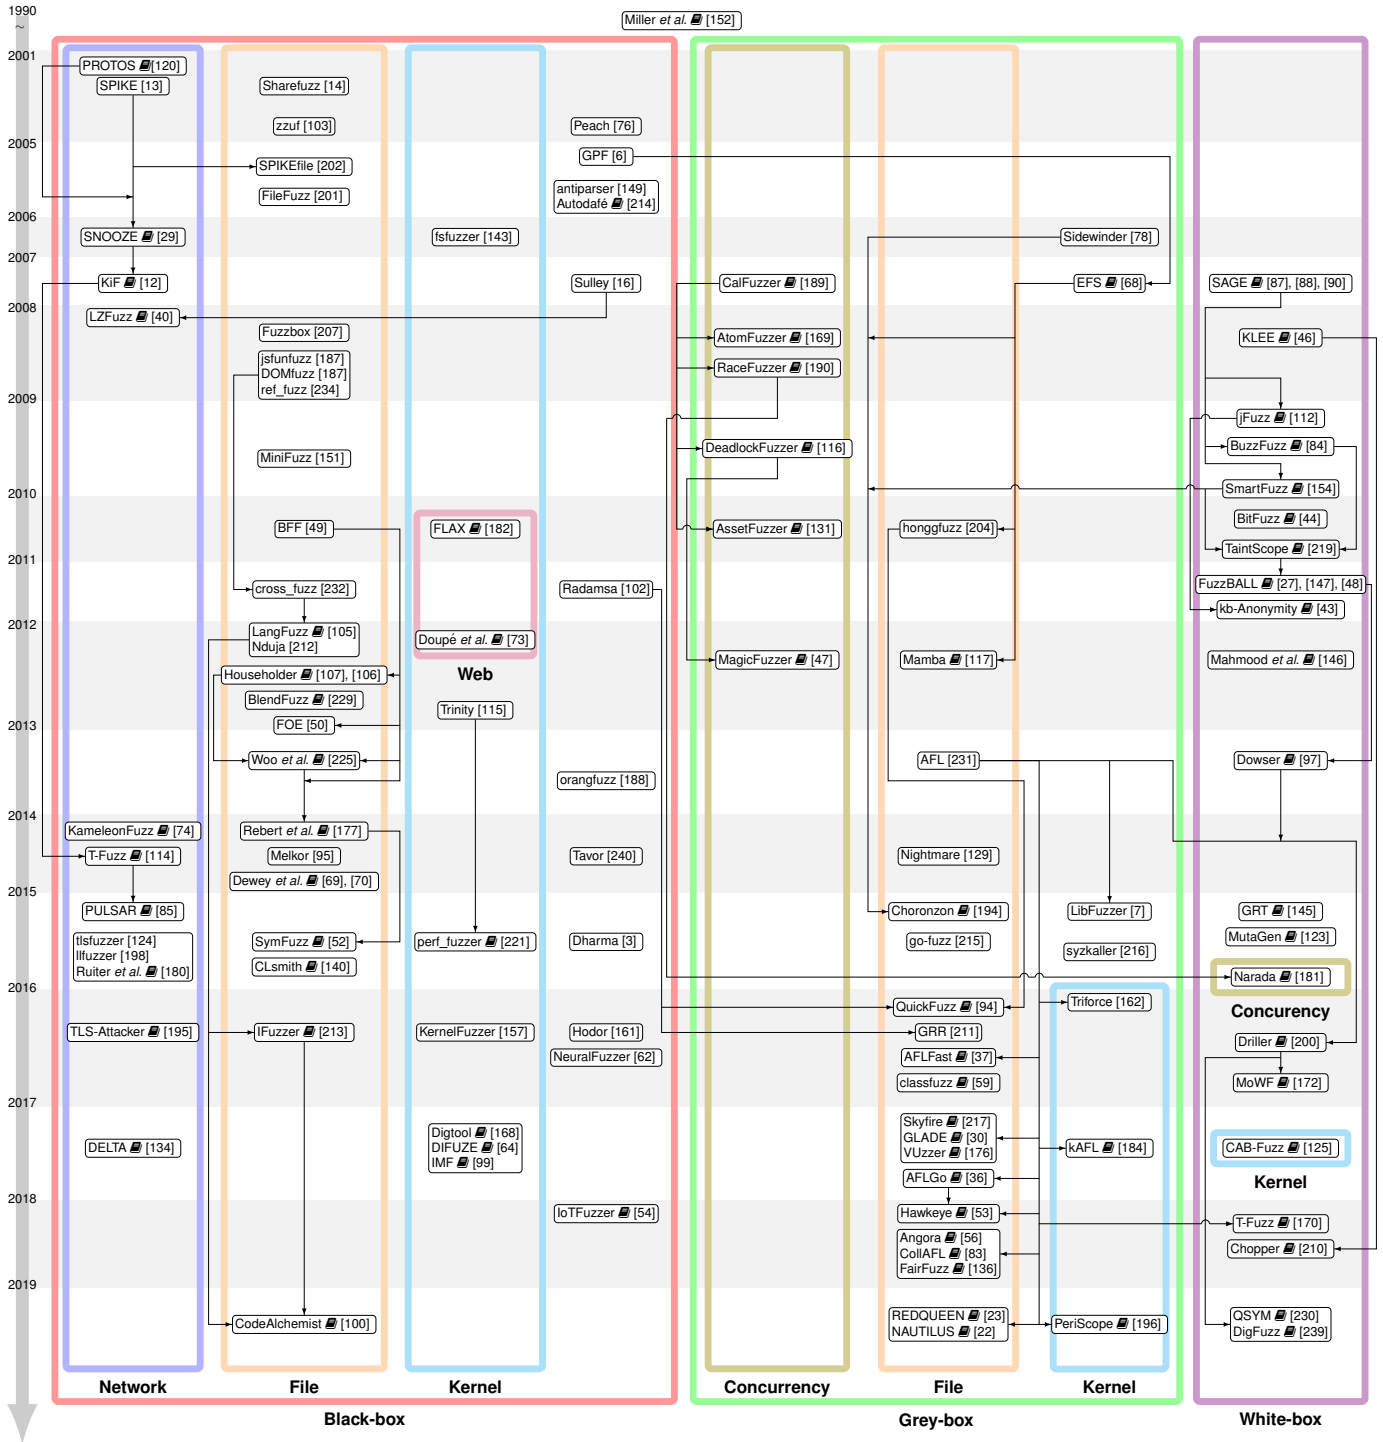



Fig. 1: Genealogy tracing significant fuzzers' lineage back to Miller *et al.*'s seminal work. Each node in the same row represents a set of fuzzers appeared in the same year. A solid arrow from X to Y indicates that Y cites, references, or otherwise uses techniques from X .  denotes that a paper describing the work was published.

TABLE 1: Overview of fuzzers sorted by their instrumentation granularity and their name. ●, ◐, and ○ represent black-, grey-, and white-box, respectively.

Fuzzer	Misc.			PREPROCESS			SCHEDULE	INPUTGEN				INPUTEVAL		CONFUPDATE		
	Feedback Gathering Granularity	Open-Sourced	Source Code Required	Support In-memory Fuzzing	Model Construction	Program Analysis	Seed Scheduling	Mutation	Model-based	Constraint-based	Taint Analysis	Crash Triage: Stack Hash	Crash Triage: Coverage	Evolutionary Seed Pool Update	Model Update	Seed Pool Culling
BFF [49]	●	◐					◐	●								
CodeAlchemist [100]	●	◐			◐			●	◐							
CLSmith [140]	●	◐						●	◐							
DELTA [134]	●	◐						●	◐							
DIFUZE [64]	●	◐	◐		○			●	◐							
Digtool [168]	●	◐						●								
Doupé <i>et al.</i> [73]	●	◐						●	◐						●	
FOE [50]	●	◐					◐	●			◐				●	
GLADE [30]	●	◐			●		◐	●	◐						●	
IMF [99]	●	◐			●			●	◐							
jsfunfuzz [187]	●	◐			●			●	◐		◐					
LangFuzz [105]	●	◐						●	◐							
Miller <i>et al.</i> [152]	●	◐						●	◐							
Peach [76]	●	◐						●	◐		◐					
PULSAR [85]	●	◐			●			●	◐						●	
Radamsa [102]	●	◐						●	◐						●	
Ruiter <i>et al.</i> [180]	●	◐						●	◐						●	
TLS-Attacker [195]	●	◐						●	◐							
zuff [103]	●	◐						●	◐							
FLAX [182]	●+○	◐	◐			◐		●		◐						
IoTFuzzer [54]	●+○	◐			●	◐		●	◐							
SymFuzz [52]	●+○	◐				◐		●	◐		◐					
AFL [231]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
AFLFast [37]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
AFLGo [36]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
AssetFuzzer [131]	◐	◐		◐		◐		●				◐	◐	◐	◐	◐
AtomFuzzer [169]	◐	◐		◐		◐		●				◐	◐	◐	◐	◐
CalFuzzer [189]	◐	◐		◐		◐		●				◐	◐	◐	◐	◐
classfuzz [59]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
CollAFL [83]	◐ [†]	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
DeadlockFuzzer [116]	◐	◐		◐		◐		●				◐	◐	◐	◐	◐
FairFuzz [136]	◐	◐		◐		◐	◐	◐ [†]				◐	◐	◐	◐	◐
go-fuzz [215]	◐	◐		◐		◐	◐	●	◐			◐	◐	◐	◐	◐
Hawkeye [53]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
honggfuzz [204]	◐	◐		◐		◐	◐	●			◐	◐	◐	◐	◐	◐
kAFL [184]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
LibFuzzer [7]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
MagicFuzzer [47]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
Nautilus [22]	◐	◐		◐		◐	◐	◐	◐			◐	◐	◐	◐	◐
RaceFuzzer [190]	◐	◐		◐		◐	◐	●				◐	◐	◐	◐	◐
RedQueen [23]	◐	◐		◐		◐	◐	◐				◐	◐	◐	◐	◐
Steelix [138]	◐ [†]	◐		◐		◐	◐	◐				◐	◐	◐	◐	◐
Syzkaller [216]	◐	◐		◐		◐	◐	●	◐			◐	◐	◐	◐	◐
Angora [56]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
Cyberdyne [92]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
DigFuzz [239]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
Driller [200]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
QSYM [230]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
T-Fuzz [170]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
VUzzer [176]	◐+○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
BitFuzz [44]	○	◐		◐		◐	◐	●		◐		◐	◐	◐	◐	◐
BuzzFuzz [84]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
CAB-Fuzz [125]	○	◐		◐		◐	◐	◐		◐		◐	◐	◐	◐	◐
Chopper [210]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
Dewey <i>et al.</i> [70]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
Dowser [97]	○	◐		◐		◐	◐	◐		◐		◐	◐	◐	◐	◐
GRT [145]	○	◐		◐		◐	◐	◐		◐		◐	◐	◐	◐	◐
KLEE [46]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
MoWF [172]	○	◐		◐		◐	◐	◐		◐		◐	◐	◐	◐	◐
MutaGen [123]	○	◐			●			●		◐		◐	◐	◐	◐	◐
Narada [181]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
SAGE [90]	○	◐	◐			◐	◐	◐		◐		◐	◐	◐	◐	◐
TaintScope [219]	○	◐		◐		◐	◐	◐		◐		◐	◐	◐	◐	◐

[†] The corresponding fuzzer is derived from AFL, and it changed this part of the fuzzing algorithm.

Program instrumentation can be either static or dynamic—the former happens before the PUT runs (`PREPROCESS`), whereas the latter happens while the PUT is running (`INPUTEVAL`). But for the reader’s convenience, we summarize both static and dynamic instrumentation in this section.

Static instrumentation is often performed at compile time on either source code or intermediate code. Since it occurs before runtime, it generally imposes less runtime overhead than dynamic instrumentation. If the PUT relies on libraries, these have to be separately instrumented, commonly by recompiling them with the same instrumentation. Beyond source-based instrumentation, researchers have also developed binary-level static instrumentation (i.e., binary rewriting) tools [238], [132], [77].

Although it has higher overhead than static instrumentation, dynamic instrumentation has the advantage that it can easily instrument dynamically linked libraries, because the instrumentation is performed at runtime. There are several well-known dynamic instrumentation tools such as DynInst [173], DynamoRIO [42], Pin [144], Valgrind [163], and QEMU [33].

A given fuzzer can support more than one type of instrumentation. For example, AFL supports static instrumentation at the source code level with a modified compiler, or dynamic instrumentation at the binary level with the help of QEMU [33]. When using dynamic instrumentation, AFL can either instrument (1) executable code in the PUT itself, which is the default setting, or (2) executable code in the PUT and any external libraries (with the `AFL_INST_LIBS` option). The second option—instrumenting all encountered code—can report coverage information for code in external libraries, and thus provides a more complete picture of coverage. However, this will also encourage AFL to fuzz additional paths in external library functions.

3.1.1 Execution Feedback

Grey-box fuzzers typically take execution feedback as input to evolve test cases. AFL and its descendants compute branch coverage by instrumenting every branch instruction in the PUT. However, they store the branch coverage information in a bit vector, which can cause path collisions. CollAFL [83] recently addressed this issue by introducing a new path-sensitive hash function. Meanwhile, LibFuzzer [7] and Syzkaller [216] use node coverage as their execution feedback. Honggfuzz [204] allows users to choose which execution feedback to use.

3.1.2 In-Memory Fuzzing

When testing a large program, it is sometimes desirable to fuzz *only a portion* of the PUT without re-spawning a process for each fuzz iteration in order to minimize execution overhead. For example, complex (e.g., GUI) applications often require several seconds of processing before they accept input. One approach to fuzzing such programs is to take a snapshot of the PUT after the GUI is initialized. To fuzz a new test case, one can then restore the memory snapshot before writing the new test case directly into memory and executing it. The same intuition applies to fuzzing network applications that involve heavy interaction between client and server. This technique is called in-memory fuzzing [104]. As an example, GRR [211], [92] creates a snapshot before loading any input

bytes. This way, it can skip over unnecessary startup code. AFL also employs a fork server to avoid some of the process startup costs. Although it has the same motivation as in-memory fuzzing, a fork server involves forking off a new process for every fuzz iteration (see §6).

Some fuzzers [7], [231] perform in-memory fuzzing on a function without restoring the PUT’s state after each iteration. We call such a technique as an *in-memory API fuzzing*. For example, AFL has an option called persistent mode [233], which repeatedly performs in-memory API fuzzing in a loop without restarting the process. In this case, AFL ignores potential side effects from the function being called multiple times in the same execution.

Although efficient, in-memory API fuzzing suffers from unsound fuzzing results: bugs (or crashes) found from in-memory fuzzing may *not* be reproducible, because (1) it is not always feasible to construct a valid calling context for the target function, and (2) there can be side-effects that are not captured across multiple function calls. Notice that the soundness of in-memory API fuzzing mainly depends on the entry point function, and finding such a function is a challenging task.

3.1.3 Thread Scheduling

Race condition bugs can be difficult to trigger because they rely on non-deterministic behaviors which may only occur infrequently. However, instrumentation can also be used to trigger different non-deterministic program behaviors by explicitly controlling how threads are scheduled [189], [190], [169], [116], [131], [47], [181]. Existing work has shown that even randomly scheduling threads can be effective at finding race condition bugs [189].

3.2 Seed Selection

Recall from §2 that fuzzers receive a set of fuzz configurations that control the behavior of the fuzzing algorithm. Unfortunately, some parameters of fuzz configurations, such as seeds for mutation-based fuzzers, have large value domains. For example, suppose an analyst fuzzes an MP3 player that accepts MP3 files as input. There is an unbounded number of valid MP3 files, which raises a natural question: which seeds should we use for fuzzing? This problem of decreasing the size of the initial seed pool is known as the *seed selection problem* [177].

There are several approaches and tools that address the seed selection problem [177], [76]. A common approach is to find a minimal set of seeds that maximizes a coverage metric, e.g., node coverage, and this process is called computing a *minset*. For example, suppose the current set of configurations \mathbb{C} consists of two seeds s_1 and s_2 that cover the following addresses of the PUT: $\{s_1 \rightarrow \{10, 20\}, s_2 \rightarrow \{20, 30\}\}$. If we have a third seed $s_3 \rightarrow \{10, 20, 30\}$ that executes roughly as fast as s_1 and s_2 , one could argue it makes sense to fuzz s_3 instead of s_1 and s_2 , since it intuitively tests more program logic for half the execution time cost. This intuition is supported by Miller’s report [153], which showed that a 1% increase in code coverage increased the percentage of bugs found by .92%. As is noted in §7.2, this step can also be part of `CONFUPDATE`, which is useful for fuzzers that can introduce new seeds into the seed pool throughout the campaign.

Fuzzers use a variety of different coverage metrics in practice. For example, AFL’s minset is based on branch coverage with a logarithmic counter on each branch. The rationale behind this decision is to allow branch counts to be considered different only when they differ in orders of magnitude. Honggfuzz [204] computes coverage based on the number of executed instructions, executed branches, and unique basic blocks. This metric allows the fuzzer to add longer executions to the minset, which can help discover denial of service vulnerabilities or performance problems.

3.3 Seed Trimming

Smaller seeds are likely to consume less memory and entail higher throughput. Therefore, some fuzzers attempt to reduce the size of seeds prior to fuzzing them, which is called *seed trimming*. Seed trimming can happen prior to the main fuzzing loop in `PREPROCESS` or as part of `CONFUPDATE`. One notable fuzzer that uses seed trimming is AFL [231], which uses its code coverage instrumentation to iteratively remove a portion of the seed as long as the modified seed achieves the same coverage. Meanwhile, Rebert *et al.* [177] reported that their size minset algorithm, which selects seeds by giving higher priority to smaller seeds in size, results in fewer unique bugs compared to a random seed selection. For the specific case of fuzzing Linux system call handlers, MoonShine [167] extends syzkaller [216] to reduce the size of seeds while preserving the dependencies between calls which are detected using a static analysis.

3.4 Preparing a Driver Application

When it is difficult to directly fuzz the PUT, it makes sense to prepare a driver for fuzzing. This process is largely manual in practice although this is done only once at the beginning of a fuzzing campaign. For example, when our target is a library, we need to prepare for a driver program that calls functions in the library. Similarly, kernel fuzzers may fuzz userland applications to test kernels [125], [165], [31]. IoTFuzzer [54] targets IoT devices by letting the driver communicate with the corresponding smartphone application.

4 SCHEDULING

In fuzzing, scheduling means selecting a fuzz configuration for the next fuzz iteration. As we have explained in §2.1, the content of each configuration depends on the type of the fuzzer. For simple fuzzers, scheduling can be straightforward—for example, zzuf [103] in its default mode allows only one configuration and thus there is simply no decision to make. But for more advanced fuzzers such as BFF [49] and AFLFast [37], a major factor to their success lies in their innovative scheduling algorithms. In this section, we will discuss scheduling algorithms for black- and grey-box fuzzing only; scheduling in white-box fuzzing requires a complex setup unique to symbolic executors and we refer the reader to another source [38].

4.1 The Fuzz Configuration Scheduling (FCS) Problem

The goal of scheduling is to analyze the currently-available information about the configurations and pick one that is

likely to lead to the most favorable outcome, e.g., finding the most number of unique bugs, or maximizing the coverage attained by the set of generated inputs. Fundamentally, every scheduling algorithm confronts the same *exploration vs. exploitation* conflict—time can either be spent on gathering more accurate information on each configuration to inform future decisions (explore), or on fuzzing the configurations that are currently believed to lead to more favorable outcomes (exploit). Woo *et al.* [225] dubbed this inherent conflict the Fuzz Configuration Scheduling (FCS) Problem.

In our model fuzzer (Algorithm 1), the function `SCHEDULE` selects the next configuration based on (i) the current set of fuzz configurations \mathbb{C} , (ii) the current time t_{elapsed} , and (iii) the total time budget t_{limit} . This configuration is then used for the next fuzz iteration. Notice that `SCHEDULE` is only about decision-making. The information on which this decision is based is acquired by `PREPROCESS` and `CONFUPDATE`, which augment \mathbb{C} with this knowledge.

4.2 Black-box FCS Algorithms

In the black-box setting, the only information an FCS algorithm can use is the fuzz outcomes of a configuration—the number of crashes and bugs found with it and the amount of time spent on it so far. Householder and Foote [107] were the first to study how such information can be leveraged in the CERT BFF black-box mutational fuzzer [49]. They postulated that a configuration with a higher observed success rate ($\# \text{bugs} / \# \text{runs}$) should be preferred. Indeed, after replacing the uniform-sampling scheduling algorithm in BFF, they observed 85% more unique crashes over 5 million runs of ffmpeg, demonstrating the potential benefit of more advanced FCS algorithms.

Shortly after, the above idea was improved on multiple fronts by Woo *et al.* [225]. First, they refined the mathematical model of black-box mutational fuzzing from a sequence of Bernoulli trials [107] to the *Weighted Coupon Collector’s Problem with Unknown Weights* (WCCP/UW). Whereas the former assumes each configuration has a fixed eventual success probability and learns it over time, the latter explicitly maintains an upper-bound on this probability as it decays. Second, the WCCP/UW model naturally leads Woo *et al.* to investigate algorithms for *multi-armed bandit* (MAB) problems, which is a popular formalism to cope with the exploration vs. exploitation conflict in decision science [34]. To this end, they were able to design MAB algorithms to accurately exploit configurations that are not known to have decayed yet. Third, they observed that, all else being equal, a configuration that is faster to fuzz allows a fuzzer to either collect more unique bugs with it, or decrease the upperbound on its future success probability more rapidly. This inspired them to normalize the success probability of a configuration by the time that has been spent on it, thus causing a faster configuration to be more preferable. Fourth, they changed the orchestration of fuzz runs in BFF from a fixed number of fuzz iterations per configuration selection (“epochs” in BFF parlance) to a fixed amount of time per selection. With this change, BFF is no longer forced to spend more time in a slow configuration before it can re-select. By combining the above, the evaluation [225] showed a $1.5\times$ increase in the number of unique bugs found using the same amount of time as the existing BFF.

4.3 Grey-box FCS Algorithms

In the grey-box setting, an FCS algorithm can choose to use a richer set of information about each configuration, e.g., the coverage attained when fuzzing a configuration. AFL [231] is the forerunner in this category and it is based on an evolutionary algorithm (EA). Intuitively, an EA maintains a population of configurations, each with some value of “fitness”. An EA selects fit configurations and applies them to genetic transformations such as mutation and recombination to produce offspring, which may later become new configurations. The hypothesis is that these produced configurations are more likely to be fit.

To understand FCS in the context of an EA, we need to define (i) what makes a configuration fit, (ii) how configurations are selected, and (iii) how a selected configuration is used. As a high-level approximation, among the configurations that exercise a control-flow edge, AFL considers the one that contains the fastest and smallest input to be fit (“favorite” in AFL parlance). AFL maintains a queue of configurations, from which it selects the next fit configuration *essentially* as if the queue is circular. Once a configuration is selected, AFL fuzzes it for essentially a constant number of runs. From the perspective of FCS, notice that the preference for fast configurations is common for the black-box setting [225].

Recently, AFLFast by Böhme *et al.* [37] has improved upon AFL in each of the three aspects above. First, AFLFast adds two overriding criteria for an input to become a “favorite”: (i) Among the configurations that exercise a control-flow edge, AFLFast favors the one that has been chosen least. This has the effect of cycling among configurations that exercise this edge, thus increasing exploration. (ii) When there is a tie in (i), AFLFast favors the one that exercises a path that has been exercised least. This has the effect of increasing the exercise of rare paths, which may uncover more unobserved behavior. Second, AFLFast forgoes the round-robin selection in AFL and instead selects the next fit configuration based on a priority. In particular, a fit configuration has a higher priority than another if it has been chosen less often or, when tied, if it exercises a path that has been exercised less often. In the same spirit as the first change, this has the effect of increasing the exploration among fit configurations and the exercising of rare paths. Third, AFLFast fuzzes a selected configuration a variable number of times as determined by a *power schedule*. The FAST power schedule in AFLFast starts with a small “energy” value to ensure initial exploration among configurations and increases exponentially up to a limit to quickly ensure sufficient exploitation. In addition, it also normalizes the energy by the number of generated inputs that exercise the same path, thus promoting explorations of less-frequently fuzzed configurations. The overall effect of these changes is very significant—in a 24-hour evaluation, Böhme *et al.* observed AFLFast discovered 3 bugs that AFL did not, and was on average $7\times$ faster than AFL on 6 other bugs that were discovered by both.

AFLGo [36] extends AFLFast by modifying its priority attribution in order to target specific program locations. Hawkeye [53] further improves directed fuzzing by leveraging a static analysis in its seed scheduling and input generation. FairFuzz [136] guides the campaign to exercise

rare branches by employing a mutation mask for each pair of a seed and a rare branch. QTEP [218] uses static analysis to infer which part of the binary is more ‘faulty’ and prioritize configurations that cover them.

5 INPUT GENERATION

Since the content of a test case directly controls whether or not a bug is triggered, the technique used for *input generation* is naturally one of the most influential design decisions in a fuzzer. Traditionally, fuzzers are categorized into either generation- or mutation-based fuzzers [203]. Generation-based fuzzers produce test cases based on a given model that describes the inputs expected by the PUT. We call such fuzzers *model-based* fuzzers in this paper. On the other hand, mutation-based fuzzers produce test cases by mutating a given *seed* input. Mutation-based fuzzers are generally considered to be *model-less* because seeds are merely example inputs and even in large numbers they do not completely describe the expected input space of the PUT. In this section, we explain and classify the various input generation techniques used by fuzzers based on the underlying test case generation (`INPUTGEN`) mechanism.

5.1 Model-based (Generation-based) Fuzzers

Model-based fuzzers generate test cases based on a given model that describes the inputs or executions that the PUT may accept, such as a grammar precisely characterizing the input format or less precise constraints such as magic values identifying file types.

5.1.1 Predefined Model

Some fuzzers use a model that can be configured by the user. For example, Peach [76], PROTOS [120], and Dharma [3] take in a specification provided by the user. Autodafé [214], Sulley [16], SPIKE [13], and SPIKEfile [202] expose APIs that allow analysts to create their own input models. Tavor [240] also takes in an input specification written in Extended Backus-Naur form (EBNF) and generates test cases conforming to the corresponding grammar. Similarly, network protocol fuzzers such as PROTOS [120], SNOOZE [29], KiF [12], and T-Fuzz [114] also take in a protocol specification from the user. Kernel API fuzzers [115], [216], [157], [162], [221] define an input model in the form of system call templates. These templates commonly specify the number and types of arguments a system call expects as inputs. The idea of using a model in kernel fuzzing originated in Koopman *et al.*’s seminal work [128] where they compared the robustness of OSes with a finite set of manually chosen test cases for system calls. Nautilus [22] employs grammar-based input generation for general-purpose fuzzing, and also uses its grammar for seed trimming (see §3.3).

Other model-based fuzzers target a specific language or grammar, and the model of this language is built in to the fuzzer itself. For example, `cross_fuzz` [232] and `DOMfuzz` [187] generate random Document Object Model (DOM) objects. Likewise, `jsfunfuzz` [187] produces random, but syntactically correct JavaScript code based on its own grammar model. `QuickFuzz` [94] utilizes existing Haskell libraries that describe file formats when generating test cases.

Some network protocol fuzzers such as Frankencerts [41], TLS-Attacker [195], tlsfuzzer [124], and llfuzzer [198] are designed with models of specific network protocols such as TLS and NFC. Dewey *et al.* [69], [70] proposed a way to generate test cases that are not only grammatically correct, but also semantically diverse by leveraging constraint logic programming. LangFuzz [105] produces code fragments by parsing a set of seeds that are given as input. It then randomly combines the fragments, and mutates seeds with the fragments to generate test cases. Since it is provided with a grammar, it always produces syntactically correct code. LangFuzz was applied to JavaScript and PHP. BlendFuzz [229] is based on similar ideas as LangFuzz, but targets XML and regular expression parsers.

5.1.2 Inferred Model

Inferring the model rather than relying on a predefined or user-provided model has recently been gaining traction. Although there is an abundance of published research on the topic of automated input format and protocol reverse engineering [66], [45], [141], [63], [28], only a few fuzzers leverage these techniques. Similar to instrumentation (§3.1), model inference can occur in either `PREPROCESS` or `CONFUPDATE`.

5.1.2.1 Model Inference in `PREPROCESS`: Some fuzzers infer the model as a preprocessing step. Test-Miner [67] searches for the data available in the PUT, such as literals, to predict suitable inputs. Given a set of seeds and a grammar, Skyfire [217] uses a data-driven approach to infer a probabilistic context-sensitive grammar, and then uses it to generate a new set of seeds. Unlike previous works, it focuses on generating semantically valid inputs. IMF [99] learns a kernel API model by analyzing system API logs, and it produces C code that invokes a sequence of API calls using the inferred model. CodeAlchemist [100] breaks JavaScript code into “code bricks”, and computes assembly constraints, which describe when distinct bricks can be assembled or merged together to produce semantically valid test cases. These constraints are computed using both a static analysis and dynamic analysis. Neural [62] and Learn&Fuzz [91] use a neural network-based machine learning technique to learn a model from a given set of test files, and generate test cases from the inferred model. Liu *et al.* [142] proposed a similar approach specific to text inputs.

5.1.2.2 Model Inference in `CONFUPDATE`: Other fuzzers can potentially update their model after each fuzz iteration. PULSAR [85] automatically infers a network protocol model from a set of captured network packets generated from a program. The learned network protocol is then used to fuzz the program. PULSAR internally builds a state machine, and maps which message token is correlated with a state. This information is later used to generate test cases that cover more states in the state machine. Doupé *et al.* [73] propose a way to infer the state machine of a web service by observing the I/O behavior. The inferred model is then used to scan for web vulnerabilities. The work of Ruiter *et al.* [180] is similar, but targets TLS and bases its implementation on LearnLib [174]. GLADE [30] synthesizes a context-free grammar from a set of I/O samples, and fuzzes the PUT using the inferred grammar. Finally, go-fuzz [215] is a grey-box fuzzer, which builds a model for each of the seed it adds

to its seed pool. This model is used to generate new inputs from this seed.

5.1.3 Encoder Model

Fuzzing is often used to test *decoder* programs which parse a certain file format. Many file formats have corresponding *encoder* programs, which can be thought of as an implicit model of the file format. MutaGen [123] is a unique fuzzer that leverages this implicit model contained in an encoder program to generate new test cases. MutaGen leverages mutation to generate test cases, but unlike most mutation-based fuzzers, which mutate an existing *test case* (see §5.2), MutaGen mutates the *encoder program*. Specifically, to produce a new test case MutaGen computes a dynamic program slice of the encoder program and runs it. The underlying idea is that the program slices will slightly change the behavior of the encoder program so that it produces testcases that are slightly malformed.

5.2 Model-less (Mutation-based) Fuzzers

Classic random testing [20], [98] is not efficient in generating test cases that satisfy specific path conditions. Suppose there is a simple C statement: `if (input == 42)`. If `input` is a 32-bit integer, the probability of randomly guessing the right input value is $1/2^{32}$. The situation gets worse when we consider well-structured input such as an MP3 file. It is extremely unlikely that random testing will generate a valid MP3 file as a test case in a reasonable amount of time. As a result, the MP3 player will mostly reject the generated test cases from random testing at the parsing stage before reaching deeper parts of the program.

This problem motivates the use of seed-based input generation as well as white-box input generation (see §5.3). Most model-less fuzzers use a *seed*, which is an input to the PUT, in order to generate test cases by modifying the seed. A seed is typically a well-structured input of a type supported by the PUT: a file, a network packet, or a sequence of UI events. By mutating only a fraction of a valid file, it is often possible to generate a new test case that is mostly valid, but also contains abnormal values to trigger crashes of the PUT. There are a variety of methods used to mutate seeds, and we describe the common ones below.

5.2.1 Bit-Flipping

Bit-flipping is a common technique used by many model-less fuzzers [231], [204], [103], [6], [102]. Some fuzzers simply flip a fixed number of bits, while others determine the number of bits to flip at random. To randomly mutate seeds, some fuzzers employ a user-configurable parameter called the *mutation ratio*, which determines the number of bit positions to flip for a single execution of `INPUTGEN`. Suppose a fuzzer wants to flip K random bits from a given N -bit seed. In this case, a mutation ratio of the fuzzer is K/N .

SymFuzz [52] showed that fuzzing performance is sensitive to the mutation ratio, and that there is not a single ratio that works well for all PUTs. There are several approaches to find a good mutation ratio. BFF [49] and FOE [50] use an exponentially scaled set of mutation ratios for each seed and allocate more iterations to mutation ratios that prove to be statistically effective [107]. SymFuzz [52] leverages a

white-box program analysis to infer a good mutation ratio for each seed.

5.2.2 Arithmetic Mutation

AFL [231] and honggfuzz [204] contain another mutation operation where they consider a selected byte sequence as an integer, and perform simple arithmetic on that value. The computed value is then used to replace the selected byte sequence. The key intuition is to bound the effect of mutation by a small number. For example, AFL selects a 4-byte value from a seed, and treats the value as an integer i . It then replaces the value in the seed with $i \pm r$, where r is a randomly generated small integer. The range of r depends on the fuzzer, and is often user-configurable. In AFL, the default range is: $0 \leq r < 35$.

5.2.3 Block-based Mutation

There are several block-based mutation methodologies, where a block is a sequence of bytes of a seed: (1) insert a randomly generated block into a random position of a seed [231], [7]; (2) delete a randomly selected block from a seed [231], [102], [204], [7]; (3) replace a randomly selected block with a random value [231], [204], [102], [7]; (4) randomly permute the order of a sequence of blocks [102], [7]; (5) resize a seed by appending a random block [204]; and (6) take a random block from a seed to insert/replace a random block of another seed [231], [7].

5.2.4 Dictionary-based Mutation

Some fuzzers use a set of predefined values with potentially significant semantic weight, e.g., 0 or -1 , and format strings for mutation. For example, AFL [231], honggfuzz [204], and LibFuzzer [7] use values such as 0, -1 , and 1 when mutating integers. Radamsa [102] employs Unicode strings and GPF [6] uses formatting characters such as $\%x$ and $\%s$ to mutate strings [55].

5.3 White-box Fuzzers

White-box fuzzers can also be categorized into either model-based or model-less fuzzers. For example, traditional dynamic symbolic execution [90], [112], [27], [147], [200] does not require any model as in mutation-based fuzzers, but some symbolic executors [88], [172], [125] leverage input models such as an input grammar to guide the symbolic executor.

Although many white-box fuzzers including the seminal work by Godefroid *et al.* [90] use dynamic symbolic execution to generate test cases, not all white-box fuzzers are dynamic symbolic executors. Some fuzzers [219], [52], [145], [182] leverage a white-box program analysis to find information about the inputs a PUT accepts in order to use it with black- or grey-box fuzzing. In the rest of this subsection, we briefly summarize the existing white-box fuzzing techniques based on their underlying test case algorithm. Please note that we intentionally omit dynamic symbolic executors such as [89], [191], [60], [46], [209], [51] unless they explicitly call themselves as a fuzzer as mentioned in §2.2.

5.3.1 Dynamic Symbolic Execution

At a high level, classic symbolic execution [126], [39], [108] runs a program with symbolic values as inputs, which represents all possible values. As it executes the PUT, it builds symbolic expressions instead of evaluating concrete values. Whenever it reaches a conditional branch instruction, it conceptually forks two symbolic interpreters, one for the true branch and another for the false branch. For every path, a symbolic interpreter builds up a path formula (or path predicate) for every branch instruction it encountered during an execution. A path formula is satisfiable if there is a concrete input that executes the desired path. One can generate concrete inputs by querying an SMT solver [155] for a solution to a path formula. Dynamic symbolic execution is a variant of traditional symbolic execution, where both symbolic execution and concrete execution operate at the same time. Thus, we often refer to dynamic symbolic execution as concolic (concrete + symbolic) testing. The idea is that concrete execution states can help reduce the complexity of symbolic constraints. An extensive review of the academic literature of dynamic symbolic execution, beyond its application to fuzzing, is out of the scope of this paper. However, a broader treatment of dynamic symbolic execution can be found in other sources [17], [185].

Dynamic symbolic execution is slow compared to grey-box or black-box approaches as it involves instrumenting and analyzing every single instruction of the PUT. To cope with the high cost, a common strategy has been to narrow its usage; for instance, by letting the user to specify uninteresting parts of the code [210], or by alternating between concolic testing and grey-box fuzzing. Driller [200] and Cyberdyne [92] have shown the usefulness of this technique at the DARPA Cyber Grand Challenge. QSYM [230] seeks to improve the integration between grey- and white-box fuzzing by implementing a fast concolic execution engine. DigFuzz [239] optimizes the switch between grey- and white-box fuzzing by first estimating the probability of exercising each path using grey-box fuzzing, and then having its white-box fuzzer focus on the paths that are believed to be most challenging for grey-box fuzzing.

5.3.2 Guided Fuzzing

Some fuzzers leverage static or dynamic program analysis techniques for enhancing the effectiveness of fuzzing. These techniques usually involve fuzzing in two phases: (i) a costly program analysis for obtaining useful information about the PUT, and (ii) test case generation with the guidance from the previous analysis. This is denoted in the sixth column of Table 1 (p. 6). For example, TaintScope [219] uses a fine-grained taint analysis to find “hot bytes”, which are the input bytes that flow into critical system calls or API calls. A similar idea is presented by other security researchers [75], [110]. Dowser [97] performs a static analysis during compilation to find loops that are likely to contain bugs based on a heuristic. Specifically, it looks for loops containing pointer dereferences. It then computes the relationship between input bytes and the candidate loops with a taint analysis. Finally, Dowser runs dynamic symbolic execution while making only the critical bytes to be symbolic hence improving performance. VUzzer [176] and GRT [145] leverage both static and dynamic

analysis techniques to extract control- and data-flow features from the PUT and use them to guide input generation.

Angora [56] and RedQueen [23] decrease the cost of their analysis by first running for each seed with a costly instrumentation and using this information for generating inputs which are run with a lighter instrumentation. Angora improves upon the “hot bytes” idea by using taint analysis to associate each path constraint to corresponding bytes. It then performs a search inspired by gradient descent algorithm to guide its mutations towards solving these constraints. On the other hand, RedQueen tries to detect how inputs are used in the PUT by instrumenting all comparisons and looking for correspondence between their operands and the given input. Once a match is found, it can be used to solve a constraint.

5.3.3 PUT Mutation

One of the practical challenges in fuzzing is bypassing a checksum validation. For example, when a PUT computes a checksum of an input before parsing it, many test cases will be rejected by the PUT. To handle this challenge, TaintScope [219] proposed a checksum-aware fuzzing technique, which identifies a checksum test instruction with a taint analysis, and patches the PUT to bypass the checksum validation. Once they find a program crash, they generate the correct checksum for the input to generate a test case that crashes the unmodified PUT. Caballero *et al.* [44] suggested a technique called stitched dynamic symbolic execution that can generate test cases in the presence of checksums.

T-Fuzz [170] extends this idea to efficiently penetrate all kind of conditional branches with grey-box fuzzing. It first builds a set of Non-Critical Checks (NCC), which are branches that can be transformed without modifying the program logic. When the fuzzing campaign stops discovering new paths, it picks an NCC, transforms it, and then restarts a fuzzing campaign on the modified PUT. Finally, when a crash is found fuzzing a transformed program, T-Fuzz tries to reconstruct it on the original program using symbolic execution.

6 INPUT EVALUATION

After an input is generated, the fuzzer executes the PUT on the input, and decides what to do with the resulting execution. This process is called *input evaluation*. Although the simplicity of executing a PUT is one of the reasons that fuzzing is attractive, there are many optimizations and design decisions related to input evaluation that effect the performance and effectiveness of a fuzzer, and we explore these considerations in this section.

6.1 Bug Oracles

The canonical security policy used with fuzz testing considers every program execution terminated by a fatal signal (such as a segmentation fault) to be a violation. This policy detects many memory vulnerabilities, since a memory vulnerability that overwrites a data or code pointer with an invalid value will usually cause a segmentation fault or abort when it is dereferenced. In addition, this policy is efficient and simple to implement, since operating systems allow such exceptional situations to be trapped by the fuzzer without any instrumentation.

However, the traditional policy of detecting crashes will not detect every memory vulnerability that is triggered. For example, if a stack buffer overflow overwrites a pointer on the stack with a valid memory address, the program might run to completion with an invalid result rather than crashing, and the fuzzer would not detect this. To mitigate this, researchers have proposed a variety of efficient program transformations that detect unsafe or unwanted program behaviors and abort the program. These are often called *sanitizers*.

6.1.1 Memory and Type Safety

Memory safety errors can be separated into two classes: spatial and temporal. Informally, spatial memory errors occur when a pointer is dereferenced outside of the object it was intended to point to. For example, buffer overflows and underflows are canonical examples of spatial memory errors. Temporal memory errors occur when a pointer is accessed after it is no longer valid. For example, a use-after-free vulnerability, in which a pointer is used after the memory it pointed to has been deallocated, is a typical temporal memory error.

Address Sanitizer (ASan) [192] is a fast memory error detector that instruments programs at compile time. ASan can detect spatial and temporal memory errors and has an average slowdown of only 73%, making it an attractive alternative to a basic crash harness. ASan employs a shadow memory that allows each memory address to be quickly checked for validity before it is dereferenced, which allows it to detect many (but not all) unsafe memory accesses, even if they would not crash the original program. MEDS [101] improves on ASan by leveraging the large memory space available on 64-bit platforms to create large chunks of inaccessible memory *redzones* in between allocated objects. These redzones make it more likely that a corrupted pointer will point to invalid memory and cause a crash.

SoftBound/CETS [159], [160] is another memory error detector that instruments programs during compilation. Rather than tracking valid memory addresses like ASan, however, SoftBound/CETS associates bounds and temporal information with each pointer, and can theoretically detect all spatial and temporal memory errors. However, as expected, this completeness comes with a higher average overhead of 116% [160]. CaVer [133], TypeSan [96] and HexType [113] instrument programs during compilation so that they can detect *bad-casting* in C++ type casting. Bad casting occurs when an object is cast to an incompatible type, such as when an object of a base class is cast to a derived type. CaVer has been shown to scale to web browsers, which have historically contained this type of vulnerability, and imposes between 7.6 and 64.6% overhead.

Another class of memory safety protection is *Control Flow Integrity* [10], [11] (CFI), which detects control flow transitions at runtime that are not possible in the original program. CFI can be used to detect test cases that have illegally modified the control flow of a program. A recent project focused on protecting against a subset of CFI violations has landed in the mainstream `gcc` and `clang` compilers [208].

6.1.2 Undefined Behaviors

Languages such as C contain many behaviors that are left undefined by the language specification. The compiler is free to handle these constructs in a variety of ways. In many cases, a programmer may (intentionally or otherwise) write their code so that it is only correct for some compiler implementations. Although this may not seem overly dangerous, many factors can impact how a compiler implements undefined behaviors, including optimization settings, architecture, compiler, and even compiler version. Vulnerabilities and bugs often arise when the compiler’s implementation of an undefined behavior does not match the programmer’s expectation [220].

Memory Sanitizer (MSan) is a tool that instruments programs during compilation to detect undefined behaviors caused by uses of uninitialized memory in C and C++ [199]. Similar to ASan, MSan uses a shadow memory that represents whether each addressable bit is initialized or not. Memory Sanitizer has approximately 150% overhead. Undefined Behavior Sanitizer (UBSan) [71] modifies programs at compile-time to detect undefined behaviors. Unlike other sanitizers which focus on one particular source of undefined behavior, UBSan can detect a wide variety of undefined behaviors, such as using misaligned pointers, division by zero, dereferencing null pointers, and integer overflow. Thread Sanitizer (TSan) [193] is a compile-time modification that detects data races with a trade-off between precision and performance. A data race occurs when two threads concurrently access a shared memory location and at least one of the accesses is a write. Such bugs can cause data corruption and can be extremely difficult to reproduce due to non-determinism.

6.1.3 Input Validation

Testing for input validation vulnerabilities such as XSS (cross site scripting) and SQL injection vulnerabilities is a challenging problem, as it requires understanding the behavior of the very complicated parsers that power web browsers and database engines. KameleonFuzz [74] detects successful XSS attacks by parsing test cases with a real web browser, extracting the Document Object Model tree, and comparing it against manually specified patterns that indicate a successful XSS attack. μ ASQLi [18] uses a similar trick to detect SQL injections. Because it is not possible to reliably detect SQL injections from a web applications response, μ ASQLi uses a database proxy that intercepts communication between the target web application and the database to detect whether an input triggered harmful behavior.

6.1.4 Semantic Difference

Semantic bugs are often discovered using a technique called *differential testing* [148], which compares the behavior of similar (but not identical) programs. Several fuzzers [41], [171], [59] have used differential testing to identify discrepancies between similar programs, which are likely to indicate a bug. Jung *et al.* [118] introduced *black-box differential fuzz testing*, which uses differential testing of multiple inputs on a single program to map mutations from the PUT’s input to its output. These mappings are used to identify information leaks.

6.2 Execution Optimizations

Our model considers individual fuzz iterations to be executed sequentially. While the straightforward implementation of such an approach would simply load the PUT every time a new process is started at the beginning of a fuzz iteration, the repetitive loading processes can be significantly reduced. To this end, modern fuzzers provide functionalities that skip over these repetitive loading processes. For example, AFL [231] provides a fork-server that allows each new fuzz iteration to fork from an already initialized process. Similarly, in-memory fuzzing is another way to optimize the execution speed as discussed in §3.1.2. Regardless of the exact mechanism, the overhead of loading and initializing the PUT is amortized over many iterations. Xu *et al.* [228] further lower the cost of an iteration by designing a new system call that replaces `fork()`.

6.3 Triage

Triage is the process of analyzing and reporting test cases that cause policy violations. Triage can be separated into three steps: deduplication, prioritization, and test case minimization.

6.3.1 Deduplication

Deduplication is the process of pruning any test case from the output set that triggers the same bug as another test case. Ideally, deduplication would return a set of test cases in which each triggers a unique bug.

Deduplication is an important component of most fuzzers for several reasons. As a practical implementation manner, it avoids wasting disk space and other resources by storing duplicate results on the hard drive. As a usability consideration, deduplication makes it easy for users to understand roughly how many different bugs are present, and to be able to analyze an example of each bug. This is useful for a variety of fuzzer users; for example, attackers may want to look only for “home run” vulnerabilities that are likely to lead to reliable exploitation.

There are currently three major deduplication implementations used in practice: stack backtrace hashing, coverage-based deduplication, and semantics-aware deduplication.

6.3.1.1 Stack Backtrace Hashing: Stack backtrace hashing [154] is one of the oldest and most widely used methods for deduplicating crashes, in which an automated tool records a stack backtrace at the time of the crash, and assigns a *stack hash* based on the contents of that backtrace. For example, if the program crashed while executing a line of code in function `foo`, and had the call stack `main → d → c → b → a → foo` (see Fig. 2), then a stack backtrace hashing implementation with $n = 5$ would group that test case with other crashing executions whose backtrace ended with `d → c → b → a → foo`.

Stack hashing implementations vary widely, starting with the number of stack frames that are included in the hash. Some implementations use one [19], three [154], [225], five [82], [49], or do not have any limit [123]. Implementations also differ in the amount of information included from each stack frame. Some implementations will only hash the function’s name or address, but other implementations will hash both the name and the offset or line. Neither option

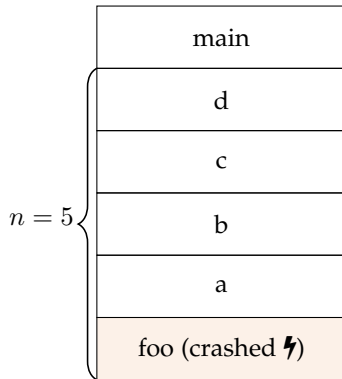


Fig. 2: Stack backtrace hashing example.

works well all the time, so some implementations [150], [82] produce two hashes: a major and minor hash. The major hash is likely to group dissimilar crashes together as it only hashes the function name, whereas the minor hash is more precise since it uses the function name and line number, and also includes an unlimited number of stack frames.

Although stack backtrace hashing is widely used, it is not without its shortcomings. The underlying hypothesis of stack backtrace hashing is that similar crashes are caused by similar bugs, and vice versa, but, to the best of our knowledge, this hypothesis has never been directly tested. There is some reason to doubt its veracity: some crashes do not occur near the code that caused the crash. For example, a vulnerability that causes heap corruption might only crash when an unrelated part of the code attempts to allocate memory, rather than when the heap overflow occurred.

6.3.1.2 Coverage-based Deduplication: AFL [231] is a popular grey-box fuzzer that employs an efficient source-code instrumentation to record the edge coverage of each execution of the PUT, and also measure coarse hit counts for each edge. As a grey-box fuzzer, AFL primarily uses this coverage information to select new seed files. However, it also leads to a fairly unique deduplication scheme as well. As described by its documentation, AFL considers a crash to be unique if either (i) the crash covered a previously unseen edge, or (ii) the crash did *not* cover an edge that was present in all earlier crashes.

6.3.1.3 Semantics-aware Deduplication: RETracer [65] performs crash triage based on the semantics recovered from a reverse data-flow analysis from each crash. Specifically, RETracer checks which pointer caused the crash by analyzing a crash dump (core dump), and recursively identifies which instruction assigned the bad value to it. It eventually finds a function that has the maximum frame level, and “blames” the function. The blamed function can be used to cluster crashes. The authors showed that their technique successfully deduped millions of Internet Explorer bugs into one. In contrast, stack hashing categorized the same crashes into a large number of different groups.

6.3.2 Prioritization and Exploitability

Prioritization, *a.k.a.* the fuzzer taming problem [58], is the process of ranking or grouping violating test cases according to their severity and uniqueness. Fuzzing has traditionally been used to discover memory vulnerabilities, and in this context

prioritization is better known as determining the *exploitability* of a crash. Exploitability informally describes the likelihood of an adversary being able to write a practical exploit for the vulnerability exposed by the test case. Both defenders and attackers are interested in exploitable bugs. Defenders generally fix exploitable bugs before non-exploitable ones, and attackers are interested in exploitable bugs for obvious reasons.

One of the first exploitability ranking systems was Microsoft’s !exploitable [150], which gets its name from the !exploitable WinDbg command name that it provides. !exploitable employs several heuristics paired with a simplified taint analysis [164], [185]. It classifies each crash on the following severity scale: EXPLOITABLE > PROBABLY_EXPLOITABLE > UNKNOWN > NOT_LIKELY_EXPLOITABLE, in which $x > y$ means that x is more severe than y . Although these classifications are not formally defined, !exploitable is informally intended to be conservative and error on the side of reporting something as more exploitable than it is. For example, !exploitable concludes that a crash is EXPLOITABLE if an illegal instruction is executed, based on the assumption that the attacker was able to coerce control flow. On the other hand, a division by zero crash is considered NOT_LIKELY_EXPLOITABLE.

Since !exploitable was introduced, other, similar rule-based heuristics systems have been proposed, including the exploitable plugin for GDB [82] and Apple’s CrashWrangler [19]. However, their correctness has not been systematically studied and evaluated yet.

6.3.3 Test case minimization

Another important part of triage is *test case minimization*. Test case minimization is the process of identifying the portion of a violating test case that is necessary to trigger the violation, and optionally producing a test case that is smaller and simpler than the original, but still causes a violation. Although test case minimization and seed trimming (see 3.3, p. 8) are conceptually similar in that they aim at reducing the size of an input, they are distinct because a minimizer can leverage a bug oracle.

Some fuzzers use their own implementation and algorithms for this. BFF [49] includes a minimization algorithm tailored to fuzzing [106] that attempts to minimize the number of bits that are different from the original seed file. AFL [231] also includes a test case minimizer, which attempts to simplify the test case by opportunistically setting bytes to zero and shortening the length of the test case. Lithium [179] is a general purpose test case minimization tool that minimizes files by attempting to remove “chunks” of adjacent lines or bytes in exponentially descending sizes. Lithium was motivated by the complicated test cases produced by JavaScript fuzzers such as jsfunfuzz [187].

There are also a variety of test case reducers that are not specifically designed for fuzzing, but can nevertheless be used for test cases identified by fuzzing. These include format agnostic techniques such as delta debugging [236], and specialized techniques for specific formats such as C-Reduce [178] for C/C++ files. Although specialized techniques are obviously limited in the types of files they can reduce, they have the advantage that they can be significantly

more efficient than generic techniques, since they have an understanding of the grammar they are trying to simplify.

7 CONFIGURATION UPDATING

The `CONFUPDATE` function plays a critical role in distinguishing the behavior of black-box fuzzers from grey- and white-box fuzzers. As discussed in Algorithm 1, the `CONFUPDATE` function can modify the set of configurations (\mathbb{C}) based on the configuration and execution information collected during the current fuzzing run. In its simplest form, `CONFUPDATE` returns the \mathbb{C} parameter unmodified. Black-box fuzzers do not perform any program introspection beyond evaluating the bug oracle O_{bug} , and so they typically leave \mathbb{C} unmodified because they do not have any information collected that would allow them to modify it¹.

In contrast, grey- and white-box fuzzers are distinguished by their more sophisticated implementations of the `CONFUPDATE` function, which allows them to incorporate new fuzz configurations, or remove old ones that may have been superseded. `CONFUPDATE` enables information collected during one fuzzing iteration to be used by all future fuzzing iterations. For example, white-box fuzzers typically create a new fuzz configuration for every new test case produced, since they produce relatively few test cases compared to black- and grey-box fuzzers.

7.1 Evolutionary Seed Pool Update

An Evolutionary Algorithm (EA) is a heuristic-based approach that involves biological evolution mechanisms such as mutation, recombination, and selection. In the context of fuzzing, an EA maintains a *seed pool* of promising individuals (i.e., seeds) that evolves over the course of a fuzzing campaign as new individuals are discovered. Although the concept of EAs is relatively simple, it forms the basis of many grey-box fuzzers [231], [7], [216]. The process of choosing the seeds to be mutated and the mutation process itself were detailed in §4.3 and §5 respectively.

Arguably, the most important step of an EA is to add a new configuration to the set of configurations \mathbb{C} , which occurs during the `CONFUPDATE` step of fuzzing. Most EA-based fuzzers use node or branch coverage as a fitness function: if a new node or branch is discovered by a test case, it is added to the seed pool. As the number of reachable paths can be orders of magnitude larger than the number of seeds, the seed pool is intended to be a *diverse* subselection of all reachable paths in order to represent the current exploration of the PUT. Also note that seed pools of different size can have the same coverage (as mentioned in §3.2, p. 7).

A common strategy in EA fuzzers is to refine the fitness function so that it can detect more subtle and granular indicators of improvements. For example, AFL [231] refines its fitness function definition by recording the number of times a branch has been taken. STADS [35] presents a statistical framework inspired by ecology to estimate how many new configurations will be discovered if the fuzzing campaign continues. Another common strategy is to measure the fraction of conditions that are met when complex branch

conditions are evaluated. For example, `LAF-INTEL` [130] simply breaks multi-byte comparison into several branches, which allows it to detect when a new seed passes an intermediate byte comparison. `LibFuzzer` [7], `Honggfuzz` [204], `go-fuzz` [215] and `Steelix` [138] instrument all comparisons, and add any test case that makes progress on a comparison to the seed pool. A similar idea was also released as a stand-alone instrumentation module for `clang` [119]. Additionally, `Steelix` [138] checks which input offsets influence comparison instructions. `Angora` [170] improves the fitness criteria of AFL by considering the calling context of each branch taken.

`VUzzer` [176] is an EA-based fuzzer whose fitness function relies on the results of a custom program analysis that determines weights for each basic block. Specifically, `VUzzer` first uses a built-in program analysis to classify basic blocks as either normal or error handling (EH). For a normal block, its weight is inversely proportional to the probability that a random walk on the CFG containing this block visits it according to transition probabilities defined by `VUzzer`. This encourages `VUzzer` to prefer configurations that exercise normal blocks deemed rare by the aforementioned random walk. The weight of EH blocks is *negative*, and its magnitude is the ratio of the number of basic blocks compared to the number of EH blocks exercised by this configuration. These negative weights are used to discourage the execution of error handling (EH) blocks, based on the hypothesis that traversing an EH block signals a lower chance of exercising a vulnerability since bugs often coincide with unhandled errors.

7.2 Maintaining a Minset

With the ability to create new fuzzing configurations comes the risk of creating too many configurations. A common strategy used to mitigate this risk is to maintain a *minset*, or a minimal set of test cases that maximizes a coverage metric. Minsetting is also used during `PREPROCESS`, and is described in more detail in §3.2.

Some fuzzers use a variant of maintaining a minset that is specialized for configuration updates. As one example, rather than completely removing configurations that are not in the minset, which is what `Cyberdyne` [92] does, AFL [231] uses a *culling* procedure to mark minset configurations as being *favorable*. Favorable fuzzing configurations are given a significantly higher chance of being selected for fuzzing by the `SCHEDULE` function. The author of AFL notes that “this provides a reasonable balance between queue cycling speed and test case diversity” [235].

8 RELATED WORK

The literature on fuzzing had an early bloom in 2007–2008, when three trade-books on the subject were published within the two-year period [79], [203], [205]. These books took a more practical approach by presenting the different tools and techniques available at the time and their usages on a variety of targets. We note that `Takanen et al.` [205] already distinguished among black-, white- and grey-box fuzzers, although no formal definitions were given. Most recently, [205] had been revised after a decade. The second edition [206] contained many updates to include modern tools such as AFL [231] and `ClusterFuzz` [61].

1. Some fuzzers add violating test cases to the set of seeds. For example, `BFF` [49] calls this feature crash recycling.

We are aware of two fuzzing surveys that are concurrent to ours [137], [139]. However, the goals of both of these surveys are more focused than ours, which is to provide a comprehensive study on recent developments covering the entire area. In particular, Li *et al.* [137] provided a thorough review of many recent advances in fuzzing, though the authors have also decided to focus on the detail of coverage-based fuzzing and not others. More similar to ours, Liang *et al.* [139] proposed an informal model to describe various fuzzing techniques. However, their model is not flexible enough to encompass some of the fuzzing approaches we cover in this paper, such as model inference (see §5.1.2) and hybrid fuzzing (see §5.3).

Klees *et al.* [127] recently found that there has been no coherent way of evaluating fuzzing techniques, which can hamper our ability to compare the effectiveness of fuzzing techniques. In addition, they provided several useful guidelines for evaluating fuzzing algorithms. We consider their work to be orthogonal to ours as the evaluation of fuzzing algorithms is beyond the scope of this paper.

9 CONCLUDING REMARKS

As we have set forth in §1, our goal for this paper is to distill a comprehensive and coherent view of modern fuzzing literature. To this end, we first present a general-purpose model fuzzer to facilitate our effort to explain the many forms of fuzzing in current use. Then, we illustrate a rich taxonomy of fuzzers using Figure 1 (p. 5) and Table 1 (p. 6). We have explored every stage of our model fuzzer by discussing the design decisions as well as showcasing the many achievements by the community at large.

REFERENCES

- [1] “Binspector: Evolving a security tool,” <https://blogs.adobe.com/security/2015/05/binspector-evolving-a-security-tool.html>.
- [2] “Cisco secure development lifecycle,” <https://www.cisco.com/c/en/us/about/security-center/security-programs/secure-development-lifecycle/sdl-process/validate.html>.
- [3] “dharma,” <https://github.com/MozillaSecurity/dharma>.
- [4] “The fuzzing project,” <https://fuzzing-project.org/software.html>.
- [5] “Google chromium security,” <https://www.chromium.org/Home/chromium-security/bugs>.
- [6] “GPF,” http://www.vdalabs.com/tools/efs_gpf.html.
- [7] “LibFuzzer,” <http://llvm.org/docs/LibFuzzer.html>.
- [8] “Microsoft Security Development Lifecycle, verification phase,” <https://www.microsoft.com/en-us/sdl/process/verification.aspx>.
- [9] “Reddit: Iama mayhem, the hacking machine that won darpa’s cyber grand challenge. ama!” https://www.reddit.com/r/IAMa/comments/4x9yn3/iama_mayhem_the_hacking_machine_that_won_darpas/.
- [10] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti, “Control-flow integrity,” in *Proceedings of the ACM Conference on Computer and Communications Security*, 2005, pp. 340–353.
- [11] —, “Control-flow integrity principles, implementations, and applications,” *ACM Transactions on Information and Systems Security*, vol. 13, no. 1, pp. 4:1–4:40, 2009.
- [12] H. J. Abdelnur, R. State, and O. Festor, “KiF: A stateful sip fuzzer,” in *Proceedings of the International Conference on Principles*, 2007, pp. 47–56.
- [13] D. Aitel, “An introduction to SPIKE, the fuzzer creation kit,” in *Proceedings of the Black Hat USA*, 2001.
- [14] —, “Sharefuzz,” <https://sourceforge.net/projects/sharefuzz/>, 2001.
- [15] M. Aizatsky, K. Serebryany, O. Chang, A. Arya, and M. Whittaker, “Announcing OSS-Fuzz: Continuous fuzzing for open source software,” Google Testing Blog, 2016.
- [16] P. Amini, A. Portnoy, and R. Sears, “sulley,” <https://github.com/OpenRCE/sulley>.
- [17] S. Anand, E. K. Burke, T. Y. Chen, J. Clark, M. B. Cohen, W. Grieskamp, M. Harman, M. J. Harrold, and P. McMinn, “An orchestrated survey of methodologies for automated software test case generation,” *Journal of Systems and Software*, vol. 86, no. 8, pp. 1978–2001, 2013.
- [18] D. Appelt, C. D. Nguyen, L. C. Briand, and N. Alshahwan, “Automated testing for sql injection vulnerabilities: An input mutation approach,” in *Proceedings of the International Symposium on Software Testing and Analysis*, 2014, pp. 259–269.
- [19] Apple Inc., “Accessing crashwrangler to analyze crashes for security implications,” Technical Note TN2334.
- [20] A. Arcuri, M. Z. Iqbal, and L. Briand, “Random testing: Theoretical results and practical implications,” *IEEE Transactions on Software Engineering*, vol. 38, no. 2, pp. 258–277, 2012.
- [21] Ars Technica, “Pwn2own: The perfect antidote to fanboys who say their platform is safe,” <http://arstechnica.com/security/2014/03/pwn2own-the-perfect-antidote-to-fanboys-who-say-their-platform-is-safe/>, 2014.
- [22] C. Aschermann, P. Jauernig, T. Frassetto, A.-R. Sadeghi, T. Holz, and D. Teuchert, “NAUTILUS: Fishing for deep bugs with grammars,” in *Proceedings of the Network and Distributed System Security Symposium*, 2019.
- [23] C. Aschermann, S. Schumilo, T. Blazytko, R. Gawlik, and T. Holz, “REDQUEEN: Fuzzing with input-to-state correspondence,” in *Proceedings of the Network and Distributed System Security Symposium*, 2019.
- [24] K. W. Y. Au, Y. F. Zhou, Z. Huang, and D. Lie, “PScout: Analyzing the android permission specification,” in *Proceedings of the ACM Conference on Computer and Communications Security*, 2012, pp. 217–228.
- [25] T. Avgerinos, A. Rebert, S. K. Cha, and D. Brumley, “Enhancing symbolic execution with Veritesting,” in *Proceedings of the International Conference on Software Engineering*, 2014, pp. 1083–1094.
- [26] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, “Basic concepts and taxonomy of dependable and secure computing,” *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.
- [27] D. Babic, L. Martignoni, S. McCamant, and D. Song, “Statically-directed dynamic automated test generation,” in *Proceedings of the International Symposium on Software Testing and Analysis*, 2011, pp. 12–22.
- [28] G. Bai, J. Lei, G. Meng, S. S. Venkatraman, P. Saxena, J. Sun, Y. Liu, and J. S. Dong, “AUTHSCAN: Automatic extraction of web authentication protocols from implementations,” in *Proceedings of the Network and Distributed System Security Symposium*, 2013.
- [29] G. Banks, M. Cova, V. Felmetzger, K. Almeroth, R. Kemmerer, and G. Vigna, “SNOOZE: Toward a stateful network protocol fuzzer,” in *Proceedings of the International Conference on Information Security*, 2006, pp. 343–358.
- [30] O. Bastani, R. Sharma, A. Aiken, and P. Liang, “Synthesizing program input grammars,” in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2017, pp. 95–110.
- [31] I. Beer, “pwn4fun spring 2014-safari-part ii,” <http://googleprojectzero.blogspot.com/2014/11/pwn4fun-spring-2014-safari-part-ii.html>, 2014.
- [32] B. Beizer, *Black-box Testing: Techniques for Functional Testing of Software and Systems*. John Wiley & Sons, 1995.
- [33] F. Bellard, “QEMU, a fast and portable dynamic translator,” in *Proceedings of the USENIX Annual Technical Conference*, 2005, pp. 41–46.
- [34] D. A. Berry and B. Fristedt, *Bandit Problems: Sequential Allocation of Experiments*. Springer Netherlands, 1985.
- [35] M. Böhme, “STADS: Software testing as species discovery,” *ACM Transactions on Software Engineering and Methodology*, vol. 27, no. 2, pp. 7:1–7:52, 2018.
- [36] M. Böhme, V.-T. Pham, M.-D. Nguyen, and A. Roychoudhury, “Directed greybox fuzzing,” in *Proceedings of the ACM Conference on Computer and Communications Security*, 2017, pp. 2329–2344.

- [37] M. Böhme, V.-T. Pham, and A. Roychoudhury, "Coverage-based greybox fuzzing as markov chain," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2016, pp. 1032–1043.
- [38] E. Bounimova, P. Godefroid, and D. Molnar, "Billions and billions of constraints: Whitebox fuzz testing in production," in *Proceedings of the International Conference on Software Engineering*, 2013, pp. 122–131.
- [39] R. S. Boyer, B. Elspas, and K. N. Levitt, "SELECT—a formal system for testing and debugging programs by symbolic execution," *ACM SIGPLAN Notices*, vol. 10, no. 6, pp. 234–245, 1975.
- [40] S. Bratus, A. Hansen, and A. Shubina, "LZfuzz: a fast compression-based fuzzer for poorly documented protocols," Dartmouth College, Tech. Rep. TR2008-634, 2008.
- [41] C. Brubaker, S. Janapa, B. Ray, S. Khurshid, and V. Shmatikov, "Using frankencerts for automated adversarial testing of certificate validation in SSL/TLS implementations," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2014, pp. 114–129.
- [42] D. L. Bruening, "Efficient, transparent, and comprehensive runtime code manipulation," Ph.D. dissertation, Massachusetts Institute of Technology, 2004.
- [43] A. Budi, D. Lo, L. Jiang, and Lucia, "kb-Anonymity: A model for anonymized behavior-preserving test and debugging data," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2011, pp. 447–457.
- [44] J. Caballero, P. Poosankam, S. McCamant, D. Babić, and D. Song, "Input generation via decompilation and re-stitching: Finding bugs in malware," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2010, pp. 413–425.
- [45] J. Caballero, H. Yin, Z. Liang, and D. Song, "Polyglot: Automatic extraction of protocol message format using dynamic binary analysis," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2007, pp. 317–329.
- [46] C. Cadar, D. Dunbar, and D. Engler, "KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs," in *Proceedings of the USENIX Symposium on Operating System Design and Implementation*, 2008, pp. 209–224.
- [47] Y. Cai and W. Chan, "MagicFuzzer: Scalable deadlock detection for large-scale applications," in *Proceedings of the International Conference on Software Engineering*, 2012, pp. 606–616.
- [48] D. Caselden, A. Bazhanyuk, M. Payer, L. Szekeres, S. McCamant, and D. Song, "Transformation-aware exploit generation using a HI-CFG," University of California, Tech. Rep. UCB/EECS-2013-85, 2013.
- [49] CERT, "Basic Fuzzing Framework," <https://www.cert.org/vulnerability-analysis/tools/bff.cfm>.
- [50] —, "Failure Observation Engine," <https://www.cert.org/vulnerability-analysis/tools/foe.cfm>.
- [51] S. K. Cha, T. Avgerinos, A. Rebert, and D. Brumley, "Unleashing mayhem on binary code," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2012, pp. 380–394.
- [52] S. K. Cha, M. Woo, and D. Brumley, "Program-adaptive mutational fuzzing," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2015, pp. 725–741.
- [53] H. Chen, Y. Xue, Y. Li, B. Chen, X. Xie, X. Wu, and Y. Liu, "Hawkeye: Towards a desired directed grey-box fuzzer," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2018, pp. 2095–2108.
- [54] J. Chen, W. Diao, Q. Zhao, C. Zuo, Z. Lin, X. Wang, W. C. Lau, M. Sun, R. Yang, and K. Zhang, "IoTfuzzer: Discovering memory corruptions in IoT through app-based fuzzing," in *Proceedings of the Network and Distributed System Security Symposium*, 2018.
- [55] K. Chen and D. Wagner, "Large-scale analysis of format string vulnerabilities in debian linux," in *Proceedings of the Workshop on Programming Languages and Analysis for Security*, 2007, pp. 75–84.
- [56] P. Chen and H. Chen, "Angora: Efficient fuzzing by principled search," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2018, pp. 855–869.
- [57] T. Y. Chen, F.-C. Kuo, R. G. Merkel, and T. H. Tse, "Adaptive random testing: The ART of test case diversity," *Journal of Systems and Software*, vol. 83, no. 1, pp. 60–66, 2010.
- [58] Y. Chen, A. Groce, C. Zhang, W.-K. Wong, X. Fern, E. Eide, and J. Regehr, "Taming compiler fuzzers," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2013, pp. 197–208.
- [59] Y. Chen, C. Su, C. Sun, S. Su, and J. Zhao, "Coverage-directed differential testing of jvm implementations," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2016, pp. 85–99.
- [60] V. Chipounov, V. Kuznetsov, and G. Candea, "S2E: A platform for in-vivo multi-path analysis of software systems," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2011, pp. 265–278.
- [61] Chrome Security Team, "Clusterfuzz," <https://code.google.com/p/clusterfuzz/>.
- [62] CIFASIS, "Neural fuzzer," <http://neural-fuzzer.org>.
- [63] P. Comparetti, G. Wondracek, C. Kruegel, and E. Kirda, "Prospex: Protocol specification extraction," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2009, pp. 110–125.
- [64] J. Corina, A. Machiry, C. Salls, Y. Shoshitaishvili, S. Hao, C. Kruegel, and G. Vigna, "DIFUZE: Interface aware fuzzing for kernel drivers," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2017, pp. 2123–2138.
- [65] W. Cui, M. Peinado, S. K. Cha, Y. Fratantonio, and V. P. Kemerlis, "RETracer: Triaging crashes by reverse execution from partial memory dumps," in *Proceedings of the International Conference on Software Engineering*, 2016, pp. 820–831.
- [66] W. Cui, M. Peinado, K. Chen, H. J. Wang, and L. Irun-Briz, "Tupni: Automatic reverse engineering of input formats," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2008, pp. 391–402.
- [67] L. Della Toffola, C. A. Staicu, and M. Pradel, "Saying 'hi' is not enough: Mining inputs for effective test generation," in *Proceedings of the International Conference on Automated Software Engineering*, 2017, pp. 44–49.
- [68] J. D. DeMott, R. J. Enbody, and W. F. Punch, "Revolutionizing the field of grey-box attack surface testing with evolutionary fuzzing," in *Proceedings of the Black Hat USA*, 2007.
- [69] K. Dewey, J. Roesch, and B. Hardekopf, "Language fuzzing using constraint logic programming," in *Proceedings of the International Conference on Automated Software Engineering*, 2014, pp. 725–730.
- [70] —, "Fuzzing the rust typechecker using clp," in *Proceedings of the International Conference on Automated Software Engineering*, 2015, pp. 482–493.
- [71] W. Dietz, P. Li, J. Regehr, and V. Adve, "Understanding integer overflow in C/C++," in *Proceedings of the International Conference on Software Engineering*, 2012, pp. 760–770.
- [72] B. Dolan-Gavitt, A. Srivastava, P. Traynor, and J. Giffin, "Robust signatures for kernel data structures," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2009, pp. 566–577.
- [73] A. Doupé, L. Cavedon, C. Kruegel, and G. Vigna, "Enemy of the State: A state-aware black-box web vulnerability scanner," in *Proceedings of the USENIX Security Symposium*, 2012, pp. 523–538.
- [74] F. Duchene, S. Rawat, J.-L. Richier, and R. Groz, "Kameleonfuzz: Evolutionary fuzzing for black-box XSS detection," in *Proceedings of the ACM Conference on Data and Application Security and Privacy*, 2014, pp. 37–48.
- [75] D. Duran, D. Weston, and M. Miller, "Targeted taint driven fuzzing using software metrics," in *Proceedings of the CanSecWest*, 2011.
- [76] M. Eddington, "Peach fuzzing platform," <http://community.peachfuzzer.com/WhatIsPeach.html>.
- [77] A. Edwards, A. Srivastava, and H. Vo, "Vulcan: Binary transformation in a distributed environment," Microsoft Research, Tech. Rep. MSR-TR-2001-50, 2001.
- [78] S. Embleton, S. Sparks, and R. Cunningham, "'sidewinder': An evolutionary guidance system for malicious input crafting," in *Proceedings of the Black Hat USA*, 2006.
- [79] G. Evron, N. Rathaus, R. Fly, A. Jenik, D. Maynor, C. Miller, and Y. Naveh, *Open Source Fuzzing Tools*. Syngress, 2007.
- [80] A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner, "Android permissions demystified," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2011, pp. 627–638.
- [81] S. Fewer, "A collection of burpsuite intruder payloads, fuzz lists and file uploads," <https://github.com/1N3/IntruderPayloads>.
- [82] J. Foote, "Gdb exploitable," <https://github.com/jfoote/exploitable>.
- [83] S. Gan, C. Zhang, X. Qin, X. Tu, K. Li, Z. Pei, and Z. Chen, "CollAFL: Path sensitive fuzzing," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2018, pp. 660–677.

- [84] V. Ganesh, T. Leek, and M. Rinard, "Taint-based directed whitebox fuzzing," in *Proceedings of the International Conference on Software Engineering*, 2009, pp. 474–484.
- [85] H. Gascon, C. Wressnegger, F. Yamaguchi, D. Arp, and K. Rieck, "PULSAR: Stateful black-box fuzzing of proprietary network protocols," in *Proceedings of the International Conference on Security and Privacy in Communication Systems*, 2015, pp. 330–347.
- [86] GitHub, "Public fuzzers," <https://github.com/search?q=fuzzing&type=Repositories>.
- [87] P. Godefroid, "Random testing for security: Blackbox vs. whitebox fuzzing," in *Proceedings of the International Workshop on Random Testing*, 2007, pp. 1–1.
- [88] P. Godefroid, A. Kiezun, and M. Y. Levin, "Grammar-based whitebox fuzzing," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2008, pp. 206–215.
- [89] P. Godefroid, N. Klarlund, and K. Sen, "DART: Directed automated random testing," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2005, pp. 213–223.
- [90] P. Godefroid, M. Y. Levin, and D. A. Molnar, "Automated whitebox fuzz testing," in *Proceedings of the Network and Distributed System Security Symposium*, 2008, pp. 151–166.
- [91] P. Godefroid, H. Peleg, and R. Singh, "Learn&Fuzz: Machine learning for input fuzzing," in *Proceedings of the International Conference on Automated Software Engineering*, 2017, pp. 50–59.
- [92] P. Goodman and A. Dinaburg, "The past, present, and future of cyberdyne," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2018, pp. 61–69.
- [93] GrammaTech, "Grammatech blogs: The cyber grand challenge," <http://blogs.grammatech.com/the-cyber-grand-challenge>.
- [94] G. Grieco, M. Ceresa, and P. Buiras, "Quickfuzz: An automatic random fuzzer for common file formats," in *Proceedings of the 9th International Symposium on Haskell*, 2016, pp. 13–20.
- [95] A. H. H., "Melkor_elf_fuzzer," https://github.com/IOActive/Melkor_ELF_Fuzzer.
- [96] I. Haller, Y. Jeon, H. Peng, M. Payer, C. Giuffrida, H. Bos, and E. Van Der Kouwe, "TypeSan: Practical type confusion detection," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2016, pp. 517–528.
- [97] I. Haller, A. Slowinska, M. Neugschwandtner, and H. Bos, "Dowsing for overflows: A guided fuzzer to find buffer boundary violations," in *Proceedings of the USENIX Security Symposium*, 2013, pp. 49–64.
- [98] D. Hamlet, "When only random testing will do," in *Proceedings of the International Workshop on Random Testing*, 2006, pp. 1–9.
- [99] H. Han and S. K. Cha, "IMF: Inferred model-based fuzzer," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2017, pp. 2345–2358.
- [100] H. Han, D. Oh, and S. K. Cha, "CodeAlchemist: Semantics-aware code generation to find vulnerabilities in javascript engines," in *Proceedings of the Network and Distributed System Security Symposium*, 2019.
- [101] W. Han, B. Joe, B. Lee, C. Song, and I. Shin, "Enhancing memory error detection for large-scale applications and fuzz testing," in *Proceedings of the Network and Distributed System Security Symposium*, 2018.
- [102] A. Helin, "radamsa," <https://github.com/aoh/radamsa>.
- [103] S. Hocevar, "zzuf," <https://github.com/samhocevar/zzuf>.
- [104] G. Hoglund, "Runtime decompilation," in *Proceedings of the Black Hat USA*, 2003.
- [105] C. Holler, K. Herzig, and A. Zeller, "Fuzzing with code fragments," in *Proceedings of the USENIX Security Symposium*, 2012, pp. 445–458.
- [106] A. D. Householder, "Well there's your problem: Isolating the crash-inducing bits in a fuzzed file," CERT, Tech. Rep. CMU/SEI-2012-TN-018, 2012.
- [107] A. D. Householder and J. M. Foote, "Probability-based parameter selection for black-box fuzz testing," CERT, Tech. Rep. CMU/SEI-2012-TN-019, 2012.
- [108] W. E. Howden, "Methodology for the generation of program test data," *IEEE Transactions on Computers*, vol. C, no. 5, pp. 554–560, 1975.
- [109] InfoSec Institute, "Charlie Miller reveals his process for security research," <http://resources.infosecinstitute.com/how-charlie-miller-does-research/>, 2011.
- [110] V. Iozzo, "0-knowledge fuzzing," in *Proceedings of the Black Hat USA*, 2010.
- [111] S. Jana and V. Shmatikov, "Abusing file processing in malware detectors for fun and profit," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2012, pp. 80–94.
- [112] K. Jayaraman, D. Harvison, V. Ganesh, and A. Kiezun, "jFuzz: A concolic whitebox fuzzer for java," in *Proceedings of the First NASA Forma Methods Symposium*, 2009, pp. 121–125.
- [113] Y. Jeon, P. Biswas, S. Carr, B. Lee, and M. Payer, "HexType: Efficient detection of type confusion errors for c++," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2017, pp. 2373–2387.
- [114] W. Johansson, M. Svensson, U. E. Larson, M. Almgren, and V. Gulisano, "T-Fuzz: Model-based fuzzing for robustness testing of telecommunication protocols," in *Proceedings of the IEEE International Conference on Software Testing, Verification and Validation*, 2014, pp. 323–332.
- [115] D. Jones, "Trinity," <https://github.com/kernelslacker/trinity>.
- [116] P. Joshi, C.-S. Park, K. Sen, and M. Naik, "A randomized dynamic program analysis technique for detecting real deadlocks," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2009, pp. 110–120.
- [117] R. L. S. Jr., "A framework for file format fuzzing with genetic algorithms," Ph.D. dissertation, University of Tennessee, 2012.
- [118] J. Jung, A. Sheth, B. Greenstein, D. Wetherall, G. Maganis, and T. Kohno, "Privacy oracle: A system for finding application leaks with black box differential testing," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2008, pp. 279–288.
- [119] M. Jurczyk, "CompareCoverage," <https://github.com/googleprojectzero/CompareCoverage>.
- [120] R. Kaksonen, M. Laakso, and A. Takanen, "Software security assessment through specification mutations and fault injection," in *Proceedings of the IFIP TC 6/TC 11 International Conference on Communications and Multimedia Security*, 2001, pp. 173–183.
- [121] A. Kanade, R. Alur, S. Rajamani, and G. Ramanalingam, "Representation dependence testing using program inversion," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2010, pp. 277–286.
- [122] A. Kapravelos, C. Grier, N. Chachra, C. Kruegel, G. Vigna, and V. Paxson, "Hulk: Eliciting malicious behavior in browser extensions," in *Proceedings of the USENIX Security Symposium*, 2014, pp. 641–654.
- [123] U. Kargén and N. Shahmehri, "Turning programs against each other: High coverage fuzz-testing using binary-code mutation and dynamic slicing," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2015, pp. 782–792.
- [124] H. Kario, "tlsfuzzer," <https://github.com/tomato42/tlsfuzzer>.
- [125] S. Y. Kim, S. Lee, I. Yun, W. Xu, B. Lee, Y. Yun, and T. Kim, "CAB-Fuzz: Practical concolic testing techniques for COTS operating systems," in *Proceedings of the USENIX Annual Technical Conference*, 2017, pp. 689–701.
- [126] J. C. King, "Symbolic execution and program testing," *Communications of the ACM*, vol. 19, no. 7, pp. 385–394, 1976.
- [127] G. Klees, A. Ruef, B. Cooper, S. Wei, and M. Hicks, "Evaluating fuzz testing," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2018, pp. 2123–2138.
- [128] P. Koopman, J. Sung, C. Dingman, D. Siewiorek, and T. Marz, "Comparing operating systems using robustness benchmarks," in *Proceedings of the Symposium on Reliable Distributed Systems*, 1997, pp. 72–79.
- [129] J. Koret, "Nightmare," <https://github.com/joxeankoret/nightmare>.
- [130] lafintel, "Circumventing fuzzing roadblocks with compiler transformations," <https://lafintel.wordpress.com/2016/08/15/circumventing-fuzzing-roadblocks-with-compiler-transformations/>, 2016.
- [131] Z. Lai, S. Cheung, and W. Chan, "Detecting atomic-set serializability violations in multithreaded programs through active randomized testing," in *Proceedings of the International Conference on Software Engineering*, 2010, pp. 235–244.
- [132] M. Laurenzano, M. M. Tikir, L. Carrington, and A. Snaveley, "PEBIL: Efficient static binary instrumentation for linux," in *Proceedings of the IEEE International Symposium on Performance Analysis of Systems & Software*, 2010, pp. 175–183.

- [133] B. Lee, C. Song, T. Kim, and W. Lee, "Type casting verification: Stopping an emerging attack vector," in *Proceedings of the USENIX Security Symposium*, 2015, pp. 81–96.
- [134] S. Lee, C. Yoon, C. Lee, S. Shin, V. Yegneswaran, and P. Porras, "DELTA: A security assessment framework for software-defined networks," in *Proceedings of the Network and Distributed System Security Symposium*, 2017.
- [135] C. Lemieux, R. Padhye, K. Sen, and D. Song, "PerfFuzz: Automatically generating pathological inputs," in *Proceedings of the International Symposium on Software Testing and Analysis*, 2018, pp. 254–265.
- [136] C. Lemieux and K. Sen, "FairFuzz: A targeted mutation strategy for increasing greybox fuzz testing coverage," in *Proceedings of the International Conference on Automated Software Engineering*, 2018, pp. 475–485.
- [137] J. Li, B. Zhao, and C. Zhang, "Fuzzing: a survey," *Cybersecurity*, vol. 1, no. 1, 2018.
- [138] Y. Li, B. Chen, M. Chandramohan, S.-W. Lin, Y. Liu, and A. Tiu, "Steelix: Program-state based binary fuzzing," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2017, pp. 627–637.
- [139] H. Liang, X. Pei, X. Jia, W. Shen, and J. Zhang, "Fuzzing: State of the art," *IEEE Transactions on Reliability*, vol. 67, no. 3, pp. 1199–1218, 2018.
- [140] C. Lidbury, A. Lascu, N. Chong, and A. F. Donaldson, "Many-core compiler fuzzing," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2015, pp. 65–76.
- [141] Z. Lin and X. Zhang, "Deriving input syntactic structure from execution," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2008, pp. 83–93.
- [142] P. Liu, X. Zhang, M. Pistoia, Y. Zheng, M. Marques, and L. Zeng, "Automatic text input generation for mobile testing," in *Proceedings of the International Conference on Software Engineering*, 2017, pp. 643–653.
- [143] LMH, S. Grubb, I. van Sprundel, E. Sandeen, and J. Wilson, "fsfuzzer," <http://people.redhat.com/sgrubb/files/fsfuzzer-0.7.tar.gz>.
- [144] C.-K. Luk, R. Cohn, R. Muth, H. Patil, A. Klausner, G. Lowney, S. Wallace, V. J. Reddi, and K. Hazelwood, "Pin: Building customized program analysis tools with dynamic instrumentation," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2005, pp. 190–200.
- [145] L. Ma, C. Artho, C. Zhang, H. Sato, J. Gmeiner, and R. Ramler, "GRT: Program-analysis-guided random testing," in *Proceedings of the International Conference on Automated Software Engineering*, 2015, pp. 212–223.
- [146] R. Mahmood, N. Esfahani, T. Kacem, N. Mirzaei, S. Malek, and A. Stavrou, "A whitebox approach for automated security testing of android applications on the cloud," in *Proceedings of the International Workshop on Automation of Software Test*, 2012, pp. 22–28.
- [147] L. Martignoni, S. McCamant, P. Poosankam, D. Song, and P. Maniatis, "Path-exploration lifting: Hi-fi tests for lo-fi emulators," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2012, pp. 337–348.
- [148] W. M. McKeeman, "Differential testing for software," *Digital Technical Journal*, vol. 10, no. 1, pp. 100–107, 1998.
- [149] D. McKinney, "antiparser," <http://antiparser.sourceforge.net/>.
- [150] Microsoft Corporation, "Iexploitable crash analyzer – MSEC debugger extensions," <https://msecdbg.codeplex.com>.
- [151] —, "Minifuzz," <https://msdn.microsoft.com/en-us/biztalk/gg675011>.
- [152] B. P. Miller, L. Fredriksen, and B. So, "An empirical study of the reliability of UNIX utilities," *Communications of the ACM*, vol. 33, no. 12, pp. 32–44, 1990.
- [153] C. Miller, "Fuzz by number: More data about fuzzing than you ever wanted to know," in *Proceedings of the CanSecWest*, 2008.
- [154] D. Molnar, X. C. Li, and D. A. Wagner, "Dynamic test generation to find integer bugs in x86 binary linux programs," in *Proceedings of the USENIX Security Symposium*, 2009, pp. 67–82.
- [155] L. D. Moura and N. Björner, "Satisfiability modulo theories: Introduction and applications," *Communications of the ACM*, vol. 54, no. 9, pp. 69–77, 2011.
- [156] C. Mulliner, N. Golde, and J.-P. Seifert, "SMS of death: from analyzing to attacking mobile phones on a large scale," in *Proceedings of the USENIX Security Symposium*, 2011, pp. 24–24.
- [157] MWR Labs, "KernelFuzzer," <https://github.com/mwrlabs/KernelFuzzer>.
- [158] G. J. Myers, C. Sandler, and T. Badgett, *The Art of Software Testing*. Wiley, 2011.
- [159] S. Nagarakatte, J. Zhao, M. M. K. Martin, and S. Zdancewic, "SoftBound: Highly compatible and complete spatial memory safety for C," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2009, pp. 245–258.
- [160] —, "CETS: Compiler enforced temporal safety for C," in *Proceedings of the International Symposium on Memory Management*, 2010, pp. 31–40.
- [161] NCC Group, "Hodor fuzzer," <https://github.com/nccgroup/hodor>.
- [162] —, "Triforce linux syscall fuzzer," <https://github.com/nccgroup/TriforceLinuxSyscallFuzzer>.
- [163] N. Nethercote and J. Seward, "Valgrind: a framework for heavyweight dynamic binary instrumentation," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2007, pp. 89–100.
- [164] J. Newsome and D. Song, "Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software," in *Proceedings of the Network and Distributed System Security Symposium*, 2005.
- [165] D. Oleksiuk, "Ioctl fuzzer," <https://github.com/Cr4sh/ioctlfuzzer>, 2009.
- [166] C. Pacheco, S. K. Lahiri, M. D. Ernst, and T. Ball, "Feedback-directed random test generation," in *Proceedings of the International Conference on Software Engineering*, 2007, pp. 75–84.
- [167] S. Pailoor, A. Aday, and S. Jana, "MoonShine: Optimizing OS fuzzer seed selection with trace distillation," in *Proceedings of the USENIX Security Symposium*, 2018, pp. 729–743.
- [168] J. Pan, G. Yan, and X. Fan, "Digtool: A virtualization-based framework for detecting kernel vulnerabilities," in *Proceedings of the USENIX Security Symposium*, 2017, pp. 149–165.
- [169] C.-S. Park and K. Sen, "Randomized active atomicity violation detection in concurrent programs," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2008, pp. 135–145.
- [170] H. Peng, Y. Shoshitaishvili, and M. Payer, "T-Fuzz: Fuzzing by program transformation," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2018, pp. 917–930.
- [171] T. Petsios, A. Tang, S. J. Stolfo, A. D. Keromytis, and S. Jana, "NEZHA: Efficient domain-independent differential testing," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2017, pp. 615–632.
- [172] V.-T. Pham, M. Böhme, and A. Roychoudhury, "Model-based whitebox fuzzing for program binaries," in *Proceedings of the International Conference on Automated Software Engineering*, 2016, pp. 543–553.
- [173] P. Project, "DynInst: Putting the performance in high performance computing," <http://www.dyninst.org/>.
- [174] H. Raffelt, B. Steffen, and T. Berg, "LearnLib: A library for automata learning and experimentation," in *Proceedings of the International Workshop on Formal Methods for Industrial Critical Systems*, 2005, pp. 62–71.
- [175] S. Rasthofer, S. Arzt, S. Triller, and M. Pradel, "Making malory behave maliciously: Targeted fuzzing of android execution environments," in *Proceedings of the International Conference on Software Engineering*, 2017, pp. 300–311.
- [176] S. Rawat, V. Jain, A. Kumar, L. Cojocar, C. Giuffrida, and H. Bos, "VUZZer: Application-aware evolutionary fuzzing," in *Proceedings of the Network and Distributed System Security Symposium*, 2017.
- [177] A. Rebert, S. K. Cha, T. Avgerinos, J. Foote, D. Warren, G. Grieco, and D. Brumley, "Optimizing seed selection for fuzzing," in *Proceedings of the USENIX Security Symposium*, 2014, pp. 861–875.
- [178] J. Regehr, Y. Chen, P. Cuoq, E. Eide, C. Ellison, , and X. Yang, "Test-case reduction for C compiler bugs," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2012, pp. 335–346.
- [179] J. Ruderman, "Lithium," <https://github.com/MozillaSecurity/lithium/>.
- [180] J. D. Ruiters and E. Poll, "Protocol state fuzzing of tls implementations," in *Proceedings of the USENIX Security Symposium*, 2015, pp. 193–206.

- [181] M. Samak, M. K. Ramanathan, and S. Jagannathan, "Synthesizing racy tests," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2015, pp. 175–185.
- [182] P. Saxena, S. Hanna, P. Poosankam, and D. Song, "FLAX: Systematic discovery of client-side validation vulnerabilities in rich web applications," in *Proceedings of the Network and Distributed System Security Symposium*, 2010.
- [183] F. B. Schneider, "Enforceable security policies," *ACM Transactions on Information System Security*, vol. 3, no. 1, pp. 30–50, 2000.
- [184] S. Schumilo, C. Aschermann, R. Gawlik, S. Schinzel, and T. Holz, "kAFL: Hardware-assisted feedback fuzzing for os kernels," in *Proceedings of the USENIX Security Symposium*, 2017, pp. 167–182.
- [185] E. J. Schwartz, T. Avgerinos, and D. Brumley, "All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask)," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2010, pp. 317–331.
- [186] M. Schwarz, D. Gruss, M. Lipp, C. Maurice, T. Schuster, A. Fogh, and S. Mangard, "Automated detection, exploitation, and elimination of double-fetch bugs using modern CPU features," in *Proceedings of the ACM Symposium on Information, Computer and Communications Security*, 2018, pp. 587–600.
- [187] M. Security, "funfuzz," <https://github.com/MozillaSecurity/funfuzz>.
- [188] —, "orangfuzz," <https://github.com/MozillaSecurity/orangfuzz>.
- [189] K. Sen, "Effective random testing of concurrent programs," in *Proceedings of the International Conference on Automated Software Engineering*, 2007, pp. 323–332.
- [190] —, "Race directed random testing of concurrent programs," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2008, pp. 11–21.
- [191] K. Sen, D. Marinov, and G. Agha, "CUTE: A concolic unit testing engine for C," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2005, pp. 263–272.
- [192] K. Serebryany, D. Bruening, A. Potapenko, and D. Vyukov, "AddressSanitizer: A fast address sanity checker," in *Proceedings of the USENIX Annual Technical Conference*, 2012, pp. 309–318.
- [193] K. Serebryany and T. Iskhodzhanov, "ThreadSanitizer: data race detection in practice," in *Proceedings of the Workshop on Binary Instrumentation and Applications*, 2009, pp. 62–71.
- [194] Z. Sialveras and N. Naziridis, "Introducing Choronzon: An approach at knowledge-based evolutionary fuzzing," in *Proceedings of the ZeroNights*, 2015.
- [195] J. Somorovsky, "Systematic fuzzing and testing of tls libraries," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2016, pp. 1492–1504.
- [196] D. Song, F. Hetzelt, D. Das, C. Spensky, Y. Na, S. Volckaert, G. Vigna, C. Kruegel, J.-P. Seifert, and M. Franz, "Periscope: An effective probing and fuzzing framework for the hardware-os boundary," in *Proceedings of the Network and Distributed System Security Symposium*, 2019.
- [197] W. Song, X. Qian, and J. Huang, "EHBDRoid: Beyond GUI testing for android applications," in *Proceedings of the International Conference on Automated Software Engineering*, 2017, pp. 27–37.
- [198] C. Spensky and H. Hu, "LI-fuzzer," <https://github.com/mit-ll/LL-Fuzzer>.
- [199] E. Stepanov and K. Serebryany, "MemorySanitizer: fast detector of uninitialized memory use in C++," in *Proceedings of the International Symposium on Code Generation and Optimization*, 2015, pp. 46–55.
- [200] N. Stephens, J. Grosen, C. Salls, A. Dutcher, R. Wang, J. Corbetta, Y. Shoshitaishvili, C. Kruegel, and G. Vigna, "Driller: Augmenting fuzzing through selective symbolic execution," in *Proceedings of the Network and Distributed System Security Symposium*, 2016.
- [201] M. Sutton, "Filefuzz," <http://osdir.com/ml/security.securiteam/2005-09/msg00007.html>.
- [202] M. Sutton and A. Greene, "The art of file format fuzzing," in *Proceedings of the Black Hat Asia*, 2005.
- [203] M. Sutton, A. Greene, and P. Amini, *Fuzzing: Brute Force Vulnerability Discovery*. Addison-Wesley Professional, 2007.
- [204] R. Swiecki and F. Gröbert, "honggfuzz," <https://github.com/google/honggfuzz>.
- [205] A. Takanen, J. D. DeMott, and C. Miller, *Fuzzing for Software Security Testing and Quality Assurance*. Artech House, 2008.
- [206] A. Takanen, J. D. DeMott, C. Miller, and A. Kettunen, *Fuzzing for Software Security Testing and Quality Assurance*, 2nd ed. Artech House, 2018.
- [207] D. Thiel, "Exposing vulnerabilities in media software," in *Proceedings of the Black Hat EU*, 2008.
- [208] C. Tice, T. Roeder, P. Collingbourne, S. Checkoway, U. Erlingsson, L. Lozano, and G. Pike, "Enforcing forward-edge control-flow integrity in gcc & llvm," in *Proceedings of the USENIX Security Symposium*, 2014, pp. 941–955.
- [209] N. Tillmann and J. De Halleux, "Pex—white box test generation for .NET," in *Proceedings of the International Conference on Tests and Proofs*, 2008, pp. 134–153.
- [210] D. Trabish, A. Mattavelli, N. Rinetzky, and C. Cadar, "Chopped symbolic execution," in *Proceedings of the International Conference on Software Engineering*, 2018, pp. 350–360.
- [211] Trail of Bits, "GRR," <https://github.com/trailofbits/grr>.
- [212] R. Valotta, "Taking browsers fuzzing to the next (dom) level," in *Proceedings of the DeepSec*, 2012.
- [213] S. Veggalam, S. Rawat, I. Haller, and H. Bos, "IFuzzer: An evolutionary interpreter fuzzer using genetic programming," in *Proceedings of the European Symposium on Research in Computer Security*, 2016, pp. 581–601.
- [214] M. Vuagnoux, "Autodafé: an act of software torture," in *Proceedings of the Chaos Communication Congress*, 2005, pp. 47–58.
- [215] D. Vyukov, "go-fuzz," <https://github.com/dvyukov/go-fuzz>.
- [216] —, "syzkaller," <https://github.com/google/syzkaller>.
- [217] J. Wang, B. Chen, L. Wei, and Y. Liu, "Skyfire: Data-driven seed generation for fuzzing," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2017, pp. 579–594.
- [218] S. Wang, J. Nam, and L. Tan, "QTEP: Quality-aware test case prioritization," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2017, pp. 523–534.
- [219] T. Wang, T. Wei, G. Gu, and W. Zou, "TaintScope: A checksum-aware directed fuzzing tool for automatic software vulnerability detection," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2010, pp. 497–512.
- [220] X. Wang, N. Zeldovich, M. F. Kaashoek, and A. Solar-Lezama, "Towards optimization-safe systems: Analyzing the impact of undefined behavior," in *Proceedings of the ACM Symposium on Operating System Principles*, 2013, pp. 260–275.
- [221] V. M. Weaver and D. Jones, "perf_fuzzer: Targeted fuzzing of the perf_event_open() system call," UMaine VMW Group, Tech. Rep., 2015.
- [222] J. Wei, J. Chen, Y. Feng, K. Ferles, and I. Dillig, "Singularity: Pattern fuzzing for worst case complexity," in *Proceedings of the International Symposium on Foundations of Software Engineering*, 2018, pp. 213–223.
- [223] S. Winter, C. Sârbu, N. Suri, and B. Murphy, "The impact of fault models on software robustness evaluations," in *Proceedings of the International Conference on Software Engineering*, 2011, pp. 51–60.
- [224] M. Y. Wong and D. Lie, "Intellidroid: A targeted input generator for the dynamic analysis of android malware," in *Proceedings of the Network and Distributed System Security Symposium*, 2016.
- [225] M. Woo, S. K. Cha, S. Gottlieb, and D. Brumley, "Scheduling black-box mutational fuzzing," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2013, pp. 511–522.
- [226] T. Xie, N. Tillmann, J. de Halleux, and W. Schulte, "Fitness-guided path exploration in dynamic symbolic execution," in *Proceedings of the International Conference on Dependable Systems Networks*, 2009, pp. 359–368.
- [227] D. Xu, J. Ming, and D. Wu, "Cryptographic function detection in obfuscated binaries via bit-precise symbolic loop mapping," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2017, pp. 921–937.
- [228] W. Xu, S. Kashyap, C. Min, and T. Kim, "Designing new operating primitives to improve fuzzing performance," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2017, pp. 2313–2328.
- [229] D. Yang, Y. Zhang, and Q. Liu, "Blendfuzz: A model-based framework for fuzz testing programs with grammatical inputs," in *Proceedings of the ACM Conference on Programming Language Design and Implementation*, 2012, pp. 1070–1076.
- [230] I. Yun, S. Lee, M. Xu, Y. Jang, and T. Kim, "QSYM: A practical concolic execution engine tailored for hybrid fuzzing," in *Proceedings of the USENIX Security Symposium*, 2018, pp. 745–762.

- [231] M. Zalewski, "American Fuzzy Lop," <http://lcamtuf.coredump.cx/afl/>.
- [232] —, "Crossfuzz," <https://lcamtuf.blogspot.com/2011/01/announcing-crossfuzz-potential-0-day-in.html>.
- [233] —, "New in AFL: persistent mode," <https://lcamtuf.blogspot.com/2015/06/new-in-afl-persistent-mode.html>.
- [234] —, "ref_fuzz," <http://lcamtuf.blogspot.com/2010/06/announcing-refuzz-2yo-fuzzer.html>.
- [235] —, "Technical "whitepaper" for afl-fuzz," http://lcamtuf.coredump.cx/afl/technical_details.txt.
- [236] A. Zeller and R. Hildebrandt, "Simplifying and isolating failure-inducing input," *IEEE Transactions on Software Engineering*, vol. 28, no. 2, pp. 183–200, 2002.
- [237] K. Zetter, "A famed hacker is grading thousands of programs—and may revolutionize software in the process," <https://goo.gl/LRwaVl>.
- [238] M. Zhang, R. Qiao, N. Hasabnis, and R. Sekar, "A platform for secure static binary instrumentation," in *Proceedings of the International Conference on Virtual Execution Environments*, 2014, pp. 129–140.
- [239] L. Zhao, Y. Duan, H. Yin, and J. Xuan, "Send hardest problems my way: Probabilistic path prioritization for hybrid fuzzing," in *Proceedings of the Network and Distributed System Security Symposium*, 2019.
- [240] M. Zimmermann, "Tavor," <https://github.com/zimmski/tavor>.