



No-Fuzz: Efficient Anti-fuzzing Techniques

Zhengxiang Zhou^(✉), Cong Wang, and Qingchuan Zhao

City University of Hong Kong, Hong Kong, China
zzzhou4-c@my.cityu.edu.hk, {congwang, cs.qczhao}@cityu.edu.hk

Abstract. Fuzzing is an automated software testing technique that has achieved great success in recent years. While this technique allows developers to uncover vulnerabilities avoiding consequent issues (e.g., financial loss), it can also be leveraged by attackers to find zero-day vulnerabilities. To mitigate, anti-fuzzing techniques were proposed to impede the fuzzing process by slowing down its rate, misinforming the feedback, and complicating the data flow. Unfortunately, the state-of-the-art of anti-fuzzing entirely focuses on enhancing its defensive capability but underestimates the nontrivial performance overhead and overlooks the requirement of extra manual efforts. In this paper, to advance the state-of-the-art, we propose an efficient and automatic anti-fuzzing technique and implement a prototype, called No-Fuzz. Comparing to prior works, our evaluations illustrate that No-Fuzz introduces less performance overhead, i.e., less than 15% of the storage cost for one fake block. In addition, in respect of the binary-only fuzzing, No-Fuzz can precisely determine the corresponding running environments and eliminate unnecessary storage overheads with high effectiveness. Specifically, it reduces 95% of the total storage cost compared with the prior works for the same number of branch reductions. Moreover, our study sheds light on approaches to improve the practicality of anti-fuzzing techniques.

Keywords: Anti-fuzzing · Software testing · Fuzzing

1 Introduction

Fuzzing was first introduced as a software testing technique in 1990 [25]. Typically, a fuzzer would persistently feed the target program with randomly generated inputs and observe the abnormalities of the program (e.g., segmentation faults) to identify program bugs. Recently, fuzzers have been well-developed - researchers integrate fuzzers with techniques like program instrumentation [3, 12, 28, 29, 31, 36, 38], program analysis techniques [30, 33, 35] for the efficiency in bug-finding. Besides, researchers also modify the fuzzing mechanisms of some classic works [3, 36], to meticulously reallocate the fuzzing resources on some specific tasks (e.g., directly fuzzing a particular code area [7, 8, 10, 23]). In general, fuzzers have achieved great success with plenty of bugs uncovered [14, 15, 27, 31].

However, exposing bugs in the program is a double-edged sword. Developers can find and fix bugs before they spread on the internet. Meanwhile, attackers

can also leverage fuzzers to explore zero-day vulnerabilities, which might cause financial loss to the companies. Although the adversaries can manually analyze the commercial software, recent studies [17,32] have shown that attackers lean more towards automated tools, like fuzzers, to find vulnerabilities than manual analysis. In the face of the worse situations of bug finding, anti-fuzzing techniques were proposed to hinder malicious use of fuzzers (ANTIFUZZ [16] and FUZZIFICATION [21]). The purpose of anti-fuzzing is to maintain the advantageous position of developers over adversaries on bug-finding. These techniques introduce penalties on fuzzers to disturb the fuzzing heuristics or slow down the fuzzing rate. The source codes of the protected program will be compiled into two versions - one is protected with anti-fuzzing codes, and the other is unmodified. Developers keep the original version, and they can thoroughly test the program. Adversaries only retrieve the protected version, and the anti-fuzzing codes inside the program can severely hinder the use of fuzzers. Consequently, developers are expected to uncover far more bugs than the adversaries and fix them to avoid the possible loss from zero-day vulnerabilities.

Anti-fuzzing techniques are promising, but the prototypes in prior works should be improved to more fine-grained application scope. Intuitively, the storage overhead should be taken into consideration for the practical adoption of anti-fuzzing techniques. In prior works, fake blocks are artificially inserted into the program to saturate the bitmap of fuzzers, while this technique can enlarge the size of the program by even several times the original program. Developers would be unwilling to burden such storage costs only for anti-fuzzing. Instead, they can resort to lighter obfuscation tools whenever applicable, even though these tools may not provide sufficient protection against fuzzers. Another factor that matters is the automation of the tools, i.e., existing prototypes are inconvenient to use. On the one hand, the developers have to locate some code areas manually; on the other hand, the prototypes have some dependencies with third-party tools/libraries that may be incompatible with the OS of users. We believe these factors challenge the future design and implementation of anti-fuzzing techniques. More specifically, the anti-fuzzing techniques should ideally hold the following two properties:

P1) *Both storage and performance overheads should be minimized;*

P2) *The implementation should support automation which has no modification to the development procedures of the program.*

Based on these thoughts, we propose our solution to anti-fuzzing techniques. The solution involves two categories of techniques - the passive detection methods and active disturbance methods. The passive detection methods precisely check whether the protected program is being fuzzed and launch mitigation strategies once fuzzers are detected. In our design, we integrate instrumentation checking and execution frequency checking to achieve low overhead anti-fuzzing techniques. The active disturbance methods impede fuzzers by attacking the basic fuzzing assumptions and prevent the fuzzers from working normally. We optimize the defective fake blocks. In our design, the fake blocks achieve the minimum storage overhead, which is less than 15% of that in prior works.

We implement these techniques as a fully automated tool, i.e., No-Fuzz. The anti-fuzzing techniques are directly added to the source codes of programs, i.e., no modification is needed to the compilation procedures (e.g., header files, linked libraries, compilation commands). Notably, No-Fuzz is also compatible with other anti-fuzzing techniques in prior works that are not mentioned.

In the evaluation, we measure the effectiveness of our techniques in reducing branch coverage with real-world software from Binutils and two popular benchmarks (Google FTS [1] and Magma [18]). Moreover, we show that our techniques can hinder bug findings for the LAVA-M [13] dataset. To confirm our optimizations to prior works, we have also compared No-Fuzz with the corresponding techniques in ANTIFUZZ [16] and FUZZIFICATION [21]. The results show that our design introduces less overhead and mitigates the negative effects of anti-fuzzing techniques on regular users. Specifically, we reduce about 95% of the storage cost compared with the prior works for the same number of branch reductions. Furthermore, we tackle the obstacle that there is no suitable metric to compare different anti-fuzzing techniques currently. Existing works measure the anti-fuzzing effects and the overhead separately. However, the performance and overhead are orthogonal - both of them vary according to different configurations; it is unfair to compare the performance of different works with unequal overhead directly. Therefore, in addition to measuring the performance and overhead separately, we propose a new metric - “anti-fuzzing efficacy” linking the two metrics to measure the increased defensive capability per unit overhead.

In short, this paper makes the following contributions: 1) throws light on the facts of existing anti-fuzzing prototypes and summarises the properties of the ideal anti-fuzzing techniques; 2) designs and implements automated anti-fuzzing prototype No-Fuzz which can detect and disturb run-time fuzzing mechanisms; 3) evaluate No-Fuzz and some of the prior anti-fuzzing techniques on common benchmarks, showing No-Fuzz’s negligible overhead to the protected binary and its effectiveness at impeding binary-only fuzzing. The source codes of all implemented tools are available at <https://github.com/CongGroup/No-Fuzz>.

2 Technical Background of Anti-fuzzing

The purpose of anti-fuzzing techniques is to combat fuzzers to reduce the number of bugs reported on protected binaries. Existing techniques can be majorly summarized as anti-fast-execution, anti-feedback, and anti-hybrid techniques based on the affected fuzzing mechanisms. We will briefly introduce them in the following parts of this section.

Anti-fast-execution: Introduce Latency to Binary. One of the fuzzers’ assumptions is that more trials with different inputs are expected to explore more paths in the binary. Fuzzers are usually designed with accelerating techniques feeding thousands of seeds per second to the program under test (PUT) [37]. Anti-fast-execution techniques insert latency into the binary to prevent fast-execution. However, the challenge is that the latency can also affect regular

users. ANTIFUZZ [16] inserts delay functions in the error handling codes manually; FUZZIFICATION [21] inserts the latency functions in cold blocks. Both techniques are trying to delay the areas that regular users rarely reach, but fuzzers are easy to fall into.

Anti-feedback: Disturb the Feedback Information. Modern fuzzers majorly rely on two feedbacks to decide fuzzing heuristics - coverage-feedback and error signals. The coverage information is stored in a bitmap of limited size, and fuzzers are expected to make decisions on seeds and mutations to maximize the coverage. The error signals inform fuzzers to save and report the seeds triggering bugs, which is also the ultimate goal of using fuzzers.

Anti-feedback techniques insert fake blocks into the protected binary to disturb the coverage feedback. These blocks contain codes irrelevant to the program logic but are recorded as valid blocks in the bitmap. If most space of the bitmap has been taken up by these fake blocks, fuzzers will be unable to update new coverage. ANTIFUZZ hijacks the control flows to randomly generated fake functions in the protected program. FUZZIFICATION inserts a fixed number of constraints and functions into the binary, and builds ROP chains as fake paths on the assembly code snippets.

For the error signals, ChaffBugs [19] suggests inserting non-exploitable bugs into the binary to confuse the segmentation faults reported to fuzzers. ANTIFUZZ proposed an approach to hinder the crash discovery by installing a signal handler. The handler hides signals from fuzzers with elegant exits, and thus fuzzers are unaware of crashes.

Anti-hybrid: Impede Program Analysis. Hybrid fuzzers [11, 20, 30, 35] generally rely on taint analysis and symbolic execution to accelerate fuzzing. Anti-hybrid techniques embed complex data flows in protected binary to hinder both of the techniques. The idea is based on the fact that program analysis techniques have difficulty in dealing with complex data flows due to the limited CPU resources. ANTIFUZZ encrypts and decrypts the inputs and transforms variables in critical comparisons to their hash values. Similarly, FUZZIFICATION adds extra copy operations to the operand string to complicate the data flows and mislead taint analysis engines to a wrong tag map.

3 No-Fuzz Design

No-Fuzz includes passive detection methods and the optimized active technique (fake blocks) of prior works. Passive detection methods detect whether the protected binary is being fuzzed in binary-only-fuzzing (BOF) mode. Once the fuzzers are found, the protection will trigger fuzzing mitigation mechanisms (e.g., introducing latency). As aforementioned, the active methods in prior works are not practical due to the storage overhead. We optimize the fake blocks and design the landing space exploiting the block-identification mechanism of binary-only instrumentation. It reduces the storage overhead of a fake block to only one byte.

3.1 Passive Detection Methods

A fundamental requirement for anti-fuzzing is to avoid the negative effects of the inserted anti-fuzzing techniques on regular users. Ideally, if the protected program itself can accurately perceive the fuzzers, we can impose severe penalties on adversaries covertly. Based on this thinking, we introduce passive detection methods to identify the running environments of protected binaries. Once fuzzers are detected, we carry out mitigation mechanisms such as delaying the execution and aborting the program to prevent fuzzing.

Detect Binary-Only Instrumentation. In the scenario of using anti-fuzzing techniques, adversaries are not able to retrieve the source codes of the protected binary - they rely on the binary-only mode of fuzzers. The key is that no matter what techniques they are using, they have to collect coverage information of the target program. Techniques such as dynamic instrumentation, hardware assistance, and binary rewriting are the most used for this observation, but all of them cause significant latency (a timing gap) to the PUT. We can detect the timing gap between the native execution and the execution with coverage-collecting codes to determine the running environment.

Timing-related techniques are common in the scope of malware detection [5, 22, 24]. We learn from existing works and design the detection on binary instrumentations. In the native execution environment (real CPU), the control flow directly falls into the block after a branch-taken instruction. On the contrary, with BOF, the instrumented program executes the additional instructions collecting coverage at the beginning of a block. We detect BOF by checking the edge instructions count (EIC). EIC is the estimated number of instructions executed when entering a function or a block (instructions of an edge). According to the experiments, the EIC for BOF can be about ten times larger than that in the native execution. If we detect the timing gap in a protected program, we acknowledge the existence of BOF and carry out mitigation techniques.

Mitigation - Introduce Latency. Due to the performance variation of the CPU, there will be false positives in the detection. A few portions of executions can have a relatively large timing gap, even if they are in the native environment. This usually happens when the CPU is conducting context switching, and extra overhead is counted as a part of the timing gap. According to our experiments, 0.03% of the executions are false positives in a stable environment, and the false-positive rate will be 0.1% in a busy environment where too many parallel tasks are executed simultaneously.

As a consequence, the mitigation mechanisms to fuzzing should be moderate in case the false positives affect regular users. We add a one-second latency to the program by triggering an IO blocking if the BOF instrumentation is detected. Although one second is insufficiently long for a fuzzer, the general effectiveness of the penalty can be guaranteed if we insert more than one detection function into the protected program.

Examining Execution Frequency. The nature of fuzzing is that the PUT will be re-executed a large number of times in a short period. This can be leveraged to detect whether the program is being fuzzed. We come up with the idea from “many a little makes a mickle”. If the PUT leaves some vestiges every fuzzing round, the vestiges will accumulate during the fast re-executions. After a while, they grow large enough and inform the PUT that it is being fuzzed. In our design, the protected program creates a temporary file every time it runs. If more than 60 files are created in a minute (the threshold is according to configurations), the program will be alerted to the existence of a fuzzer. However, the challenge is that the program only creates the files, but it is difficult to manage them. It can be time-consuming to traverse the temporary files and check which are created by the protected program. Besides, not deleting them can mess up the file systems for regular users. To cope with this program, we find that the daemon process is suitable for the management of temporary files.

A daemon process is a process that runs as a background process. It detaches from the parent process and keeps running after the termination of its parent. We designed the daemon process to patrol the temporary files - to prevent the temporary files from being unintentionally deleted and delete them after the patrolling. The patrolling daemon process detaches itself from the protected program during the execution. A temporary file with an ID to indicate its order will be created by the daemon process. These files are created in ascending order, and the largest order is the detection threshold for fuzzers. It then locks the file for a period which we call “patrolling time”. After the patrolling, it checks whether the file it locks is correct and deletes the file it creates. The protected program seeks temporary files with the threshold order for every execution. Once the file is found, it means the program has been executed more than the threshold number of times in the patrolling time, and the BOF is likely to exist; so we can apply the mitigation techniques.

Mitigation - Aborting Program. Different from the timing gap, the results of the daemon process are accurate, and there are no false positives. We can adopt a more severe penalty in this method. The protected PUT can abort the execution or trigger an artificially inserted bug to misinform the crashes to fuzzers. To further avoid the mitigation strategy affecting the regular users in some unexpected situations, the developers can set a long patrolling time (5 min) and a large threshold (1000 files). Regular users can rarely execute the program at such a high frequency.

3.2 Active Methods: Minimum Fake Blocks

Existing fake blocks impede fuzzers at a non-optimal storage cost. ANTIFUZZ [16] (default configuration) introduces about 20 MB of anti-fuzzing codes, while based on the OSS-Fuzz [4] project, most commercial software occupies no more than 100 MB. Besides, small programs are more sensitive to storage costs and are hard to burden high storage costs. Unfortunately, they are more likely to be chosen as the fuzzing targets by attackers due to their faster execution speed and less program logic.

This defect comes to the fore as future fuzzers enhance the capability of fuzzing mechanisms (e.g., size of the bitmap) or improve fuzzing heuristics (e.g., scheduling seeds to avoid triggering anti-fuzzing mechanisms). Anti-fuzzing techniques have to insert more protection codes in proportion to the increased fuzzer capability to handle the intensified arms race. Under this condition, the fundamental storage overhead should be as low as possible; otherwise, the proportioned overhead needed in the arms race will be unsustainable.

The landing space is designed to minimize the extra storage overhead of fake blocks. Existing approaches pile up function calls and constraints, disturbing fuzzers at the function level, where one fake block takes up about nine bytes (one *cmp* and one *jmp* instructions). However, we observe that some of the storage overhead of fake blocks is unnecessary at the assembly level. For example, C compilers generate function frames for each function that controls the base pointer and stack pointer (e.g., *push ebp*). These assembly codes have nothing to do with anti-fuzzing, and eliminating these codes can further reduce the storage overhead.

As adversarial cannot retrieve the source codes of targets, they use BOF with the assistance of external tools to collect coverage feedback. These tools insert codes before entering a new block. If a control-flow-changing instruction (*jmp*, *call* and *ret*) is encountered, they generate a new block record as the updated coverage. Theoretically, the minimum block is the instruction only occupying one byte (opcode of the minimum size), which should be at most 15% (from nine bytes to one byte) of the storage cost of prior works. To achieve the minimum fake block, we instrument each function with a code segment called *landing space*. The landing space contains instructions that have no effect on the normal execution. They are either one-byte instructions or two-byte instructions with an immediate value which is the opcode of a one-byte instruction. This ensures that each byte in the landing space can be translated into a valid instruction.

We further modify the destination address of function calls to the address of a random byte in the landing space. The rationale is that when the modified

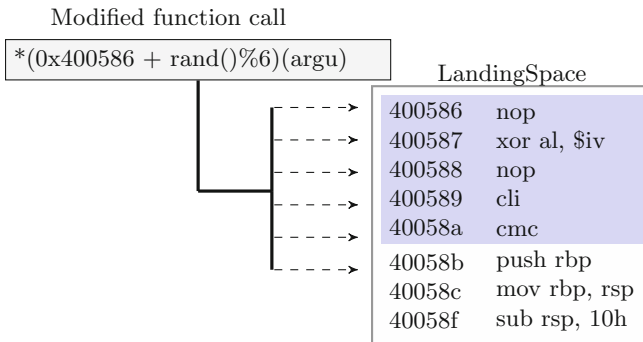


Fig. 1. A function with landing space. \$iv is the immediate value, in this example, it will be 0×90 which is the opcode for nop.

function call is invoked, the control flow “land” at a random instruction. The fuzzer considers this instruction the start of a new block and records the address as new coverage. Since the control flow can “land” at any byte in the landing space, fuzzers will record most of the possible addresses after sufficient rounds of fuzzing. The corresponding fake coverage can overwhelm the fuzzer’s bitmap.

Figure 1 illustrates the assembly codes of a function and the landing space. The original destination address of the function is $0x40058b$, and we insert the landing space before this address in the text section between $0x400586$ and $0x40058a$. The modified function call $(0x400586 + rand()\%6)$ jumps to the landing space or the original start of the function. In this example, a fuzzer will record six fake blocks at the cost of six bytes.

Optimizations. The naive implementation of the landing space seems able to disturb the coverage feedback of BOF. However, we found that the size of the landing space is restricted. A too large size introduces non-negligible latency to the protected program. Moreover, fuzzers like AFL calculate the hash value by exclusive-or operations on the addresses of two blocks. The problem is that the addresses of fake blocks in the landing space are close, and so are the calculated hash values. The chance of hash collision in the landing space is higher than that in normal functions. Intuitively, the more hash collisions happen in the landing space, the less bitmap can be saturated by fake blocks. In other words, the fuzzer will be more powerful in discovering real branches in the protected program. As a means of coping strategy, we propose two optimizations to mitigate the limitations.

Jump Over Unnecessary Instructions. If the control flow lands on the first few bytes in the landing space, it has to execute the rest instructions, which incurs significant latency for a large landing space. To avoid executing the unnecessary instructions, we modify some one-byte instructions to a short jump, and the jumping offset is the opcode of the next instruction. As shown in Fig. 2, if control flow lands at $0x400500$, the assembly code is translated as a two-byte

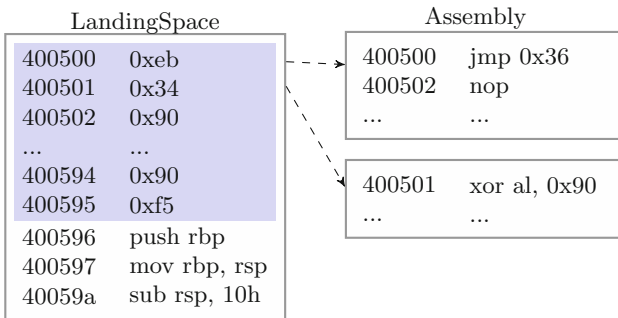


Fig. 2. Jump over unnecessary blocks

short jump with offset $0x36$. However, if it lands at the next byte $0x400501$, the corresponding assembly code is “xor al, 0x90”. The functionality of the landing space still remains as every byte can be disassembled correctly and recorded as a new block. Yet the jump instructions reduce the performance overhead to 5% of the original landing space.

Spray LandingSpace at Different Addresses. To reduce the hash collision rates, we increase the blocks of different addresses. As shown in Fig. 2, we wrap functions in the original program with several intermediate blocks. The intermediate blocks only redirect the control flows in the protected binaries but have no effect on the program execution. We artificially keep wide disparities in the addresses of intermediate blocks; thus, these blocks are likely to generate more hash values than those generated from a single landing space. The size of the landing spaces in these functions is accordingly reduced, and they will be distributed to the intermediate blocks.

4 Evaluation

We evaluate No-Fuzz to answer the following four research questions (**RQs**):

- **RQ 1.** Can No-Fuzz hinder fuzzers from exploring new branches?
- **RQ 2.** How effective are the anti-fuzzing techniques at preventing fuzzers from finding bugs?
- **RQ 3.** What are the storage and performance overhead to deploy anti-fuzzing techniques?
- **RQ 4.** What is the suitable metric to compare different anti-fuzzing techniques?

For **RQ 1**, coverage is considered orthogonal to the bug-finding abilities of fuzzers [9]. The more coverage a fuzzer can reach, the more likely it finds bugs inside the target program. We evaluate the coverage reduction on real-world binaries after applying No-Fuzz to show the defense effectiveness of anti-fuzzing techniques. To stress **RQ 2**, we evaluate the anti-fuzzing techniques on the LAVA-M [13] benchmark and measure the shortest time needed to find a bug. The benchmark consists of four buggy binaries (`base64`, `md5sum`, `who`, `uniq`) with dozens to thousands of artificially inserted bugs.

The **RQ 3**. and **RQ 4**. are both related to the overhead evaluation of anti-fuzzing techniques. To answer **RQ 3**., we evaluate the storage and performance overhead of anti-fuzzing techniques on real-world programs of different sizes. The **RQ 4**. is based on the concern that the overhead alone is not able to judge an anti-fuzzing technique comprehensively. The problem is that if the number of defensive codes added to the protected programs increases, the overhead is likely to increase accordingly. The defensive capability is orthogonal to the extra overhead as the more defensive codes are inserted into the protected binary, the safer it is able to be. Thus judging a defensive technique only based on one

metric (anti-fuzzing effect or the overhead) is not enough. We suggest combining these two metrics and unifying the judging criteria of anti-fuzzing techniques by measuring the defensive ability at a unit cost of storage or execution rate. We define a new metric anti-fuzzing efficacy as the number of reduced branches per byte of extra storage cost and that per millisecond of latency. It measures the capability of anti-fuzzing techniques with respect to the introduced overhead.

In all experiments, the latency mitigation is set to be one second; the daemon process will patrol for one minute and alert if there are more than 60 times of executions; the landing space is configured to occupy 100 bytes, and the functions will be wrapped in 50 intermediate functions. AFL and AFL-based fuzzers (AFLFast and QSYM) use AFL-QEMU. Honggfuzz supports Intel-PT and QEMU, and we adopt both binary-only modes. Each fuzzing campaign runs with three CPU cores. Notably, QSYM runs two AFL instances with two CPU cores and an SMT solver using one core. Fuzzing campaigns on the LAVA-M dataset are kept running for 48 h, while the others last for 24 h. Due to the nondeterministic fuzzing behaviors, we repeat fuzzing campaigns ten times for each fuzzer x target combination.

4.1 Reducing Code Coverage

We evaluated the branch coverage of five fuzzers against eight real-world binaries from Binutils, Magma, and Google FTS. Figure 3 shows the average branches covered by each fuzzer on the binaries with and without No-Fuzz protection. Each technique is separately evaluated to avoid the effect of one technique covering up others. From the figure, the combo of all No-Fuzz techniques can severely hinder the branch explorations of fuzzers. The fuzzers can only discover 36.9% of branches on average that should have been, and most of the branches are just for initializations and input correctness checks.

A single passive detection technique reduces 34.8%–52.2% of the branch coverage. The effect variation of the detection should also be attributed to the choices of mitigation techniques and the difference in fuzzers. As the results show, aborting the PUT (the purple columnar) is more effective than introducing latency (the gray ones) to the protected programs. However, introducing latency affects users more slightly than aborting the program. This is the trade-off between effectiveness and impacts on users. It will be reckless to conclude that one mitigation technique outperforms the others. Our suggestion is to apply the more severe mitigation techniques to the more precise detection techniques.

The landing space hinders, on average, 32.4% of the branches. We observe that it is less effective against Honggfuzz. It is because Honggfuzz records the coverage in a temporary file and the size of the bitmap is 16 M, while the bitmap of AFL only occupies 64 K. We consider the current configuration of the landing space too small to saturate the bitmap of Honggfuzz. In the actual situation, the developers can configure a larger landing space for better protection against Honggfuzz.




Fig. 3. Branch covered by four fuzzers against eight binaries with and without different protections. The techniques are Timing Gap, Daemon Process, Landing Space. The fuzzers are AFL, AFLFast, Honggfuzz-QEMU, Honggfuzz-PT, and QSYM .

4.2 Preventing Fuzzers from Finding Bugs

LAVA-M Benchmark. Despite many recent works suggesting using the up-to-date benchmarks in bug-finding experiments (e.g., Google fuzzer test suite [1] and Magma [18]), we find they are not suitable for the evaluation of BOF. The problem is that these benchmarks heavily rely on sanitizers, but unfortunately, most BOF techniques are not able to support sanitizers. There are some works like QEMU-AddressSanitizer [2] to fill the gap in BOF and sanitizers, but our evaluation covers different BOF techniques, and not all of them have such complementary tools. Due to this limitation, we decided to use LAVA-M, whose bugs directly trigger segmentation faults and can be caught by the BOF.

Moreover, the efficiency of BOF severely degrades compared with static instrumentation, which is almost a quarter of that of the latter. Even the LAVA-M benchmark contains thousands of bugs; only a few unique bugs can be uncovered for each buggy binary in BOF. Measuring the number of bugs found will be insignificant even though it is considered the ground truth for fuzzer evaluation. Instead of the number of bugs found, we measure the time that fuzzers need to find the first bug in each buggy binary within 48h. This metric can better illustrate the bug-finding capabilities of BOF in the LAVA-M benchmark.

Results. The average time of five fuzzers to find one bug in the LAVA-M benchmark is shown in Table 1. From the table, we find that all fuzzers are able to find at least one bug in the unprotected programs in 48h. Note that QSYM finds the bug in only several minutes, which is quite faster than other fuzzers. It can be attributed to the design of the LAVA-M benchmark. The bugs in LAVA-M are all designed based on an integer comparison. If an input can bypass the comparison, the corresponding bug will be triggered. This mechanism is

essentially more beneficial to fuzzers that are able to solve constraints. It makes sense that QSYM, as a hybrid fuzzer, outperforms other mutational fuzzers due to its symbolic execution engine.

On the other hand, with the anti-fuzzing defenses, some fuzzing campaigns exceed the 48-h time limit and fail to find any bug. The rest campaigns uncover some bugs, but it takes much more time than that spent on the corresponding unprotected programs. Notably, we found Honggfuzz is incompatible with the target md5sum in LAVA-M as it misjudges the handled errors as crash signals. Honggfuzz generates millions of crash seed files, and the majority of them are false positives. It is difficult to distinguish the correct crash seeds in the millions of files, and we have to discard this fuzzer x target combination.

Generally, the passive detection techniques and the landing space successfully impede all fuzzing campaigns, as the time to find a bug extends after applying

Table 1. Time of fuzzers to find a bug in native and protected LAVA-M. \checkmark means the fuzzing campaign fails to find a bug within 48 h.

	native	TG	DP	LS	all
base64					
AFL	12 h 54 m	\checkmark	\checkmark	\checkmark	\checkmark
AFLFast	13 h 8 m	\checkmark	\checkmark	\checkmark	\checkmark
Hfz-QEMU	1 h 22 m	\checkmark	\checkmark	\checkmark	\checkmark
Hfz-PT	5 h 23 m	\checkmark	\checkmark	\checkmark	\checkmark
QSYM	2 m	\checkmark	\checkmark	\checkmark	\checkmark
md5sum					
AFL	36 h 32 m	\checkmark	\checkmark	\checkmark	\checkmark
AFLFast	9 h 20 m	\checkmark	\checkmark	\checkmark	\checkmark
Hfz-QEMU	–	–	–	–	–
Hfz-PT	–	–	–	–	–
QSYM	51 m	\checkmark	\checkmark	\checkmark	\checkmark
who					
AFL	3 h 8 m	\checkmark	\checkmark	\checkmark	\checkmark
AFLFast	6 h 37 m	\checkmark	\checkmark	\checkmark	\checkmark
Hfz-QEMU	37 m	2 h 9 m	\checkmark	6 h 32 m	\checkmark
Hfz-PT	4 h 31 m	11 h 45 m	\checkmark	\checkmark	\checkmark
QSYM	1 m	\checkmark	\checkmark	\checkmark	\checkmark
uniq					
AFL	7 h 19 m	\checkmark	\checkmark	\checkmark	\checkmark
AFLFast	6 h 47 m	23 h 59 m	\checkmark	\checkmark	\checkmark
Hfz-QEMU	4 m	10 h 55 m	\checkmark	9 h 3 m	\checkmark
Hfz-PT	2 h 48 m	16 h 20 m	\checkmark	17 h 56 m	\checkmark
QSYM	5 m	\checkmark	\checkmark	\checkmark	\checkmark

these techniques. If all of the anti-fuzzing techniques in No-Fuzz are applied, none of the fuzzing campaigns can find a bug. Overall, the evaluation confirms that No-Fuzz is effective at preventing the BOF of different fuzzers from discovering bugs in the protected programs.

4.3 Performance and Storage Overhead of No-Fuzz

We are inspired by the fact that the size of input files can affect the performance overhead accordingly. Generally, a larger input will invoke more functions and be processed for a longer time. Considering the overhead of the landing space is also accordingly proportional to the functions executed, for fairness, we prepare two sets of input files. One set only contains small invalid files, which can fastly trigger errors in the programs, while the other set consists of valid samples of different sizes to trigger the normal functionalities. The results adopt the average time of the executions with both sets of input samples. Another consideration is that the overhead will be less significant for large and complex programs. The programs are categorized into two groups based on their size and average execution time to mitigate the bias in the evaluations, as shown in Appendix 6. The evaluation results of each group will be analyzed separately.

Performance Overhead. As shown in Table 2, the performance overhead for passive detection techniques is about 10–20% for small binaries and less than 1% for large binaries. Although the overhead is relatively large for small programs, the absolute latency is only around 5 ms, which is usually unnoticeable for regular users. The timing gap detection introduces a little more latency than the daemon process. We think it should be attributed to the false positives of the timing gap. Similarly, the landing space introduces the overhead proportioned to the size and complexity of the programs. Small programs incur 40% latency, while for large programs, the proportion decreases to 2%.

Storage Overhead. From Table 2, No-Fuzz introduces negligible storage overhead to protected binaries. The passive detection techniques take up storage ranging from about 1 KB to 50 KB (10 KB on average), but they are all less than 1% of the original size of the protected programs. The landing space inserts fake blocks according to the number of functions in the protected binary. Basically, the fewer functions in the original program, the less overhead it has. The storage cost can range from 0.3 MB to 1 MB (0.8 M on average) for different programs.

Comparisons with Prior Works. We evaluate the existing anti-coverage techniques in prior works to show that the landing space is worth it. The default configurations of ANTIFUZZ and FUZZIFICATION insert a fixed number of fake blocks; thus the storage overhead is stable - 20 MB for ANTIFUZZ and 1.2 MB for FUZZIFICATION. Clearly, both techniques take up more space than the landing

space, and the storage advantage of No-Fuzz is more evident for small binaries due to the accordingly fewer fake blocks in smaller binaries. Note that the storage overhead of FUZZIFICATION is much smaller than that of ANTIFUZZ. However, according to our experiments (Appendix 4), we find the default configuration of the FUZZIFICATION is not enough to saturate the bitmaps of fuzzers. The real storage overhead of the effective configuration of FUZZIFICATION should be even larger than the current 1.2 MB.

Table 2. Overhead (CPU) of No-Fuzz and anti-coverage techniques of ANTIFUZZ and FUZZIFICATION on real-world programs.

	TG	DP	LS	All	AF(cov)	FZ(cov)	Reference
CPU							
Small	6.8 ms (21.4%)	3.4 ms (10.6%)	13.1 ms (40.9%)	47.2 ms (147.5%)	11.3 ms (35.3%)	14.7 ms (45.9%)	32.0 ms
Large	23.5 ms (1.1%)	10.7 ms (0.5%)	43.8 ms (2.0%)	64.7 ms (3.0%)	15.6 ms (0.7%)	44.6 ms (2.1%)	2156.2 ms
Storage							
Small	8.4 K (0.2%)	10.5 K (0.3%)	0.8 M (25.9%)	1.1 M (35.3%)	22.2 M (696.6%)	1.25 M (39.1%)	3.2 M
Large	43.0 K (0.04%)	27.1 K (0.02%)	2.0 M (1.9%)	2.4 M (2.2%)	22.3 M (21.0%)	1.27 M (1.2%)	106.5 M

4.4 Anti-fuzzing Efficacy

To compare the merits of different anti-fuzzing techniques, we introduce a new metric - anti-fuzzing efficacy. The efficacy illustrates the associations between the anti-fuzzing effects (coverage reduction) and storage/performance overhead. As anti-fuzzing techniques improve the defensive capability by inserting extra codes into the protected programs, more defensive codes promise more defensive effects. A well-designed anti-fuzzing technique should introduce as low overhead as possible while possessing effective defensive capability. The defense capability per unit storage/performance cost should be a better metric to describe the anti-fuzzing effects. Specifically, we calculate the efficacy by the number of coverage reductions every kilobyte of defensive codes and the reduction of every millisecond of latency. As a reference, we also calculate the efficacy of ANTIFUZZ and FUZZIFICATION based on the results in Appendix 4.

From Table 3, the passive detection techniques have both the highest performance and storage anti-fuzzing efficacy. Passive detection techniques are the most affordable anti-fuzzing techniques that can hinder the BOF at the lowest cost. Moreover, these techniques insert fixed defensive codes into the protected programs. Thus, the efficacy is stable and more or less has the same order of magnitude among all evaluated programs. The performance efficacy of the landing space is about 20% of the passive detection techniques, and the storage efficacy is much smaller (only about 1% of the passive detection methods). Despite the fact that the landing space seems to be less efficient, it can be deemed as a complement to the passive detection techniques. As we will discuss in Sect. 5, a single defensive technique is weak against adversaries, and it is worth the development of different techniques.

For reference, we evaluate the anti-coverage techniques of ANTIFUZZ and FUZZIFICATION. The performance efficacy of landing space and ANTIFUZZ is close, while FUZZIFICATION is inefficient due to the insufficient number of blocks. The storage efficacy of the landing space is 14–28 times that of these anti-coverage techniques, i.e., on average, we reduce about 95% of the initial storage cost. It illustrates that No-Fuzz is more practical, which can better utilize the storage while keeping adequate anti-fuzzing protection.

Table 3. Space and performance efficacy of different anti-fuzzing techniques against four fuzzers.

	TG	DP	LS	All	AF(cov)	FZ(cov)
Performance (#branches/ms)						
AFL	183.9	559.3	84.2	46.6	144.1	11.5
AFLFast	165.0	509.3	61.9	42.3	121.7	5.3
Hfz-Q	158.5	530.1	86.2	47.0	133.9	17.0
Hfz-P	134.5	428.4	62.6	36.9	91.9	5.0
QSYM	226.8	684.9	120.7	58.6	174.6	14.8
Avg	173.7	542.4	83.1	46.3	133.2	10.7
Storage (#branches/kB)						
AFL	148.8	181.1	1.4	2.0	0.07	0.14
AFLFast	133.6	164.9	1.0	1.8	0.06	0.15
Hfz-Q	128.3	171.6	1.4	2.0	0.07	0.20
Hfz-P	108.9	138.7	1.0	1.6	0.05	0.06
QSYM	183.6	221.8	2.0	2.5	0.09	0.17
Avg	140.6	175.6	1.4	2.0	0.07	0.14

5 Discussion

While our design performs effectively and efficiently, it could be further improved. We consider our design a complement to prior works, and we appreciate some designs in prior works which also deserve adoption in the future development of anti-fuzzing tools (e.g., installing the signal handler to hide crashes). We will also stress these concerns in our discussion. In the following, we will discuss the robustness of the anti-fuzzing technique and the advantages of anti-fuzzing over obfuscation because the potential inherent problems (i.e., robustness) and the seemingly possibility that they can be substituted by obfuscation techniques in certain cases.

Robustness of Anti-fuzzing Techniques. A primary concern about anti-fuzzing is its robustness. In particular, a dedicated attacker can perform manual analysis to disarm the defense if the details are known, and several prior works suggested applying obfuscation techniques as a countermeasure against reverse engineering. However, for experienced attackers, obfuscation techniques can also be disarmed. We want to reflect the necessity for anti-fuzzing techniques to be robust against manual analysis. Anti-fuzzing techniques are proposed to introduce extra efforts (time, resources, knowledge, etc.) adversaries need to fuzz a protected program. Particularly, they are suitable to defend the untargeted fuzzing tasks on a large scale that are not worth enough to analyze a single binary for attackers manually. Moreover, anti-fuzzing techniques enhance the basic requirements of BOF, from knowing nothing to at least understanding anti-fuzzing and reverse-engineering techniques. The defensive techniques can further reduce the chance that protected binaries are chosen as the targets of BOF. Due to the above reasons, we consider that reverse engineering does not contradict the ultimate purpose of anti-fuzzing techniques.

Anti-fuzzing or Obfuscation. Obfuscation is originally considered a potential solution to anti-fuzzing. Compared with the emerging anti-fuzzing techniques, it is well-developed with profound community support. Prior works have conducted some experiments to show the ineffectiveness of obfuscation in anti-fuzzing [16, 21]. However, we will revise their arguments with more experiments and show that obfuscation techniques can be effective at times.

In fact, there are already obfuscation techniques designed to confront symbolic executions, which are similar to anti-hybrid techniques [6, 34]. In addition, the obfuscation techniques involving self-modifying codes can be even more powerful against BOF. The self-modifying code is a common technique used in packing and encryption. It reuses the memory space by overwriting the existing opcodes with those of new instructions. The problem is that no matter how many functions are overwritten to a self-modifying block, they possess the same memory address. Most fuzzers record the hash values of function block addresses, and the self-modifying block will be identified as only one function; thus, the coverage information of overwritten functions is lost.

We have conducted some experiments with BOF and dummy programs where the programs are protected by self-modifying codes. The results (appendix 7 and 7) show that BOF cannot be correctly performed on the program with self-modifying codes. Besides, [26] shows that self-modifying codes will severely slow down the translations in emulators, which confirms the anti-fuzzing effectiveness of obfuscation against BOF like `af1-qemu` and `honggfuzz-qemu`. Fortunately, self-modifying codes are usually not welcomed by developers. Commercial software rarely uses self-modifying codes due to the possibility of false positives as malicious attempts and the new bugs introduced by risky modifications. Overall, anti-fuzzing techniques cover the shortage of static obfuscation techniques against fuzzers, and we argue that future anti-fuzzing works should be designed in the scope of static techniques without self-modifying codes.

Future Work. We consider that the future work of anti-fuzzing can focus on the following two aspects. On the one hand, we can keep reducing the overhead for existing anti-fuzzing techniques, which increases the number of defensive codes inserted into a program in a disguised way. A possible direction is that the anti-hybrid techniques in prior works are overqualified to disturb the program analysis. They use cryptography functions (e.g., hash, CRC) to wrap the variables. We have conducted some experiments, and it turns out that only hundreds of calculations are enough to overwhelm the symbolic executions, and they are cheaper than the heavy cryptography functions. On the other hand, future anti-fuzzing works can embed anti-fuzzing mechanisms into the program logic. For instance, similar to the flatten technique in obfuscation, we can split a block into several new blocks, and each new block contains a part of the assembly codes of the original block. Thus, each new block is logically dependent on the program and cannot be easily eliminated. Indeed, there are some challenges that need to be solved for this idea, e.g., how to maintain the context among different blocks and how to ensure the number of fake blocks is large enough.

6 Conclusion

In this paper, we design several practical and fully-automated anti-fuzzing techniques and integrate them into a prototype tool No-Fuzz. We optimize the storage cost of fake blocks as prior works insert them at the function level occupying an unrealistic storage room. In addition to the active anti-fuzzing techniques that disturb the fuzzing mechanisms, we also design the passive detection methods which precisely determine the running environments of the protected programs and launch mitigation techniques when binary-only-fuzzing exists. The evaluations demonstrate that No-Fuzz significantly reduces the branch coverage of fuzzers. Furthermore, we have also shown that No-Fuzz can impede bug findings in the LAVA-M dataset, i.e., fuzzers have to spend much more time finding a bug. We propose a new metric, the anti-fuzzing efficacy, to measure the defensive capability of an anti-fuzzing technique at a unit overhead cost. Based on this metric, we illustrate that No-Fuzz achieves the same or higher level of protection against fuzzers with even lower overhead than prior works.

In summary, we enhance the awareness of overhead and the importance of automation in an anti-fuzzing arms race. Inspired by this, we summarize the desired properties for future anti-fuzzing techniques - be with less overhead and automated. We have moved one step toward practical anti-fuzzing techniques and hope our efforts can further promote this topic.

Acknowledgements. This work was supported by the Research Grants Council of Hong Kong under Grants CityU 11217819, 11217620, 11218521, N_CityU139/21, RFS2122-1S04, C2004-21GF, R1012-21, and R6021-20F.

A Appendix

(See Tables 4, 5, 6, 7 and 8)

Table 4. The branch coverage of **ANTIFUZZ** and **FUZZIFICATION**

	readelf	objdump	size	nm	xmllint	libpng	guetzli	sqlite3
AFL								
AF(cov)	1745	1449	839	706	1691	919	2103	3840
FZ(cov)	4498	2630	1247	1529	5556	1702	2805	4994
AFLFast								
AF(cov)	1913	1345	656	774	1503	958	2207	3757
FZ(cov)	3383	2848	1102	1411	5463	1492	2727	5065
Hfz-Q								
AF(cov)	1049	875	622	894	4117	1073	2499	4103
FZ(cov)	2422	2812	1277	1523	6122	1996	3040	6144
Hfz-P								
AF(cov)	1884	1067	899	632	3555	788	1931	3097
FZ(cov)	2612	2374	1148	1190	5312	1540	2563	4833
QSYM								
AF(cov)	1946	1276	814	932	4255	1287	2493	4441
FZ(cov)	5072	3731	1448	1740	6442	2084	3538	7433

Table 5. The overhead comparisons between upx and existing anti-fuzzing techniques. As emphasized in the table, the overhead of existing anti-fuzzing techniques is more or less close to that of packing techniques.

	readelf	objdump
Exec time	125.4 ms	2156.2 ms
upx	+13.0%	+24.3%
AF(cov)	+37.7%	+2.0%
FZ(cov)	+32.9%	+14.6%
Storage cost	3.22 M	9.94 M
upx	-2.24 M (-69.6%)	-7.44 M (-74.8%)
AF(cov)	21.27 M (+660.6%)	21.25 M (+213.8%)
FZ(cov)	1.29 M (+40.1%)	1.28 M (+12.9%)

Table 6. Real-world programs of different size and execution time.

size & exec time	files
Small	readelf, objdump, size, nm guetzli, libpng, sqlite3, xmllint
Large	ffmpeg_g, nomacs, calc, impress

Table 7. We fuzz a dummy program as well as the obfuscated versions. The dummy program contains several magic-byte checks and will crash if the constraints are satisfied. Tigress(S) only applies static obfuscation, while Tigress(D) adopts self-modifying codes which are dynamic.

	First crash	Fuzz rate
Native	45 s	1805 exc/s
UPX	1 m 27 s	1800 exc/s
Tigress(D)	+ ∞	0 exc/s
Tigress(S)	2 m 4 s	1667 exc/s
llvm-obfuscator	1 m 48 s	1400 exc/s

Table 8. Evaluations to launch BOF on obfuscated/packed binaries. \times means BOF cannot initialize on the binary within 30 min for all four fuzzers. \checkmark means all fuzzers succeed in launching the BOF. As the table suggests, self-modifying codes (upx & obfuscation with JIT) can completely prevent the BOF from initialization.

	native	upx	llvm-obf	Tigress (D)	Tigress (S)
dummy	\checkmark	\checkmark	\checkmark	\times	\checkmark
binutils	\checkmark	\times	\checkmark	\times	\checkmark
libjpeg	\checkmark	\times	\checkmark	\times	\checkmark
libpng	\checkmark	\times	\checkmark	\times	\checkmark
libtiff	\checkmark	\times	\checkmark	\times	\checkmark
ffmpeg	\checkmark	\times	\checkmark	\times	\checkmark
gzip	\checkmark	\times	\checkmark	\times	\checkmark

References

1. Binary-only fuzzing of honggfuzz. <https://github.com/google/fuzzer-test-suite>. Accessed 12 Mar 2022
2. QASan (QEMU-AddressSanitizer). <https://github.com/andreaforaldi/qasan>. Accessed 12 Mar 2021
3. A library for coverage-guided fuzz testing. <https://llvm.org/docs/LibFuzzer.html> (2017) Accessed 23 Oct 2020
4. Aizatsky, M., Serebryany, K., Chang, O., Arya, A.: Announcing oss-fuzz: continuous fuzzing for open source software. Google Testing Blog, Announcing OSS-Fuzz (2016)

5. Balzarotti, D., et al.: Efficient detection of split personalities in malware. In: Proceedings of the NDSS (2010)
6. Banescu, S., Collberg, C., Ganesh, V., Newsham, Z., Pretschner, A.: Code obfuscation against symbolic execution attacks. In: Proceedings of the ACSAC (2016)
7. Böhme, M., Pham, V.-T., Roychoudhury, A.: Coverage-based greybox fuzzing as Markov chain. In: Proceedings of the ACM CCS (2016)
8. Böhme, M., Pham, V.-T., Nguyen, M.-D., Roychoudhury, A.: Directed greybox fuzzing. In: Proceedings of the ACM CCS (2017)
9. Böhme, M., Szekeres, L., Metzman, J.: On the reliability of coverage-based fuzzer benchmarking. In: Proceedings of the IEEE ICSE (2022)
10. Chen, H., et al.: Hawkeye: towards a desired directed grey-box fuzzer. In: Proceedings of the ACM CCS (2018)
11. Chen, Y., et al.: SAVIOR: towards bug-driven hybrid testing. In: Proceedings of the IEEE S&P (2020)
12. Dinesh, S., Burow, N., Xu, D., Payer, M.: RetroWrite: statically instrumenting cots binaries for fuzzing and sanitization. In: Proceedings of the IEEE S&P (2020)
13. Dolan-Gavitt, B.: LAVA: large-scale automated vulnerability addition. In: Proceedings of the IEEE S&P (2016)
14. Google: A scalable fuzzing infrastructure. <https://github.com/google/clusterfuzz>. Accessed 23 Oct 2020
15. Google: syzkaller found bugs - Linux kernel. https://github.com/google/syzkaller/blob/master/docs/linux/found_bugs.md
16. Güler, E., Aschermann, C., Abbasi, A., Holz, T.: ANTIFUZZ: impeding fuzzing audits of binary executables. In: Proceedings of the USENIX Security (2019)
17. Hafiz, M., Fang, M.: Game of detections: how are security vulnerabilities discovered in the wild? In: Proceedings of ACM ESE (2015)
18. Hazimeh, A., Herrera, A., Payer, M.: Magma: a ground-truth fuzzing benchmark. In: Proceedings of ACM Measurement and Analysis of Computing Systems (2020)
19. Zhenghao, H., Yu, H., Dolan-Gavitt, B.: Chaff bugs: deterring attackers by making software buggier. arXiv (2018)
20. Huang, H., Yao, P., Wu, R., Shi, Q., Zhang, C.: Pangolin: Incremental hybrid fuzzing with polyhedral path abstraction. In: Proceedings of IEEE S&P (2020)
21. Jung, J., Hong, H., Solodukhin, D., Pagan, D., Hyung Lee, K., Kim, T.: Fuzzification: anti-fuzzing techniques. In: Proceedings of USENIX Security (2019)
22. Kang, M.G., Yin, H., Hanna, S., McCamant, S., Song, D.: Emulating emulation-resistant malware. In: Proceedings of ACM Workshop on Virtual Machine Security (2009)
23. Lemieux, C., Sen, K.: Fairfuzz: A targeted mutation strategy for increasing greybox fuzz testing coverage. In: Proceedings of the ACM/IEEE ASE (2018)
24. Lindorfer, M., Kolbitsch, C., Comporetti, P.M.: Detecting environment-sensitive malware. In: Proceedings of the International Workshop on Recent Advances in Intrusion Detection (2011)
25. Miller, B.P., Fredriksen, L., So, B.: An empirical study of the reliability of unix utilities. In: Proceedings of the ACM Communication (1990)
26. Raffetseder, T., Kruegel, C., Kirda, E.: Detecting system emulators. In: Proceedings of the International Conference on Information Security (2007)
27. Rash, M.: A collection of vulnerabilities discovered by the afl fuzzer. <https://github.com/mrash/afl-cve>. Accessed 13 Sept 2020
28. Rawat, S., et al.: Application-aware evolutionary fuzzing. In: Proceedings of the NDSS, Vuzzer (2017)

29. Schumilo, S., Aschermann, C., Gawlik, R., Schinzel, S., Holz, T.: KAFLE: Hardware-assisted feedback fuzzing for OS kernels. In: Proceedings of the USENIX Security (2017)
30. Stephens, N., et al.: Driller: augmenting fuzzing through selective symbolic execution. In: Proceedings of the NDSS (2016)
31. Swiecki, R.: Honggfuzz (2020). <https://github.com/google/honggfuzz>
32. Votipka, D., Stevens, R., Redmiles, E., Hu, J., Mazurek, M.: Hackers vs. testers: a comparison of software vulnerability discovery processes. In: Proceedings of the IEEE S&P (2018)
33. Wang, T., Wei, T., Gu, G., Zou, W.: TaintScope: a checksum-aware directed fuzzing tool for automatic software vulnerability detection. In: Proceedings of the IEEE S&P (2010)
34. Wang, Z., Ming, J., Jia, C., Gao, D.: Linear obfuscation to combat symbolic execution. In: Proceedings of the ESORICS (2011)
35. Yun, I., Lee, S., Xu, M., Jang, Y., Kim, T.: QSYM: a practical concolic execution engine tailored for hybrid fuzzing. In: Proceedings of USENIX Security (2018)
36. Zalewski, M.: American fuzzy lop (2019). <http://lcamtuf.coredump.cx/afll>
37. Zalewski, M.: Technical “whitepaper” for afl-fuzz (2019). http://lcamtuf.coredump.cx/afl/technical_details.txt
38. Zhang, Z., You, W., Tao, G., Aafer, Y., Liu, X., Zhang, X.: STOCHFUEZZ: sound and cost-effective fuzzing of stripped binaries by incremental and stochastic rewriting. In: Proceedings of IEEE S&P (2021)