ATFuzzer: Dynamic Analysis Framework of AT Interface for Android Smartphones

IMTIAZ KARIM and FABRIZIO CICALA, Purdue University
 SYED RAFIUL HUSSAIN, Pennsylvania State University
 OMAR CHOWDHURY, University of Iowa
 ELISA BERTINO, Purdue University

Application processors of modern smartphones use the AT interface for issuing high-level commands (or AT-commands) to the baseband processor for performing cellular network operations (e.g., placing a phone call). Vulnerabilities in this interface can be leveraged by malicious USB or Bluetooth peripherals to launch pernicious attacks. In this article, we propose ATFuzzer, which uses a grammar-guided evolutionary fuzzing approach that mutates production rules of the AT-command grammar instead of concrete AT commands to evaluate the correctness and robustness of the AT-command execution process. To automate each step of the analysis pipeline, ATFuzzer first takes as input the 3GPP and other vendor-specific standard documents and, following several heuristics, automatically extracts the seed AT command grammars for the fuzzer. ATFuzzer uses the seed to generate both valid and invalid grammars, following our cross-over and mutation strategies to evaluate both the integrity and execution of AT-commands. Empirical evaluation of ATFuzzer on 10 Android smartphones from 6 vendors revealed 4 invalid AT command grammars over Bluetooth and 14 over USB with implications ranging from DoS, downgrade of cellular protocol version, to severe privacy leaks. The vulnerabilities along with the invalid AT-command grammars were responsibly disclosed to affected vendors and assigned CVE's.

CCS Concepts: • **Security and privacy** → **Mobile and wireless security**; *Distributed systems security*; Denial-of-service attacks;

Additional Key Words and Phrases: Android smartphone security and privacy, vulnerabilities, attack

ACM Reference format:

Imtiaz Karim, Fabrizio Cicala, Syed Rafiul Hussain, Omar Chowdhury, and Elisa Bertino. 2020. ATFuzzer: Dynamic Analysis Framework of AT Interface for Android Smartphones. *Digit. Threat.: Res. Pract.* 1, 4, Article 23 (December 2020), 29 pages. <https://doi.org/10.1145/3416125>

1 INTRODUCTION

Early mobile phones used to have only a single processor, called baseband processor, implementing the cellular modem to control the radio communications for different cellular networks, such as GSM, UMTS, and LTE. Those

Authors' addresses: I. Karim, F. Cicala, and E. Bertino, Purdue University, 305 N University St, West Lafayette, IN 47907; emails: {karim7, ficiala, bertino}@purdue.edu; S. R. Hussain, Pennsylvania State University, University Park, PA 16802; email: hussain1@psu.edu; O. Chowdhury, University of Iowa, Iowa City, IA 52242; email: omar-chowdhury@uiowa.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2576-5337/2020/12-ART23 \$15.00

<https://doi.org/10.1145/3416125>

early devices were, therefore, capable of supporting only phone calls and SMS. With the recent advancements in micro-architectures and operating systems, modern smartphones run on an interconnected dual CPU architecture consisting of a general purpose application processor and a baseband processor. The application processor provides a self-contained operating environment that delivers necessary system capabilities required to support user applications, including parallel processing, I/O and memory management, graphics processing, and multi-media decoding. The application processor can issue ATtention (AT) commands [59] through the radio interface layer (RIL, also called AT interface) to interact with the baseband processor for performing different cellular network operations (e.g., placing a phone call). Most of the modern smartphones also accept AT commands issued by peripherals connected via Bluetooth or USB.

Problem and scope. The AT interface is an entry point for accessing the baseband processor. Therefore, any incorrect execution behavior in processing AT commands may cause unauthorized access to private information, inconsistent system states, and crashes of the RIL daemon and the telephony stack. This article thus focuses on developing a systematic approach for analyzing the correctness and robustness of the baseband-related AT command execution process to uncover vulnerabilities that can be easily exploited to carry out attacks.

Incorrect execution of AT commands may manifest in one of the following forms: (1) *Syntactic errors*: the device accepts and processes syntactically invalid AT commands. For instance the AT command `AT + COPN; III` is processed by Nexus5, LG G3, and many other COTS smartphones and leaks sensitive information. However, the command does not conform to any command syntax from the specifications; and (2) *Semantic violations*: the device processes syntactically correct AT commands, but does not conform to the prescribed behavior. The command `AT + COPS = 0, 1, c19vf, 2` conforms to the standards, but on LG G3 it causes a crash in the RIL daemon.

A successful exploitation of such invalid commands may enable malicious peripheral devices (e.g., a headset), connected to the smartphone over Bluetooth, to an access phone's sensitive information, such as IMSI (International Mobile Subscriber Identity, unique to a subscriber) and IMEI (International Mobile Equipment Identity, unique to a device), or to downgrade the cellular protocol version or stop the cellular Internet, even when the peripheral is only allowed to access the phone's call and media audio.

Prior efforts. Existing approaches [42, 43, 55] strive to identify the types of valid AT commands (i.e., commands with valid inputs/arguments conforming to the 3GPP reference [18, 22, 24, 25, 54] or vendor-specific commands [3, 19, 20, 33] added for vendor customization) exposed through USB interfaces on modern smartphone platforms and the functionality they enable. Yet these studies have at least one of the following limitations: (A) The analyses [55] do not test the robustness of the AT interface in the face of invalid commands; (B) The analyses [55] only consider the USB interface and thus leave the Bluetooth interface exposed to invalid AT commands; and (C) The analyses [12, 42, 44, 48] are not general enough to be applicable to smartphones from different vendors.

Challenges. Conceptually, one can approach our problem using one of the following two techniques: (1) static analysis; (2) dynamic analysis. As the source code of firmware is not always available, a static analysis-based approach would have to operate on a binary level. The firmware binaries, when available, are often obfuscated and encrypted. Making such binaries amenable to static analysis requires substantial manual reverse-engineering effort. To make matters worse, such manual efforts are often firmware-version-specific and may not apply to other firmware, even from the same vendor. Dynamic analysis-based approaches also often require instrumenting the binary to obtain coverage information for guiding the search. Like static analysis, such instrumentation requires reverse-engineering effort that again is not scalable. Also, during dynamic analysis, due to the separation of the two processors, it is often difficult to programmatically detect observable RIL crashes from the application processor. Finally, in many cases, undesired AT commands are blacklisted [37–39] and hence can induce rate-limiting by completely shutting down the AT interface. The only way to recover from such a situation is to reboot the test device, which can substantially slow down the analysis.

Our approach. In this article, we propose ATFuzzer—a framework that overcomes the limitations of current static and dynamic analysis approaches to test the correctness and robustness of the AT interface of COTS android smartphones. One of the key objectives driving the design of ATFuzzer is discovering problematic input formats instead of just some misbehaving concrete AT commands. Towards this goal, ATFuzzer employs a grammar-guided evolutionary fuzzing-based approach. Generation-based fuzzers [13, 14, 56] typically use a grammar to generate syntactically correct inputs so they can be correctly parsed by the program-under-test. Consequently, this allows the fuzzer to exercise some critical functionality implemented deeply in the code. Mutation-based fuzzers [10, 28–30, 45, 62], on the contrary, blindly mutate higher-ranked concrete input instances to generate new inputs. Ranking of an input instance is often measured by its ability to uncover new code or triggering a crash. Unlike typical mutation-based and generation-based fuzzers, ATFuzzer follows a different strategy. It mutates the production rules of the AT command grammars and uses sampled instances of the generated grammar to fuzz the test programs. Our approach has the following two clear benefits: First, a production rule (respectively, grammar) describing a valid AT command can be viewed as a symbolic representation for a set of concrete AT commands. Such a symbolic representation enables ATFuzzer to efficiently navigate the input search space by generating a diverse set of concrete AT command instances for testing. The diversity of fuzzed input instances is likely achieved because mutating a grammar can move the fuzzer to test a different syntactic class of inputs with high probability. Second, if ATFuzzer can find a problematic production rule whose sampled instances can regularly trigger an incorrect behavior, the production rule can then be used as an evidence that can contribute towards the identification of the underlying flaw that causes the misbehavior.

ATFuzzer is an automated fuzzing framework that takes as input the 3GPP and other vendor-specific standard documents and following several heuristic generates the seed grammar for the fuzzer. This AT grammar of AT commands is used to generate the initial population of grammars by mutating the seed grammars. Each grammar from this initial population is sampled by ATFuzzer to generate grammar-compliant random inputs. The fitness of each grammar is evaluated through these inputs following our proposed fitness function. Since code-coverage or subtle memory corruptions are not suitable to be used as the fitness function for such vendor-specific, closed-source firmware, we leverage the execution timing information of each AT command as a loose-indicator of code-coverage information. Based on the fitness score of each grammar, ATFuzzer selects the parent grammars for crossover operation. We design a grammar-aware *two-point* crossover operation to generate a diverse set of valid and invalid grammars. After the crossover operations, we incorporate six proposed mutation strategies to include randomness within the grammar itself. The intuition behind using both crossover and mutation operations is to test the integrity of each command field as well as the command sequence.¹ Through running the fuzzer for sometime, we report the incorrect AT grammars that have the highest fitness score.

Findings. To evaluate the generality and effectiveness of our approach, we evaluated ATFuzzer on 10 Android smartphones (from 6 different vendors) with both Bluetooth and USB interfaces. ATFuzzer has been able to uncover a total of 4 erroneous AT grammars over Bluetooth and another 14 AT grammars over USB. Impacts of these errors range from complete disruption of cellular network connectivity to retrieval of sensitive personal information. We show practical attacks through Bluetooth that can downgrade or shutdown Internet connectivity, and also enable illegitimate exposure of IMSI and IMEI when such impacts are not achievable through valid AT commands. In addition, the syntactically and semantically flawed AT commands over USB can cause crashes, compound actions, and syntactically incorrect commands to be processed. For instance, an invalid AT command ATDI in LG Nexus 5 induces the program to execute two valid AT commands—ATD (dial) and ATI (display IMEI)—simultaneously. These anomalies add a new dimension to the attack surface when blacklisting or access control mechanisms are put in place to protect the devices from valid yet unsafe AT commands. The vulnerabilities along with the invalid AT command grammars were responsibly disclosed to the

¹ Achieving some cellular network operations through the AT interface (e.g., sending an SMS) may require issuing a sequence of AT commands instead of a single AT command.

affected vendors. Among the discovered vulnerabilities, two of them (i.e., DoS and privacy leak attacks) have been assigned CVEs (CVE-2019-16400 [40] and CVE-2019-16401 [41]). Moreover, to test the effectiveness of our AT grammar extraction process, we apply it to four AT command standard PDF's, extracting more than 300 AT commands compared to the manually extracted 75 grammars with a false positive rate of less than 2%.

Contributions. The article has the following contributions:

- (1) We propose ATFuzzer—an automated and systematic framework that leverages grammar-guided, evolutionary fuzzing for dynamically testing the AT command interface in modern Android smartphones. We have made our framework open-source alongside the corpus of AT command grammars we tested. The tool and its detailed documentation are publicly available at: <https://github.com/Imtiazkarimik23/ATFuzzer>.
- (2) We develop an automated AT grammar extractor that automatically extracts seed AT grammars from the specifications by leveraging two heuristics.
- (3) We show the effectiveness of our approach by uncovering 4 problematic AT grammars through Bluetooth and 14 problematic grammars through USB interface on 10 smartphones from 6 different vendors.
- (4) We demonstrate that all the anomalous behavior of the AT program exposed through Bluetooth are exploitable in practice by adversaries, whereas the anomalous behavior of AT programs exposed through USB would be effectively exploitable even when valid but unsafe AT commands are blacklisted. The impact of these vulnerabilities ranges from private information exposure to persistent denial-of-service attacks.

2 BACKGROUND

In this section, we introduce background concepts on the AT commands and the interfaces to issue such commands. We also show how to extract the AT commands supported by a specific smartphone and derive the related grammars.

2.1 AT Commands

In addition to the set of AT commands defined by the standards for cellular networks [24], vendors of cellular baseband processors and operating systems support vendor-specific AT commands [3, 17, 20] for testing and debugging purposes. Based on the functionality, different AT commands have different formats, differing in number and types of parameters. According to the specific usage of a command, there are four primary types of AT commands:

- (1) *Execution command* to execute an action, e.g., ATH, causes the device to hang up the current phone call.
- (2) *Read command* to get/read a parameter value, e.g., AT + CFUN?, returns the current parameter setting of +CFUN, which controls cellular functionalities.
- (3) *Test command* to test for allowed parameters, e.g., AT + CFUN =?, returns the allowed parameters for +CFUN command.
- (4) *Set command* to set/write a parameter, e.g., AT + CFUN = 0, turns off (on) cellular connectivity (airplane mode).

Note that, +CFUN is a variable that can be instantiated with different functionality (e.g., +CFUN=1 refers to setting up the phone with full functionality).

2.2 AT Interfaces for Smartphones

AT commands can be invoked by an application running on the smartphone or from another host machine or peripheral device connected through the smartphone's USB or Bluetooth interface (shown in Figure 1). While older generations of Android smartphones allowed installed applications to run AT commands, recent Android smartphones have restricted this feature to prevent arbitrary applications from accessing device's sensitive

Fig. 1. AT Interface for Android Smartphones connected to a host machine through USB interface.

resources illegitimately through AT commands. However, unlike the case of installed applications, nearly all Android phones allow one to execute AT commands over Bluetooth, whereas for USB devices require a minimal configuration to be set up to activate this feature. Android smartphones typically have different parsers for executing AT commands over these interfaces. The use of different parsers motivates the testing of AT interface for both USB and Bluetooth.

2.3 Issuing AT Commands over Bluetooth and USB

In this section, we introduce details pertaining to issuing AT commands over Bluetooth and USB.

2.3.1 Bluetooth. For executing AT commands over Bluetooth, the injecting host machine/peripheral device needs to be paired with the Android smartphone. The Bluetooth on a smartphone may have multiple profiles (services), but only certain profiles (e.g., hands-free profile, headset profile) support AT commands. Figure 1(b) shows the flow of AT command execution over Bluetooth.

When a device is paired with the host machine, it establishes and authorizes a channel for data communication. After receiving an AT command, the system-level component of the Bluetooth stack recognizes the AT command with the prefix “AT” and compares it against a list of permitted commands (based on the connected Bluetooth profile). When the parsing is completed, the AT command is sent to the application-level component of the Bluetooth stack in the user space where the Bluetooth API takes the action as per the AT command issued. Similar to the example through USB, if a baseband-related command is invoked (e.g., ATD <phone_no>;), the RILD is triggered to deliver the command to the baseband processor. Contrary to USB, only a subset of the AT commands related to specific profiles are accepted/processed through Bluetooth.

2.3.2 USB. If a smartphone exposes its USB Abstract Control Model (ACM) interface [48], it creates a tty device such as `/dev/ttyACM0`, which enables the phone to receive AT commands over the USB interface. However, in phones for which the USB modem interface is not included in the default USB configuration, switching to alternative USB configuration [48] enables communication to the modem over USB. The modem interface appears as `/dev/ttyACM*` device node in Linux, whereas it appears as a COM* port in Windows. Figure 1(a) shows the execution path of an AT command over USB.

When the *AT command injector* running on a host machine sends a command through `/dev/ttyACM*` or COM* to a smartphone, the `ttyABC` (ABC is a placeholder for actual name of the tty device) device in the smartphone receives the AT command and relays it to the native daemon in the Android userspace. The native daemon takes actions based on the type of command. If the command is related to baseband, for instance, ATD <phone_no>;, the RILD (Radio Interface Layer Daemon) is triggered to deliver the command to the baseband processor, which executes the command—makes a phone call to the number specified by <phone_no>. However, if the command is

operating system-specific (e.g., Android, iOS, or Windows), such as +CKPD for tapping a key, the native daemon does not invoke RILD.

2.4 AT Commands and Their Grammars

We obtain the list of valid AT commands and their grammars from the 3GPP standards [18, 21–25, 54]. Note that not every standard AT command is processed/recognized by all smartphones. This is because different smartphone vendors enforce different whitelisting and blacklisting policies for minimizing potential security risks. Also, vendors often implement several undocumented AT commands. Any problematic input instance that ATFuzzer finds, we check to see whether it is one of the vendor-specific, undocumented AT commands following the approach by Tian et al. [55]. We do not report the undocumented, vendor-specific AT commands that ATFuzzer discovers as invalid, since they have already been documented [55]. We aim at finding malformed AT command sets that are due to the parsing errors in the AT parser itself.

3 OVERVIEW OF OUR APPROACH

In this section, we first present the threat model and then formally define our problem statement. Finally, we provide a high-level overview of our proposed mechanism with a running example.

3.1 Threat Model

For Bluetooth and USB AT interfaces exposed by modern smartphones, we define the following two different threat models.

3.1.1 Bluetooth Threat Model. In the threat model we consider for Bluetooth, the adversary possesses a malicious Bluetooth peripheral, or can compromise a benign one, that can perform a complete pairing procedure with the target Android device. The adversary's equipment can be any type of Bluetooth peripheral, such as headphones, speakers, smartwatch, and it connects through its default profile. The victim's smartphone grants only specific permissions to the adversary's device based on the default profile. For instance, if the adversary uses a set of malicious headphones, the victim's smartphone only allows audio permissions. Also, it can be the case that the adversary sets up a fake peripheral device through a man-in-the-middle (MitM) attack exploiting known vulnerabilities of Bluetooth pairing and bonding [2, 32, 51] procedures. Finally, in our threat model, we assume that the adversary does not have physical access to the target device and cannot install any malicious apps on it.

3.1.2 USB Threat Model. In the threat model for USB, we consider an adversary who has control over a USB host, such as a PC or a USB charging station. Hence, the adversary can establish a USB connection with the target Android device through the malicious host. We assume that the attacker can gain access to the exposed AT interface even if the victim's smartphone is in idle state and, thus, send AT commands to the device. Similarly to the Bluetooth threat model, the adversary cannot install any malicious application on the target smartphone. Moreover, we do not require the USB debugging option of the device to be turned on.

3.2 Problem Statement

Let $\mathcal{I}$ be the set of finite strings over printable ASCII characters, $\mathcal{R} = \{\text{ok}, \text{error}\}$ be the set of parsing statuses, and $\mathcal{A}$ be a set of actions (e.g., phone-call $\in \mathcal{A}$). The AT interface of a smartphone can be viewed as a function $\mathcal{P}$ from $\mathcal{I}$ to $\mathcal{R} \times 2^{(\mathcal{A} \cup \{\text{nop}, \perp\})}$, that is, $\mathcal{P} : \mathcal{I} \rightarrow \mathcal{R} \times 2^{(\mathcal{A} \cup \{\text{nop}, \perp\})}$ in which “nop” refers to *no operation* whereas $\perp$ captures undefined behavior including a crash. nop is used to capture the behavior of $\mathcal{P}$ ignoring an AT command, possibly, due to blacklisting or parsing errors.

Given the smartphone AT interface under test $\mathcal{P}_{\text{Test}}$ and a reference AT interface induced by the standard $\mathcal{P}_{\text{Ref}}$, we aim to identify concrete *vulnerable* AT command instances $s \in \mathcal{I}$ such that $\mathcal{P}_{\text{Test}}$ and $\mathcal{P}_{\text{Ref}}$ do not agree on

Fig. 2. Partial Abstract Syntax Tree (AST) of the reference grammar (grey box denotes non-terminal symbols and white box indicates terminal symbols).

their response for s , that is, $\mathcal{P}_{\text{Ref}}(s) \neq \mathcal{P}_{\text{Test}}(s)$. Given pairs $\langle r_1, a_1 \rangle, \langle r_2, a_2 \rangle \in \mathcal{R} \times 2^{(\mathcal{A} \cup \{\text{nop}, \perp\})}$, we write $\langle r_1, a_1 \rangle = \langle r_2, a_2 \rangle$ if and only if $r_1 = r_2$ and $a_1 = a_2$. Note that, a_1 and a_2 are both sets of actions as one command can mistakenly trigger multiple actions.

Note that, there can be a reason $\mathcal{P}_{\text{Ref}}$ and $\mathcal{P}_{\text{Test}}$ can legitimately disagree on a specific input AT command $s \in \mathcal{I}$ as s can be blacklisted by $\mathcal{P}_{\text{Test}}$. Due to CVE-2016-4030 [37], CVE-2016-4031 [38], and CVE-2016-4032 [39], Samsung has locked down the exposed AT interface over USB with a command whitelist for some phones. In this case, we do not consider s to be a vulnerable input instance. Precisely, when s is a blacklisted command, we observed that $\mathcal{P}_{\text{Test}}$ often returns $\langle \text{ok}, \text{nop} \rangle$. Finally, we instantiate the oracle $\mathcal{P}_{\text{Ref}}$ through manual inspection of the standard.

3.3 Running Example

To explain ATFuzzer’s approach, we now provide a partial, example context-free grammar (CFG) of a small set of AT commands (see Figure 3 for the grammar and Figure 2 for the partial Abstract Syntax Tree (AST) of the grammar) that we adopted from the original 3GPP suggested grammar [24, 54]. In our presentation, we use the bold-faced font to denote non-terminals and regular-faced font to denote terminals. We use “.” to represent explicit concatenation, especially, to make the separation of terminals and non-terminals in a production rule clear. We use $[\dots]$ to define regular expressions in grammar production rules and $[\dots]^*$ to represent the Kleene star operation on a regular expression denoted by $[\dots]$. In our example, **Dnum** can take any alphanumeric string up to length n as an argument. Our production rules are of the form: $s \rightarrow \alpha \cdot \mathbf{B}_1\{\phi\}$ where s is a non-terminal, α denotes a (possibly empty) sequence of terminals, $\mathbf{B}_1$ represents a possibly empty sequence of non-terminals and terminals, and ϕ represents a condition that imposes additional well-formedness restrictions on the production.

In the above example, we show the correct AT command format for making a phone call. Examples of valid inputs generated from this grammar can be—ATD * 752# + 436644101453.

3.4 Overview of ATFuzzer

In this section, we first touch on the technical challenges that ATFuzzer faces and how we address them. We conclude by providing the high-level operational view of ATFuzzer.

3.4.1 Challenges. For effectively navigating the input search space and finding vulnerable AT commands, ATFuzzer has to address the following four technical challenges:

C1: (Efficient extraction of AT grammars). ATFuzzer needs to be provided with a large number of seed AT grammars from which the initial seed population is generated. This task, as performed in the initial version of ATFuzzer [55], is highly manual and labor-intensive as one has to review hundreds of pages of 3GPP [21–24] and other

```

command → AT · cmd
cmd → dgrammar | cfungrammar
cmd → ε | cmd_AT
cmd_AT → cmd ; cmd
dgrammar → D · Dnum · Darg ;
cfungrammar → +CFUN? | +CFUN = CFUNarg1, CFUNarg2
cmd → +CTFR = number, type, subaddr, satype
Dnum → [a - zA - Z0 - 9 + *#]* {Length(Dnum) ≤ n}
Darg → l | G | ε
CFUNarg1 → [0 - 9]* {CFUNarg1 ∈ ℤ and 0 ≤ CFUNarg1 ≤ 127}
CFUNarg2 → [0 - 1] {CFUNarg2 ∈ ℤ and 0 ≤ CFUNarg2 ≤ 1}
number1 → number1 | number2
number2 → [a - zA - Z0 - 9 + *#]* {if type = 145}
number → [a - zA - Z0 - 9 + *#]* {if type = 129}
type → 145 | 129
subaddr → [a - zA - Z0 - 9 + *#]*
satype → [0 - 9]* {if satype = ε, satype = 128}
:
:

```

Fig. 3. Partial reference context-free grammar for AT commands.

standards [18, 25, 54] from different vendors and manufactures with unique structures and formats to extract the seed AT grammars. The final challenge is, therefore, to develop a program that can automatically extract the AT commands and their formats as grammars from the documents to automate the design pipeline of ATFuzzer.

C2: (Problematic input representation). The second challenge is to efficiently encode the pattern of problematic inputs. It is crucial, as the problematic AT commands that have similar formats/structures but are not identical may trigger the same behavior. For instance, both ATD123 and ATD1111111111 test inputs are problematic (neither of them is a compliant AT command due to missing a trailing semicolon) and have a similar structure (i.e., ATD followed by a phone number), but are not the same concrete test inputs. While processing these problematic AT commands, one of our test devices, however, stops cellular connectivity. Mutation at the concrete input level would require the fuzzer to try a lot of inputs of the same vulnerable structure before shying away from that abstract input space. This may prevent the fuzzer from testing diverse classes of inputs.

C3: (Syntactic correctness). As shown in Figure 3, most of the AT commands have a specific number and type of arguments. For instance, +CFUN = has two arguments: CFUNarg1 and CFUNarg2. The second challenge is to effectively test this structural behavior and argument types, thoroughly by generating diverse inputs that do not comply with the command structure or the argument types.

C4: (Semantic correctness). Each argument of an AT command may have associated conditions. For instance, $Length(Dnum) \leq n$ in the fifth production rule of Figure 3. Also, arguments may correlate with each other, such as, one argument defines a type on which another argument depends. For instance, +CTFR = refers to a service that causes an incoming alert call to be forwarded to a specified number. It takes four arguments—the first two are **number** and **type**, respectively. Interestingly, the second argument defines the format of the number given as the first argument. If the dialing string includes access code character “+”, then the type should be 145, otherwise, it should be 129. These correlations are prevalent in many AT commands. Hence, the third challenge is to systematically test conditions associated with the arguments of commands to cover both syntactical and semantic bugs.

C5: (Feedback of a test input). The AT interface can be viewed as a black-box providing only limited output of the form: OK (i.e., correctly parsed) or ERROR (i.e., parsing error). The last challenge is to devise a mechanism that can provide information about the code-coverage of the AT interface for the injected test AT command and thus effectively guide us through the fuzzing process.

3.4.2 Insights on Addressing Challenges. For addressing **C1**, we collect the corpus of AT command specifications provided by different standard organizations and vendors. These documents include the complete lists of AT commands. While some of the commands are shared among all the standards and implementations, others are unique to specific vendors. We, therefore, aim at implementing a program to scan and parse the content of the documents based on a regular expression (*regex*) that matches specific sequences of characters. As we are interested in extracting both AT commands and the corresponding parameter structures, we need a *regex* formula to match the specification of the entire command.

For addressing **C2**, we use the extracted grammar itself as the seed of our evolutionary fuzzing framework rather than using a particular instance (i.e., a concrete test input) of the grammar. This is highly effective, as the mutation of a production rule can influence the fuzzer to test a diverse set of inputs. Also, when a problematic grammar is identified, it can serve as abstract evidence of the underlying flaw in the AT interface. Finally, as a grammar can be viewed as a symbolic representation of a set of concrete input AT commands, mutating a grammar can enable the fuzzer to cover large diverse classes of AT commands. The insight here is that testing diverse input classes is likely to uncover diverse types of issues. This enables us to encode a class of bugs through a problematic grammar itself.

To address challenges **C3** and **C4**, at each iteration, ATFuzzer chooses parents with the highest fitness scores and switches parts of the grammar production rules among each other. This causes changes not only to the structural and type information in the child grammars but also forms two very different grammars that try to break the correlation of the arguments. For instance, suppose that ATFuzzer has selected the following two production rules from two different parent grammars: $+CFUN = CFUNarg1, CFUNarg2$ and $+CTFR = \text{number}, \text{type}, \text{subaddr}, \text{satype}$. After applying our proposed grammar crossover mechanisms, the resultant child grammar production rules are: $+CFUN = CFUNarg1, \text{type}, \text{subaddr}, \text{satype}$ and $+CTFR = \text{number } CFUNarg2$. The production rule $+CFUN$ takes only two arguments whereas our new child grammar creates a production rule that has four arguments. The same reasoning also applies to $+CTFR$. Thus, the new grammars with modified production rules would test this structural behavior precisely. Furthermore, $+CTFR$'s first argument **number** is correlated with its second argument **type**. In the modified child grammars, **type**, however, has been replaced with $CFUNarg2$. Recall from our grammar definition, **type** takes argument from the set $\{145, 129\}$, whereas $+CFUNarg2$ takes argument from the set $\{0, 1\}$. Therefore, this single operation completes two tasks at once—it not only tests the correlation among two arguments of $+CTFR$ but also tests conditions of both $+CFUN$ and $+CTFR$. Crossing over grammar production rules creates a drastic modification in the input format, and it aims to explore the diverse portions of the input space to create highly unusual inputs. To test both the structural aspects, we use six very different mutation strategies that create little change to the grammar (compared to crossover) but prove highly effective for checking the robustness of the AT interface.

Finally, for addressing **C5**, we use the precise timing information of injecting an AT command and receiving its output. We keep an upper bound on this time, i.e., a timer ($\mathcal{T}$). If the output is not received within $\mathcal{T}$, we infer that the AT interface has become unresponsive possibly due to the blacklisting mechanism enforced by several vendors. We use this timing information as a *loose-indicator* for code-coverage information. Our intuition is to explore as much of the AT interface as possible. A high execution time loosely indicates that the test command traverses more basic-blocks than the other inputs with lower execution time. We try to leverage this simple positive correlation to design a feedback edge (i.e., a fitness function) of the closed-loop. The timing information, however, cannot help to infer how many new basic-blocks a test input could explore. Since our focus is mainly on baseband-related AT commands, an error in the AT interface has a higher probability of causing disruptions in the baseband, which also trickles down to cellular connectivity. We leverage this key insight and consider both the cellular Internet connectivity information from the target device and the device's debug information (*Logcat*, *dumpsys*, *tombstone*) as an indication of the baseband health after running an AT command. Using this information, we devise our fitness function for guiding ATFuzzer.

Fig. 4. Overview of ATFuzzer framework.

3.4.3 High-level Description of ATFuzzer. ATFuzzer comprises two modules, namely, *evolution module* and *evaluation module*, interacting in a closed-loop (see Figure 4). The evolution module is bootstrapped with a seed AT command grammar that is mutated to generate P_{size} (refers to population size—a parameter to ATFuzzer) different versions of that grammar. Concretely, new grammars are generated from parent grammar(s) by ATFuzzer through the following high-level operations: (1) Population initialization; (2) Parent selection; (3) Grammar crossover; (4) Grammar mutation. Particularly relevant is the operation of *parent selection* in which ATFuzzer uses a *fitness function* to select higher-ranked (parent) grammars for which to apply the crossover/mutation operations (i.e., steps 3 and 4) to generate new grammars. Choosing the higher-ranked grammars to apply mutation is particularly relevant for generating effective grammars in the future.

Evaluating the fitness function requires the evaluation module. For a given grammar g , the evaluation module samples several g -compliant commands to test. It uses the *AT command injector* (as shown in Figures 1 and 4) to send these test commands to the device-under-test. The fitness function uses the individual scores of the concrete g -compliant instances to assign the overall score to g .

4 DETAILED DESIGN OF ATFUZZER

In this section, we discuss our proposed crossover and mutation techniques for the *evolution module* followed by the fitness function design used by the *evaluation module*.

4.1 AT Grammar Extractor

To compile a substantial amount of valid AT commands and their structures, we first collected manuals available online in PDF format [9, 24, 52, 54]. Since these manuals are provided by different vendors, each manual has a unique format, and the lists of AT commands are structured in different ways. We, therefore, apply several heuristics, described in what follows, to effectively extract the AT grammars from the specifications:

H1: (Extraction using AT regular expression:) For most specifications, the definition of AT commands and their grammars are spread over multiple lines and are often not difficult to parse. However, for all these documents, the AT commands start with the prefix “AT.” For these cases, we extract the content of each PDF manual and convert it into *unicode-8* text. Then, we arrange the content of each page on a single line by replacing the *newline* symbol with a space. Once the full content of each page is organized on the same line, we can match AT commands through an appropriate regular expression. We leverage the structure of the general AT commands shown in Figure 5.

An AT command usually starts with the prefix “AT,” with a command extension that can be from a specific set $\{+ \& \# * ..\}$. For instance, for the command `AT + CFUN = 0, 1`, “+” is the extension. The extension can also be null. In general, based on the four types of AT commands, a command can appear in the specifications in any of the following forms:

Command Types	Prefix	Extension	Command Name	Optional Parameter
1. Execution	AT	+	CMD	
2. Read	AT	+	CMD	?
3. Test	AT	+	CMD	= ?
4. Set	AT	+	CMD	=<par ₁ >, <par ₂ >

Fig. 5. General structure of AT commands.

ALGORITHM 1: ATFuzzer

Data: $P_{size}, \mathcal{P}_{pop}, G_{AT}, T_{size}$
Result: G_{best} : Best Grammar

```

1  $P \leftarrow \text{InitializePopulation}(P_{size}, \mathcal{P}_{pop}, G_{AT});$ 
2 while stopping condition is not met do
3   for each grammar  $G_i \in P$  do
4     Generate random input  $I$ 
5     AssesFitness( $G_i, I$ )
6     if Fitness( $G_i$ ) > Fitness( $G_{best}$ ) then
7       |  $G_{best} = G_i$ ;
8     end
9   end
10   $Q = \{\}$ 
11  for  $\frac{P_{size}}{2}$  times do
12     $P_a \leftarrow \text{ParentSelection}(P, T_{size})$ 
13     $P_b \leftarrow \text{ParentSelection}(P, T_{size})$ 
14     $C_a, C_b \leftarrow \text{GrammarBasedCrossover}(P_a, P_b)$ 
15     $Q = Q \cup \{\text{Mutate}(C_a), \text{Mutate}(C_b)\}$ 
16  end
17   $P \leftarrow Q$ 
18 end

```

- Execution command: AT + CMD;
- Read command: AT + CMD?;
- Test command: AT + CMD =?;
- Set command: AT + CMD = [< par₁ > [, < par₂ > [, . . . , < par_n >]]]].

Here, CMD is the command and it can also have some optional parameters. Tian et al. [55] also utilized the general structure of AT commands for extracting AT commands from firmware images. Apart from the structure of AT commands, the specifications also follow some specific structures in defining AT commands. The general structure can be defined as $CMD = [< par_1 > [, < par_2 >]], < par_3 >$. The parameters par_1, par_2, par_3 are always included between angular brackets, whereas the square brackets are used to identify a group of parameters. Even though the structure appears to be very simple, there are several cases in which they do not match the general structure. For instance, the extended command (“+”) can be null, or there is no equal sign, or the elements of the list of parameters occur in different order. Furthermore, due to formatting errors in the conversion from PDF to text, there are also additional empty spaces, or the command itself could be concatenated with other strings irrelevant for the AT command. To address all these issues, we develop a regular expression that matches extended AT commands in the specification PDF files (see Figure 6). The regular expression syntax is from Python.

```

1 (?P<cmd>AT[!@#$%^&*+]?[A-Z\d]+)(?P<arg>(\?|\?|(|\s|\\[<>|+;|,|\s\\w\d
  ]+\\)|[<>|+;|\s\\w\d]+>|,))*)

```

Fig. 6. The regular expression derived in Python for extracting AT command grammars from specifications.

In the regular expression in Figure 6, we can identify the two main groups ($?P < cmd > \dots$) and ($?P < arg > \dots$):

- $< cmd >$ matches the AT commands of the form AT + CMD or ATCMD, that is, it matches execution commands or the first part of the other three types of AT command;
- $< arg >$ is structured as a sequence of OR options to match read (?), test (=?), or set commands. The latter matches any sequence of parameters in angular and square brackets that can be preceded by the equal symbol (e.g., $= [< par_1 > [, < par_2 >]]$, $< par_3 > \dots$). Since each parameter has to be contained within the brackets, we can avoid additional meaningless text that might be concatenated to the command due to formatting errors.

H2: (Extraction from tables): For some specifications [24], the description of each AT command does not include the prefix “AT”. Fortunately, these documents contain the list of AT commands and their parameters in tables. We can, therefore, extract the AT commands described in the document by inspecting the content of the tables in the manual using some keywords with *grep*.

After executing the AT command extractor over the specification documents, we have available the list of AT commands and associated parameters. Each parameter has specific properties according to its type and function within the command. However, since AT commands tend to have similar parameters, we can infer such attributes from the name of the argument. Therefore, from a given entry of the list of AT commands, we can reconstruct the complete grammar that is later used as a seed for ATFuzzer.

4.2 Evolution Module

Given the AT grammar (shown in Figure 3), ATFuzzer’s evolution module randomly selects at most n **cmds** to generate the initial seed AT grammar denoted as G_{AT} . The evolution module yields the grammars G_{best} with the highest scores until a certain stopping condition is met, such as total testing time or number of iterations. Algorithm 1 describes the high-level steps of ATFuzzer’s *evolution module*.

4.2.1 Initialization. The *evolution module* starts with initializing the population P (Line 1 in Algorithm 1) by applying both our proposed crossover and mutation strategies with three parameters: the population size P_{size} ; the probability $\mathcal{P}_{pop}$ of applying crossover and mutation on the grammar; the tournament size T_{size} . The key insight in using $\mathcal{P}_{pop}$ is that it correlates with the number of syntactic and semantic bugs explored. The higher the value of $\mathcal{P}_{pop}$ is, the more diverse the initial population is and vice versa. The more diverse the initial population is, the higher the number of test inputs that check syntactic correctness is and vice versa. Therefore, to explore both syntactic and semantic bugs, we vary the values of $\mathcal{P}_{pop}$, aiming to strike a balance between grammar diversity. To assess the fitness of the initial population P , the evolution module invokes the evaluation module (Lines 3–8) with the generated grammars.

4.2.2 Parent Selection for the Next Round. We use the *tournament selection* technique to get a diverse population at every round. We perform “tournaments” among P grammars (Lines 12–13 in Algorithm 1). The winner of each tournament (the one with the highest fitness score) is then selected for crossover and mutation. In what follows, we discuss in detail our *tournament selection* technique addressing the functional and structural bloating problems of evolutionary fuzzers [56].

Restraining functional bloating. We leverage another insight in selecting grammars at each round of the tournament selection procedure to reduce *functional bloating* [56]—the continuous generation of grammars

Fig. 7. Examples of two-point grammar crossover mechanism.

ALGORITHM 2: Two-Point Grammar Crossover

Data: ParentGrammar P_a , ParentGrammar P_b
Result: P_a, P_b

- 1 Randomly pick production rule R_a from P_a and R_b from P_b
- 2 $R_a \leftarrow R_{a_1}, R_{a_2}, \dots, R_{a_l}$
- 3 $R_b \leftarrow R_{b_1}, R_{b_2}, \dots, R_{b_m}$
- 4 $c_1 \leftarrow$ random integer chosen from 1 to $\min(l, m)$
- 5 $c_2 \leftarrow$ random integer chosen from 1 to $\min(l, m)$
- 6 **if** $c > d$ **then**
- 7 swap c_1 and c_2
- 8 **end**
- 9 **for** i from c to $d - 1$ **do**
- 10 swap grammar rules of R_{a_i}, R_{b_i}
- 11 **end**

containing similar mutated production rules—which adversely affects diverse input generation in evolutionary fuzzing. At each round, we *randomly* select grammars from our population. This is due to the fact that while running an evolutionary fuzzing, the range of fitness values becomes narrow and reduces the search space it focuses on. For example, at any round, if the fuzzer finds a grammar that has a mutated production rule related to `+CFUN` causing an error state in the AT interface, then all the grammars containing this mutated rule will obtain high fitness values. If we then only select parents based on the highest fitness, we would inevitably fall into functional bloating and would narrow down our focused search space with grammars that are somehow associated with this mutated version of `+CFUN` only.

To constraint this behavior, we perform the tournament selection procedure in which we randomly choose T_{size} (where T denotes the set of selected grammars for the tournament and $T_{\text{size}} \leq P_{\text{size}}$) number of grammars from the population P . The key insight of choosing *randomly* is to give chances to the lower fitness grammars in the next round to ensure a diverse pool of candidates with both higher and lower fitness scores.

Restraining structural bloating. After running ATFuzzer for a while, i.e., after a certain number of generations, the average length of individual grammar grows rapidly. This behavior is characterized as *structural bloating*. Referring to the AT grammar in Figure 3, multiple `cmds` (production rules) can contribute to generating the final commands that are sent to the AT command injector for evaluation. These commands can grow indefinitely, but do not induce any structural changes, and thus cause structural bloating. These input commands, therefore, hardly contribute to the effectiveness of the fuzzer. To limit this behavior, we restrict the grammar to have at most three `cmds` at each round to generate the input AT commands for testing.

4.2.3 Grammar Crossover. In the grammar crossover stage, ATFuzzer strives to induce changes in the grammar aiming to systematically break the correlation and structure of the grammar. For this, we take inspiration from traditional genetic programming and apply our custom two-point crossover technique to the grammars.

Fig. 8. Example of all the grammar mutation strategies.

Two-point crossover. ATFuzzer picks up two random production rules from the given parent grammars and generates two random numbers c_1 and c_2 within ℓ where ℓ is the minimum length between the two production rules. ATFuzzer then swaps the fields of the two production rules that are between points c_1 and c_2 .

Figure 7 shows how ATFuzzer performs the two-point crossover operation on production rules $+CTFN =$ and $+CTFR =$ (a subset of the AT grammar in Figure 3) used for controlling the cellular functionalities and for urgent call forwarding, respectively. By applying two-point crossover on $+CFUN = CFUNarg1, CFUNarg2$ and $+CTFR = \text{number, type, subaddr, satype}$, ATFuzzer generates $+CFUN = \text{number, CFUNarg2}$ and $+CTFR = CFUNarg1, \text{type, subaaddr, satype}$, which in turn contribute in generating versatile inputs.

4.2.4 Grammar Mutation. During crossover operation, ATFuzzer constructs grammars that may have diverse structures that are, however, not enough to test the constraints and correlations associated with a command and its arguments. This is due to the fact that AT commands have constraints not only on the fields but also on the commands itself. Therefore, generating versatile grammars that can generate such test inputs is an important aspect of ATFuzzer design. To deal with this pattern, we propose six mutation strategies—*field addition*, *field trimming*, *condition negation*, *type negation*, *fixed integers*, and *connector alteration*. We use $+CTFR = \text{number, type, subaddr, satype}; CFUN = CFUNarg1, CFUNarg2$ (one of the grammars generated from the example seed grammar presented in Figure 3) to illustrate these mutation strategies with examples shown in Figure 8.

Field addition. With our first strategy, we randomly insert/add a field chosen from the production rule of the given grammar at a random location. For instance, applying this mutation strategy (shown in Figure 8(a)) to one of the grammars $+CTFR = \text{number, type, subaddr, satype}; CFUN = CFUNarg1, CFUNarg2$ yields $+CTFR = \text{number, type, satype, subaddr, satype}; CFUN = CFUNarg1, CFUNarg2$ containing an additional argument added after the second argument of the actual grammar. The mutation has changed both the number of arguments and structure of the grammar and, thereby, tests the structure of the grammar.

Field trimming. Our second mutation strategy is to randomly trim an argument from a production rule for the given grammar. Referring to Figure 8(b), applying this to our example grammar, we obtain a production

ALGORITHM 3: Grammar Mutation

Data: Grammar G_a , Tunable parameters : $P_\alpha, P_\beta, P_\gamma, P_\delta, P_\theta, P_\epsilon$
Result: Mutated G_a

- 1 Randomly pick production rule R_a from G_a
- 2 $R_a \leftarrow R_{a_1}, R_{a_2}, \dots, R_{a_l}$
- 3 $c \leftarrow$ random integer chosen from 1 to l
- 4 $P \leftarrow$ Generate random probability from $(0, 1)$
- 5 **if** $P_\alpha \geq P$ **then**
- 6 | trim argument R_{a_c} from production rule R_a
- 7 **end**
- 8 **if** $P_\beta \geq P$ **then**
- 9 | replace $P_c\{\phi\}$ with $P_c\{\neg\phi\}$ in production rule R_a
- 10 **end**
- 11 **if** $P_\gamma \geq P$ **then**
- 12 | $d \leftarrow$ random integer chosen from 1 to l
- 13 | add argument P_{a_c} at position l in production rule R_a
- 14 **end**
- 15 **if** $P_\delta \geq P$ **then**
- 16 | change the type of R_{a_c} from production rule R_a
- 17 **end**
- 18 **if** $P_\theta \geq P$ **then**
- 19 | change R_{a_c} to a fixed value between $\{int.max, int.min, 0, 1, -1\}$
- 20 **end**
- 21 **if** $P_\epsilon \geq P$ **then**
- 22 | change the connector between AT commands to other special characters
- 23 **end**

rule $+CTFR = \text{number, type, subaddr}$; $CFUN = CFUNarg1, CFUNarg2$, which also deviates from the original grammar with respect to both the structure and type.

Condition negation. Our third mutation strategy focuses on the constraints associated with the arguments of a command. Referring to the AT grammar in Figure 3, we encode the constraints with additional conditions (denoted with $\{\dots\}$) in the grammar production rules. With the negation strategy, we randomly pick a production rule of the grammar and choose a random argument that has a condition associated with it. We negate the condition that we use to replace the original one at its original place in the production rule. Figure 8(c) demonstrates how we negate the production rule associated with number used to represent a phone number. The number is a string type with a constraint on its length. We negate this condition with the following three heuristics: (i) Generating strings that are longer than the specified length; (ii) Generating strings that contain not only alphanumeric characters but also special characters; and (iii) Generating an empty string.

Type negation. This mutation strategy aims at effectively testing the type of each argument. The grammars are encoded as having several types, e.g., strings, integers, floats. Following this mutation strategy a random production rule is picked and the type of the rule is changed from its original to any of the other types. Figure 8(d) demonstrates how the rule is applied to the example grammar. In the original grammar, satype is an integer. Following our type negation mutation strategy, the type is changed to string.

Fixed integer. Through this mutation strategy the integer handling of the AT parser is properly tested. Following the strategy, a production rule having type integer is packed and replaced with fixed integers such as: highest

possible positive number, highest possible negative number, zero, -1 , 1 . Our example in Figure 8(e) shows the application on satype, where satype is changed to the maximum possible integer.

Connector alteration. Multiple AT commands are connected with *mathsf*; “semicolon.” To effectively test the conjunction of multiple AT commands, we introduce this mutation strategy. The idea is to replace the connectors with other special characters, e.g., \$, %, *, \$, #. As to restrain structural bloating, we have restricted the grammar to have at most three AT commands. Therefore, we can have at most two connectors and modify one of them randomly. Following our same example, the impact of this mutation strategy is shown in the generated example grammar where the connector is changed from “;” to “#” (Figure 8(f)).

4.3 Evaluation Module

The primary task of the evaluation module is to generate a number of test inputs (i.e., concrete AT command instances) for the grammar received from the evolution module. It then evaluates the test inputs with the *AT command injector* and finally evaluates the grammar itself based on the scores of the generated test inputs. In what follows, we explain how the evaluation module calculates the fitness score of a grammar.

4.3.1 Fitness Evaluation. At the core of ATFuzzer is the fitness function that guides the fuzzing and acts as a liaison for the coverage information. We devise our fitness function based on the timing information and baseband-related information of the smartphone. Our fitness function comprises two parts: (1) Fitness score of the test inputs generated from a grammar; (2) Fitness score of the grammar in the population.

Fitness score of the test inputs of a grammar. The fitness evaluator of ATFuzzer generates N inputs from each grammar and calculates the score for each input. We define this fitness function for an input AT command instance x as:

$$\text{fitness}(x) = \alpha \times \text{timing}_{\text{score}} + (1 - \alpha) \times \text{disruption}_{\text{score}},$$

where α is a tunable-parameter that controls the impact of $\text{timing}_{\text{score}}$ and $\text{disruption}_{\text{score}}$. Let t_x be the time required for executing an AT command x ($0 \leq x < N$) on the smartphone under test. *Execution time* of an AT command is defined as the time between when the AT command is sent and when the output is received by the AT command injector. Note that we normalize the execution time by the input length.

Let $t_1, t_2, \dots, t_N$ be the time for executing N AT commands, we define the timing score for instance x in a population of size N as follows:

$$\text{timing}_{\text{score}} = \frac{t_i}{t_1 + t_2 + \dots + t_N}.$$

Note that while running AT commands over Bluetooth, the commands and their responses are transmitted in over-the-air (OTA) Bluetooth packets. To compute the precise execution time of the AT command on the smartphone, we take off the transmission and reception times from the total running time. Also, to make sure Bluetooth signal strength change does not interfere with the timing information, our system keeps track of the RSSI (Received Signal Strength Indication) value and carries out the fuzzing at a constant RSSI value.

We define $\text{disruption}_{\text{score}}$ based on the following four types of disruption events: (i) Complete shutdown of SIM card connectivity; (ii) Complete shutdown of cellular Internet connectivity; (iii) Partial disruption in cellular Internet connectivity; (iv) Partial disruption of SIM card connectivity with the phone. For cases (i) and (ii), complete shutdown causes denial of cellular/SIM functionality, recovery from which requires rebooting the device. ATFuzzer uses **adb reboot** command that takes $\sim 15 - 20$ seconds to restart the device without entailing any manual intervention. On the contrary, partial shutdown for the cases (iii) and (iv) induces denial of cellular/SIM functionality for $\sim 3 - 5$ seconds and thus does not call for rebooting the device to recuperate back to its normal state. These events are detected and monitored using the open-source tools available to us from Android, e.g., *logcat*, *dumpsys*, and *tombstone*. When injecting the AT commands, we use these tools to detect the events at runtime. We take into account if there is a crash in the baseband or the RIL daemon. We assign a score between

0 – 1 to a disruption event in which 0 denotes *no disruption at all* (i.e., the device is completely functional) with no adverse effects and 1 denotes complete disruption of the cellular or SIM card connectivities.

Fitness score of a grammar. After computing the fitness scores for all the concrete input instances, we calculate the grammar’s score by taking the average over all instance scores.

5 EVALUATION

Our primary goal in this section is to evaluate the effectiveness of ATFuzzer by following the best possible practices [7, 57] and guidelines [27]. We, therefore, first discuss the experiment setup and evaluation criteria and then evaluate the efficacy of our prototype against the widely used AFL [62] fuzzer—customized for our context.

5.1 Experiment Setup

ATFuzzer setup. We implemented ATFuzzer with ~4000 lines of Python code. We encoded the grammars (with JSON) for a corpus of 90 baseband-related AT commands following the specification in the 3GPP [24] documentation and extracting some of the vendor-specific AT commands following the work of Tian et al. [55]. During its initialization, ATFuzzer receives the name of the AT command as input, retrieves the corresponding grammar that will be used as the seed (G_{AT} in Algorithm 1) from the file, generates the initial grammar population, and realizes the proposed crossover and mutation strategies. Hence, our approach is general and easily adaptable to other structured inputs, since it is not bound to any specific grammar structure. Since testing a concrete AT command instance requires 15–20 seconds on average (because of checking the cellular and SIM card connectivity after executing a command and for rebooting the device in case of AT interface’s unresponsiveness for blacklisting), we set P_{size} to 10, which we found through empirical study to be the most suitable in terms of ATFuzzer’s stopping condition. Following the same procedure, we test 10 concrete AT commands in each round for a given grammar. We set the probability P_{pop} to 0.5 to ensure uniform distribution in the grammar varying ratio.

Conceptually, one can argue for testing at a “batch” mode to chop the average time for fuzzing an AT grammar. For instance, injecting 10 AT commands together and then checking the cellular and SIM connectivity at once. Although this design philosophy is intuitive, it fails to serve our purpose. The reason is that this approach may be able to detect permanent disruptions but it is unable to detect temporary disruptions to cellular or SIM connectivity. For instance, even if the second AT command in the batch induces a temporary disruption, there will be no trace of disruptions at all by the time when the 10th (i.e., the last) AT command is executed.

Target devices configuration. We tested 10 different devices (listed in Table 1) from 6 different vendors running 6 different Android versions to diversify our analysis. For Bluetooth, we do not require any configuration on the phone, whereas for running AT commands over USB, some phones require specific set up procedures.

The first step for all the Android devices we tested is to enable the USB debugging option in the Developer options menu, which is normally hidden. To unlock this menu, it is necessary to tap seven times the Build number entry in the device settings, enable the Developer options, and finally enable USB debugging. The USB debugging feature (among other features in the Developer menu) aims at providing direct access to the device’s system, functionalities for debugging operations, access to the file system, and so on. It is important to ensure that the USB configuration is set on MPT (Media Transfer Protocol). For additional details see Reference [1].

The second step is to configure the USB interface within the device’s system. Some of the devices we tested expose their modem functionality by default and therefore required no additional configuration (also listed in Table 1). However, for the devices that do not expose any modem, it was necessary to root them and set a specific type of USB configuration. The USB configuration can be changed by setting `sys.usb.config` property. Thereby, the devices can be accessed through ADB (Android Debug Bridge) and Fastboot tools. With ADB it is possible to access the device’s file system, reboot it in different modes, such as bootloader mode, rooting it, and finally change the device’s properties directly with the command `setprop <property-name> <value>`. Fastboot,

Table 1. List of the Devices We Tested, with Software Information, USB Configuration Required, and the Operating System We Used to Fuzz Each Device

Device	Android Version	Build Number	Baseband Vendor	Baseband	USB Config	OS	Interface
Samsung Note2	4.3	JSS15J. I9300XU GND5	Samsung Exynos 4412	N7100DD UFND1	None	Linux	Bluetooth and USB
Samsung Galaxy S3	4.3	JSS15J. I9300XX UGND5	Samsung Exynos 4412	I9300XX UGNA8	None	Linux	Bluetooth and USB
LG G3	6.0	MRA58K	Qualcomm Snapdragon 801	MPSS.DI.2.0.1. c1.13-00114 -M8974AA AAANPZM- 1.43646.2	None	Linux	Bluetooth and USB
HTC Desire 10 lifestyle	6.0.1	1.00.600.1 8.0_g CL800193 release-keys	Qualcomm Snapdragon 400	3.0.U205591 @60906G_01. 00.U0000. 00_F	sys.usb.config mtp,adb,diag, modem, modem_mdm, diag_mdm	Windows	Bluetooth and USB
LG Nexus 5	5.1.1	LMY48I	Qualcomm Snapdragon 800	M8974A-2.0.50.2.26	sys.usb.config diag,adb	Linux	Bluetooth and USB
Motorola Nexus 6	6.0.1	MOB30M	Qualcomm Snapdragon 805	MDM9625_ 104662.22. 05.34R	fastboot oem bp-tools-on	Windows	Bluetooth and USB
Huawei Nexus 6P	6.0	MDA89D	Qualcomm Snapdragon 810	.2.6.1.c4- 00004-M899 4FAAAN AZM-1	fastboot oem enable-bp-tools	Windows	Bluetooth and USB
Samsung Galaxy S8+	8.0.0	R16NW.G95 5USQU5CRG3	Qualcomm Snapdragon 835	G955US QU5CRG3	None	Linux	Bluetooth and USB
Huawei P8 Lite ALE-L21	5.0.1	ALE-L21CO2B140	HiSilicon Kirin 620 (28 nm)	22.126.12.00.00	None	None	Bluetooth
Pixel 2	8.0.0	OPD3.1708 16.012	Qualcomm MSM8998 Snapdragon 835	g8998-00122-1708231715	None	None	Bluetooth

instead, allows operating the device in boot loader mode, installs new partitions, and changes pre-boot settings required for rooting.

Few devices are required specific configurations. For instance, for LG Nexus 5, we have to set `sys.usb.config` from the default “mtp,adb” to “diag,adb” through `adb shell`. This setting allows accessing the phone in diagnostic mode and therefore to communicate with the AT command interface. For Motorola Nexus 6 and Huawei Nexus 6P, the USB configuration can be changed by rebooting the phone in boot loader mode by executing the command “adb reboot fastboot.” Then, issue the command “fastboot oem bp-tools-on” and “fastboot oem enable-bp-tools” for Nexus 6 and Nexus 6P, respectively, as reported in Reference [12]. The command sets the property `sys.usb.config` to `diag, serial_smd, rmnet_ipa, adb` allowing establishing a serial communication with the device and interacting with it through the AT interface.

Note that running ATFuzzer using Windows operating system might require additional settings and to install/reinstall some of the drivers. For instance, when connecting Nexus 6 and Nexus 6P after setting the phone USB properties, Windows should detect the device and correctly set the USB drivers. It is possible to check the setting through the hardware panel where the entries “Android Composite ADB Interface,” “Qualcomm HS-USB Diagnostics,” “Qualcomm HS-USB Modem,” “Qualcomm Wireless HS-USB Ethernet Adapter,” and “USB Composite Device” should be listed. The ADB interface is specifically needed to perform ADB operations, whereas the Qualcomm modem is necessary to use AT Commands. If the Qualcomm modem is not listed, one problem could be the Windows drivers. If that is the case, it is necessary to uninstall the drivers for all the Nexus 6 (Nexus 6P) devices listed (more than one could be listed), reconnect the phone to the computer, and wait for the drivers to be automatically installed.

Table 2. Summary of ATFuzzer 's Bluetooth Parser Findings

Class of Bugs	Grammar and Command Instance	action/implication	Nexus5	LG G3	Nexus6	Nexus6P	HTC	S8plus	S3	Note2	Huawei P8lite	Pixel 2
Syntactic -returns OK with action	<code>cmd → D.Dnum.Darg1.Darg2</code> <code>Dnum → [A - Z0 - 9 + #]*</code> <code>Darg1 → I G ε</code> <code>Darg2 → ;</code> <code>Darg3</code> <code>Darg3 → [A, B, C]*</code> E.g., ATD + 46420480341; AB; C	crash/internet connectivity disruption	✓									✓
	<code>cmd → D.Dnum.Darg1.Darg2</code> <code>Dnum → [A - Z0 - 9 + #]*</code> <code>Darg1 → I G ε</code> <code>Darg2 → ;</code> <code>Darg3</code> <code>Darg3 → [A, B, C]*</code> E.g., ATD + 46420480341; AB; C	crash/downgrade			✓	✓						
	<code>cmd → +CIMI.Arg1</code> <code>Arg1 → [a - zA - Z0 - 9 + #]*</code> E.g., AT + CIMI;;; abc	read/IMSI leak						✓	✓	✓		
	<code>cmd → +CGSN.Arg1 Arg1</code> <code>→ [a - zA - Z0 - 9 + #]*</code> E.g., AT + CGSN;;; abc##	read/IMEI leak						✓	✓	✓		
Correctly formatted command	<code>cmd → +CIND?</code>	read/leaks call status, call setup stage, internet service status, signal strength, current roaming status, battery level, call held status	✓	✓	✓	✓		✓	✓	✓	✓	✓
	<code>cmd → +CHUP</code>	execution (cutting phone calls)/ DoS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	<code>cmd → Arg.D.Dnum.Darg;</code> <code>Arg → [a - zA - Z]</code> <code>Dnum → [a - zA - Z0 - 9 + #]*</code> <code>Darg → I G ε</code> E.g., ATD * 61 * +1812555673 * 11 * 25#;	execution/ call forwarding, activating do not disturb mode, activating selective call blocking	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

5.2 Evaluation Criteria

ATFuzzer has three major components—grammar crossover, mutation, and feedback loop—to effectively test a target device. We, therefore, aim to answer the following research questions to evaluate ATFuzzer:

- **RQ1:** How proficient is ATFuzzer in finding erroneous AT commands over Bluetooth?
- **RQ2:** How proficient is ATFuzzer in finding erroneous AT commands over USB?
- **RQ3:** How effective is our grammar-aware crossover strategy?
- **RQ4:** How effective are our grammar-aware mutation strategies?
- **RQ5:** When using grammars, how much does the use of timing feedback improve fuzzing performance?
- **RQ6:** Is ATFuzzer more efficient than other state-of-the-art fuzzers for testing AT interface?
- **RQ7:** How effective is the AT grammar extraction module of ATFuzzer?

To tackle **RQ1–RQ2**, we let our ATFuzzer run over USB and Bluetooth each for one month to test the 10 different smartphones listed in Table 1. ATFuzzer has been able to uncover a total of 4 erroneous AT grammars inducing a crash, downgrade, and information leakage over Bluetooth and 14 erroneous AT grammars over USB. Based on the type of actions and responses to the problematic AT command instances, we initially categorize our results as syntactic- and semantic-problematic AT grammars and further categorize the syntactically problematic grammars into three separate classes: (i) responds ok with composite actions; (ii) responds ok with an action; (iii) responds error with an action. Here, an action means either a crash (i.e., any disruption event defined in Section 4), or leakage of any sensitive information, or the execution of a command (e.g., hanging up a phone call).

We summarize ATFuzzer 's findings for Bluetooth in Table 2 and for USB in Table 3. To answer the research questions **RQ3–RQ5**, we evaluate ATFuzzer by disabling one of its components at a time. We create three new

Table 3. Summary of ATFuzzer's Findings over USB

Class of Bugs	Grammar and Command Instance	action/implication	Nexus5	LG G3	Nexus6	Nexus6P	HTC	S8plus
Syntactic—returns OK with composite actions	$\text{cmd} \rightarrow \text{I.Arg.D.Dnum.Darg};$ $\text{Arg} \rightarrow [a - zA - Z]$ $\text{Dnum} \rightarrow [a - zA - Z0 - 9 + \#]^*$ $\text{Darg} \rightarrow I G \epsilon$ E.g., AT+IHD + 4642048034I;	read, execution/leaks manufacturer, model, revision, and IMEI	✓	✓	✓	✓	✓	
	$\text{cmd} \rightarrow +\text{COPN}; \text{Arg}$ $\text{Arg} \rightarrow [i I]^*$ E.g., AT + COPN; III	read/leaks list of operators, manufacturer, model, revision, and IMEI	✓	✓			✓	✓
Syntactic—returns OK with an action	$\text{cmd} \rightarrow \text{Arg1.I.Arg1}$ $\rightarrow [X H]$ E.g., ATHIX	read/leaks manufacturer, model, revision, and IMEI	✓	✓	✓	✓	✓	
	$\text{cmd} \rightarrow \text{Arg1.I.Arg1.Arg1}$ $\text{Arg1} \rightarrow X$ E.g., ATXIX	read/leaks manufacturer, model, revision, and IMEI				✓	✓	✓
	$\text{cmd} \rightarrow \text{Arg1.Arg2.Arg3}$ $\text{Arg1} \rightarrow +\text{CIMI} I +\text{CEER}$ $\text{Arg2} \rightarrow * ;$ $\text{Arg3} \rightarrow Q Z$ E.g., AT + CIMI * Q	read/leaks IMSI, manufacturer, model, revision, and IMEI	✓	✓				
	$\text{cmd} \rightarrow +\text{CLCC}; \text{Arg1}$ $\text{Arg1} \rightarrow [a - zA - Z0 - 9]^*$ E.g., AT + CLCC; ABC123	read/leaks current call list	✓	✓		✓	✓	✓
	$\text{cmd} \rightarrow \text{Arg1}; \text{Arg2}$ $\text{Arg1} \rightarrow +\text{COPN} +\text{CGMI} $ $+ \text{CGMM} + \text{CGMR}$ $\text{Arg2} \rightarrow [X E]$ E.g., AT + COPN; X	read/leaks list of operators, IMEI, model, and revision information of the device	✓	✓	✓	✓	✓	✓
Syntactic—returns ERROR with an action	$\text{cmd} \rightarrow \text{Arg1.Arg2.Arg3}$ $\text{Arg1} \rightarrow +\text{CIMI} I +\text{CEER}$ $\text{Arg2} \rightarrow * ;$ $\text{Arg3} \rightarrow [Q Z]$ E.g., AT; L	read/leaks IMSI, manufacturer, model, revision, and IMEI	✓	✓	✓		✓	
	$\text{cmd} \rightarrow \text{Arg1}; \text{Arg2}$ $\text{Arg1} \rightarrow +\text{COPN} +\text{CGMI} $ $+ \text{CGMM} + \text{CGMR}$ $\text{Arg2} \rightarrow [X E]$ E.g., AT + CGMM; O	read/leaks list of operators, IMEI, model, and revision information of the device	✓	✓	✓	✓		
	$\text{cmd} \rightarrow \text{Arg.D.Dnum.Darg};$ $\text{Arg} \rightarrow [a - zA - Z]$ $\text{Dnum} \rightarrow [a - zA - Z0 - 9 + \#]^*$ $\text{Darg} \rightarrow I G \epsilon$ E.g., ATMD + 4632048034	crash/internet connectivity disruption	✓					
	$\text{cmd} \rightarrow \text{D.Dnum.Darg}; \text{Arg}$ $\text{Arg} \rightarrow [a - zA - Z\&\# \{ \}]^*$ $\text{Dnum} \rightarrow [a - zA - Z0 - 9 + \#]^*$ $\text{Darg} \rightarrow I G \epsilon$ E.g., ATD + 4632048034; l#	crash/internet connectivity disruption	✓		✓			
	$\text{cmd} \rightarrow +\text{CUSD} =, \text{String}$ $\text{String} \rightarrow [a - zA - Z0 - 9 + \#]^*$ E.g., AT + CUSD =, ABC	crash/internet connection disruption			✓	✓		

(Continued)

Table 3. Continued

Class of Bugs	Grammar and Command Instance	action/implication	Nexus5	LG G3	Nexus6	Nexus6P	HTC	S8plus
Semantic—returns OK with an action	$\text{cmd} \rightarrow +\text{CCFC}=\text{Arg1},\text{Arg2},\text{Arg3},$ $145,32,\text{Arg4},13,27$ $\text{Arg1} \rightarrow [1-5]$ $\text{Arg2} \rightarrow [1-2]$ $\text{Arg3} \rightarrow [0-9]^*$ $\text{Arg4} \rightarrow [a-zA-z0-9]^*$ E.g., AT + CCFC = 3, 2, 732235, 145, 32, cA4{NYv, 13, 27	crash/internet connectivity disruption	✓		✓	✓		
	$\text{cmd} \rightarrow +\text{COPS} = 0,\text{Arg1},\text{Arg2},2$ $\text{Arg1} \rightarrow [0 1]$ $\text{Arg2} \rightarrow [a-zA-z]^*$ E.g., AT + COPS = 0, 1, c19v6fC, 2	crash/internet connectivity disruption		✓				

instances of ATFuzzer: ATFuzzer without crossover, ATFuzzer without mutation, and ATFuzzer without fitness evaluation. In what follows, we evaluate these three variants with the AT grammar (in Figure 3) and compare their efficacy in discovering bugs against original ATFuzzer with all capabilities switched on. Moreover, to answer the research question **RQ6**, we create our variation of AFL (American Fuzzy Lop). To perform a fair comparison, we run all our experiments on Nexus 5 for each variation of ATFuzzer and our version of AFL for three days. To tackle research question **RQ7** and test the effectiveness of the AT grammar extraction module, we apply it to four AT command standard PDF's and compare it to our previous manual AT grammar extraction.

5.3 Findings over Bluetooth (RQ1)

Unlike USB, Bluetooth does not require any pre-processing or configuration to the phone to execute AT commands. Besides this, over-the-air Bluetooth communications are inherently vulnerable to MitM attacks [2, 32, 51]. All these enable the adversary to readily exploit the vulnerabilities over Bluetooth with sophisticated attacks.

5.3.1 Results. We first discuss the results that relate to invalid AT commands and then we discuss the attacks and impacts of both invalid and valid AT commands.

(1) Syntactic errors – responds ok with actions. ATFuzzer uncovered four problematic grammars in these categories in seven different Android smartphones. We observe that the target device responds to the invalid AT command and also performs an action. For instance, ATFuzzer found a specific variant of ATD grammar ATDA; A; B in Nexus 5, which is syntactically incorrect, but returns OK and makes the cellular Internet connectivity temporarily unavailable. Beside this, the concrete instances of the same grammar also downgrade the cellular connectivity from 4G to the 3G/2G in Nexus 6 and Nexus 6P smartphones, thus entails severe security and privacy impacts.

5.3.2 Attacks with Invalid AT Commands. We now present three practical attacks that can be carried out using the invalid grammars uncovered through ATFuzzer.

Denial of service. The adversary using a malicious Bluetooth peripheral device (e.g., Bluetooth headphone with only call audio and media permissions) or a MitM instance may exploit the invalid AT command, e.g., ATDB; A; B, and temporarily disrupt the Internet connectivity of the Pixel 2 and Nexus 5 phones. To cause long-term disruptions in Internet connectivity, the adversary may inject this command intermittently and thus prevent the user from accessing the Internet. Note that there is no valid AT command that controls the Internet connectivity over Bluetooth and thus it is not possible to achieve this impact using a valid AT command.

Downgrade. The same invalid grammar (shown in Table 2) exploited in the previous DoS attack in Nexus 5 phone can also be exploited to downgrade the cellular connectivity on Nexus 6 and Nexus 6P phones. Similar to

the previous DoS attack, such downgrade of cellular connectivity is not possible with any valid AT commands running over Bluetooth. Downgrade (also known as *bidding-down*) attacks have catastrophic implications, as they open the avenue to perform over-the-air man-in-the-middle attacks in cellular networks [6, 34].

IMSI & IMEI catching. ATFuzzer uncovered the invalid variations (AT + CIMI; ; ; ; abc and AT + CGSN123df) of two valid AT commands (+CIMI and +CGSN), which enable the adversary to illegitimately capture the IMSI and IMEI of a victim device over Bluetooth. Exploiting this, any Bluetooth peripheral connected to the smart-phone can surreptitiously steal such important personal information of the device and the network. We have successfully validated this attack in Samsung Galaxy S3, Samsung Note 2, and Samsung Galaxy S8+. One thing to be noted here is that, after manual testing, we found out that the valid versions of these two commands also leak IMSI and IMEI. We argue that even if there is a blacklist/firewall policy put into place to stop the leakage through valid AT commands, it will not be sufficient, because it will leave the scope to use the invalid versions of the command (that ATFuzzer uncovered) to expose this sensitive information.

The impact of this attack is particularly more fatal than that of the previous two attacks. This is because the illegitimate exposure of IMSI and IMEI through Bluetooth provides an edge to the adversary to further track the location of the user or intercept phone calls and SMS using fake base stations [15, 16] or MitM relays [50]. Samsung has already acknowledged the vulnerabilities and is working on issuing patches to the affected devices. We also summarize the findings of ATFuzzer in Table 2. CVE-2019-16401 [41] has been assigned to this vulnerability along with other sensitive information leakage for the affected Samsung devices.

5.3.3 Attacks with Valid AT Commands. We summarize ATFuzzer's other findings, in which we demonstrate that the exposed AT interface over Bluetooth allows the adversary to run valid AT commands to attain malicious goals that may negatively affect a device's expected operations. The results are particularly interesting, as Bluetooth interface has not yet been *systematically* examined.

Information leak. The adversary can use a valid AT command to learn the whole set of private information about the phone. The malicious Bluetooth peripheral device can get the call status, call setup state, Internet service status, the signal strength of the cellular network, current roaming status, battery level, and call hold status for the phone using this valid AT commands.

DoS attacks. A malicious peripheral can exploit the AT + CHUP command to prevent the victim device from receiving any incoming phone call. From the previous *information leakage* (e.g., call status) attack, an attacker can probe periodically to detect whether there is a phone call or not. Whenever he detects there is a phone call, the attacker injects AT + CHUP to cut the phone call. To make the matters worse, the attack is transparent to the victim, i.e., there is no indication on the mobile screen that an attack is going on. The victim device user perceives either there is no incoming call or abrupt call drops due to poor signal quality or network congestions. CVE-2019-16400 [40] has been assigned for this along with other reported denial of service attacks in Samsung phones.

Call forwarding. If the victim device is subscribed to call forwarding service, the adversary may exploit the ATD command to forward victim device's incoming calls to an attacker-controlled device. Exploiting this, the adversary first prevents the victim device from receiving the incoming calls and then learns sensitive information, such as password or PIN for two-factor authentication possibly sent by an automated teller. Note that such call forwarding is also transparent to the user, since the user is unaware of any incoming calls.

Activating *do not disturb* mode. The adversary using a malicious Bluetooth peripheral or MitM instance can turn on the *do not disturb* mode of the carrier through ATD command. Similar to call forwarding attack, it is also completely transparent to the user, as no visible indication of *do not disturb* mode is displayed on the device. While the user observes all the network status bars and the Internet connectivity, the device, however, does not receive any call from the network.

Selective call blocking. A variation of the previous attack is also possible in which the adversary may allow the victim phone to receive selective calls by intermittently turning on/off the *do not disturb* mode. This may force the user to receive calls only from selective users not affecting others.

5.4 Findings over USB (RQ2)

We now discuss findings over USB.

(1) Syntactic errors – responds ok with composite actions. It is one of the interesting classes of problematic grammars for which the AT interface of the affected devices respond to invalid AT commands with ok, but performs multiple actions together. These invalid commands are compositions of invalid characters and two valid AT commands with no semicolon as their separator. For instance, ATFuzzer generated an invalid command `ATIHD + 4632048034`; using two valid grammars for `ATD` and `ATI` (as shown in Figure 3) and invalid characters for which the target device returns ok but places a phone call to 4632048034 and shows the manufacturer, model revision, and IMEI information simultaneously.

(2) Syntactic errors – responds ok with an action. In this type of syntactically problematic grammar, the target device responds to an invalid command instance with ok but performs an action. For instance, the grammar `cmd → Arg1. I. Arg2` in Table 3 can be instantiated with an invalid command instance `ATHIX`, which returns sensitive device information.

(3) Syntactic errors – responds error with an action. In this class of syntactic error, the AT interface recognizes the inputs as faulty by acknowledging with error, but it still executes the action associated with the command and even does worse by crashing the RIL daemon and inducing complete disruptions in the cellular Internet connectivity. It basically reveals a fundamental flaw in the AT interface—*if a command is considered as erroneous, it should not be executed*. For example, the grammar `cmd → D. Dnum` in Table 3 can be instantiated with `ATD + 4632048034` (a variation of the `ATD` production rule in Figure 3), which is supposed to start a cellular voice call. Instead, the grammar returns error in the form of `NO CARRIER` and induces the cellular Internet connectivity to go down completely for a certain amount of time (15–20 seconds). We have also found grammars for which the device returns other error statuses, e.g., `ERROR`, `NO CARRIER`, `CME ERROR`, `ABORTED`, and `NOT SUPPORTED`, but still executes those invalid commands.

(4) Semantic errors. This class of grammars conforms with the input pattern defined by the standards [24], but induces disruptions in the cellular connectivity for which the recovery requires rebooting the device. The grammars of this class are shown in Table 3.

Possible exploitation. It may appear that the implications of invalid AT commands over USB are negligible as compared to the valid AT commands, which may wreak havoc by taking full control of the device. We, however, argue that if AT interface exposure is restricted through blacklisting the critical and unsafe valid AT commands by the parser in the first place, the adversary will still be able to induce the device to perform same semantic functionalities using invalid AT commands. This is due to the uncovered vulnerabilities for which the parser will fail to identify the invalid AT commands as the blacklisted commands and thus allows the adversary to achieve same functionalities as the valid ones.

5.4.1 Efficacy of Grammar-aware Crossover (RQ3). ATFuzzer without crossover (by disabling the crossover in ATFuzzer) uncovered only three problematic grammars as compared to ATFuzzer with all proposed crossover and mutations (Table 4). This is due to the fact that ATFuzzer without crossover cannot induce enough changes in the structure and type of the arguments of parent grammars, as a result of which it reduces the search space.

5.4.2 Efficacy of Grammar-aware Mutation (RQ4). Since ATFuzzer without mutation cannot induce changes in the arguments and the respective conditions, it uncovered only two problematic grammars. ATFuzzer without crossover, however, performs slightly better than that of the ATFuzzer without crossover. This also justifies our

Table 4. Result Obtained with Different Fuzzing Approaches on Nexus5 over a Period of Three Days for Each Approach

Fuzzing Approach	Problematic Grammars
ATFuzzer	9
ATFuzzer w/o feedback	5
ATFuzzer w/o crossover	3
ATFuzzer w/o mutation	2
Modified AFL	2

Table 5. Comparison between Manual and Automated AT Grammar Extraction

Approach	Commands Extracted	False Positives
Manual extraction	75	0%
AT grammar extraction	300>	<2%

intuition that mutation strategies play a vital role in any fuzzer as compared to crossover techniques. Without mutation, a fuzzer unlikely generates interesting inputs for the system under test.

5.4.3 Efficacy of Timing Feedback (RQ5). We observed that ATFuzzer without feedback performs better than the other two (RQ2 and RQ3) variations. ATFuzzer without feedback uncovered five problematic grammars and thus is less effective than ATFuzzer with feedback. AT interface being a complete black box with little to no feedback, we had to resort to various creative ways including timing information to generate feedback score. However, this resorts to an upper bound for the coverage information and loosely dictates ATFuzzer.

5.4.4 Comparison with Other State-of-the-art Fuzzers (RQ6). We compare the effectiveness of ATFuzzer against AFL [62]. Since current versions of AFL require instrumenting the test programs, we implemented a modest string fuzzer that adopts five mutation strategies (*walking bit flips*, *walking byte flips*, *known integers*, *block deletion*, and *block swapping*) employed by AFL and incorporated our proposed timing-based feedback loop to it. We evaluate this AFL variant with 80 different seeds (consisting of valid and invalid command instances of 40 different randomly chosen AT reference grammars).

Table 4 shows that the AFL variant uncovered two different problematic grammars, whereas ATFuzzer uncovers nine unique grammars after running for three days. Though we decided to compare our tool with AFL, which is the best choice we had, as AFL is considered the state-of-the-art tool for fuzzing, we do not claim the comparison to be ideal. Because AFL relies heavily on code coverage information and for our case, we replaced the code coverage with the best available substitute (i.e., coarse-grained timing information as a loose indicator to code coverage). We acknowledge that this is a best-effort approach and the evaluation may be sub-optimal.

5.4.5 AT Grammar Extraction Effectiveness (RQ7). Compared to our previous manual AT grammar extraction procedure of 75 command grammars, our automated extraction procedure with two heuristics performs significantly better. We could extract 129 grammars from the 3GPP specification [24] and in the order of 300 different grammars from the other documents (each specification defines vendor-specific commands, so the total number can vary). Moreover, we want to report that through manual inspection, we encountered less than 2% of false-positive commands on average. The comparison between the manual- and our heuristic-based command extraction is summarized in Table 5.

6 RELATED WORK

In this section, we mainly discuss the relevant work on the following four topics: AT commands, mutation-based fuzzing, grammar-based mutation, grammar-based generation.

AT commands. Most of the following previous work related to AT commands investigate how an adversary can misuse valid AT commands to attack various systems. The work from Tian et al. [55] can be considered the most relevant to our work, however, it is significantly different in the following three aspects: (i) First, they only show the impact of AT commands over USB, as they consider the functionality and scope of AT commands over Bluetooth too limited to study. We, however, demonstrate the dire consequences of AT commands over Bluetooth interface with the uncovered invalid and valid AT commands. (ii) Second, they only show the impact of valid AT commands, whereas we demonstrate the impact of invalid AT commands exploring different attack surfaces. (iii) Finally, one of the primary objectives of our work is to test the robustness of the AT interface, which is a different and complementary end objective than theirs.

BlueBug [5] exploits a Bluetooth security loophole on few Bluetooth-enabled cell phones and issues AT commands via a covert channel. It, however, relies on the Bluetooth security loophole to attack and does not apply to all phones. In contrast, we have demonstrated a variety of attacks using valid and invalid AT commands running over Bluetooth that do not rely on any specific Bluetooth assumptions and also are applicable to all the modern smartphones we had in our corpus. Injecting AT commands on Android baseband was previously discussed on the XDA forum [4]. Pereira et al. [42, 44] used AT commands to flash malicious images on Samsung phones. Hay [11] discovered that AT interface can be exploited from Android bootloader and discovered new commands and attacks using the AT interface. AT commands have been used to exploit modems other than smartphones as well. Most prominently, USBswitcher [43, 48] and Reference [42] demonstrate how these commands expose actions potentially causing security vulnerabilities in smartphones. Some other work use AT commands as a part of their tool; for instance, Mulliner et al. [36] use the AT commands as feedback while fuzzing SMS of phones. Xenakis et al. [60, 61] devise a tool using AT commands to steal sensitive information from baseband. None of them, however, actually analyzes or discovers bugs in the AT parser itself.

Mutation-based fuzzers. Initial mutation-based fuzzers [35] used to mutate the test inputs randomly. To make this type of fuzzer more effective, a huge amount of work has been carried out to develop sophisticated techniques to improve mutation strategies—coverage information through instrumenting the binary [10, 29, 30, 62]; resource usage information [28, 45]; control and data flow features [47]; static vulnerability prediction models [31]; data-driven seed generation [57]; high-level structural representation of seed file [46]. There are also a few mutation-based fuzzers that incorporate the idea of grammars rather than inputs. Wang et al. [58] use grammars to guide mutation, whereas Aschermann et al. [7] rely on code coverage feedback. Simulated annealing power schedule with genetic fuzzing has also been incorporated in Reference [8]. However, due to the black-box nature of our system and structural pattern of AT command inputs, none of the existing concepts suffice fuzzing AT parser.

Generation-based fuzzers. Generation-based fuzzers generate inputs based on a model [13, 14, 53, 56], specification, or defined grammars. However, to the best of our knowledge, no fuzzer discovers a class of bugs at the grammar-level, rather generates concrete input instances. There are also some generation-based, more precisely, defined grammar-based fuzzers [49, 63] that use manually specified grammars as inputs. For instance, Mangleme is an automated broken HTML generator and fuzzer, and Jsfunfuzz [49] uses specific knowledge about past and present vulnerabilities and uses grammar rules to produce inputs that may cause problems. Both of them are, however, random fuzzers.

The current article is an extended version of our previous paper [26]. With respect to the previous paper, the current article has several major extensions. The first is the design and implementation of an AT grammar extractor that automatically retrieves AT command grammars from specification PDFs available online. By integrating such an extractor with ATFuzzer, we have been able to obtain a considerably larger set of AT

commands grammar (more than 300 compared to the 75 commands we manually extracted in the previous work), which enables more comprehensive tests. We have also extended ATFuzzer's fuzzing mutation module introducing additional mutation strategies. We have evaluated the effectiveness of the new fuzzing mutation module. We notice that the enhanced ATFuzzer is able to perform a more refined fuzzing, which allows us to mutate input AT command grammars more precisely compared with the version of ATFuzzer described in our previous paper. Therefore, we have been able to uncover a new AT command grammar that causes device crash and connectivity disruption. In addition to such major extensions, we have included in the current article more accurate details about the experiment setup and the required device configurations.

7 DISCUSSION

Defenses. Our findings show that current implementations of baseband processors and AT command interfaces fail to correctly parse and filter out some of the possible anomalous inputs. In this article, we do not explicitly explore defenses for preventing malicious users from exploiting these flaws. However, our findings signify that restricting the AT interface through access control policies, black-listing may not work due to the parsing bugs and invalid AT commands that the parser executes. Completely removing the exposure of AT modem interface over Bluetooth and USB can resolve the problem. Other than that, at a conceptual level, having a formal grammar specification of the supported AT command grammar may provide a better way to test the AT interface. Another aspect that particularly requires attention is the deployment of stricter policies that filter out anomalous AT commands.

Responsible disclosure. Given the sensitive nature of our findings, we have reported these to the relevant stakeholders (e.g., respective modems and devices vendors and manufacturers). Moreover, following the responsible disclosure policy, we have waited 90 days before making our findings public. Both Samsung and Google have released patches to resolve the reported vulnerabilities. The reported vulnerabilities have been assigned (CVE-2019-16400 and CVE-2019-16401).

8 CONCLUSION AND FUTURE WORK

This article proposes ATFuzzer for testing the correctness of the AT interface exposed by the baseband processor in a smartphone. Towards this goal, ATFuzzer leverages a grammar-guided evolutionary fuzzing-based approach. Unlike generational fuzzers that use the input grammar to generate syntactically correct inputs, ATFuzzer mutates the production rules in the grammar itself. Such an approach enables ATFuzzer to not only efficiently navigate the input search space but also allows it to exercise a diverse set of input AT commands. In our evaluation with ATFuzzer on 10 Android smartphones from 6 vendors revealed 4 invalid AT command grammars that are processed by the Bluetooth AT interface and can induce DoS, downgrade connectivity, and privacy leaks. For the USB AT interface, however, ATFuzzer identified 13 invalid AT command grammars that are equally damaging as the ones found for the Bluetooth AT interface. Our findings have been responsibly shared with the relevant stakeholders among which Samsung has acknowledged our findings and are working towards a patch. Two of our findings have also been assigned CVEs (CVE-2019-16400 and CVE-2019-16401).

As part of future work, we plan to apply hybrid fuzzing to our problem domain. In the hybrid fuzzing paradigm, a black-box fuzzer's capabilities are enhanced through the use of lightweight static analysis (e.g., dynamic symbolic execution, taint analysis). Such an approach would, however, require us to address the issues concerning firmware binaries' practice of employing obfuscation and encryption.

One option is to apply both reverse engineering and binary instrumentation approaches to examine firmware binaries and further investigate the AT execution process. This will also allow us to define a more suitable fitness evaluation function to improve the ATFuzzer's evaluation module.

Finally, we want to improve the AT command extractor to retrieve more accurate information. Currently, we can extract the list of AT commands alongside the sequence of expected parameters. However, some of the

details about such parameters are embedded and mixed into the text in the specifications. Thus, by employing text analysis strategies and natural language processing approaches, we can further improve the extraction process and increase its accuracy.

REFERENCES

- [1] Android Developers. [n.d.]. Configure on-device developer options. Retrieved from <https://developer.android.com/studio/debug/dev-options>. [Online; accessed May 1, 2020].
- [2] [n.d.]. CWE-325: Missing required cryptographic step - CVE-2018-5383. Carnegie Mellon University, CERT Coordination Center. Retrieved from <https://www.kb.cert.org/vuls/id/304725/>. [Online; accessed May 1, 2020].
- [3] [n.d.]. Neo 1973 and Neo FreeRunner GSM modem, AT Command set. Retrieved from http://wiki.openmoko.org/wiki/Neo_1973_and_Neo_FreeRunner_gsm_modem. [Online; accessed May 1, 2020].
- [4] [n.d.]. XDA Forum [online]. Retrieved from <https://forum.xda-developers.com/galaxy-s2/help/how-to-talk-to-modem-commands-t1471241>. [Online; accessed May 1, 2020].
- [5] M. Herfurt, A. Laurie, M. Holtmann. [n.d.]. The BlueBug. AL Digital Ltd. Retrieved from https://trifinite.org/trifinite_stuff_bluebug.html#introduction. [Online; accessed May 1, 2020].
- [6] Iosif Androulidakis. 2011. Intercepting mobile phone calls and short messages using a GSM tester. In *Proceedings of the International Conference on Computer Networks*. Springer, 281–288.
- [7] Cornelius Aschermann, Tommaso Frassetto, Thorsten Holz, Patrick Jauernig, Ahmad-Reza Sadeghi, and Daniel Teuchert. 2019. NAU-TILUS: Fishing for deep bugs with grammars. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'19)*.
- [8] Marcel Böhme, Van-Thuan Pham, Manh-Dung Nguyen, and Abhik Roychoudhury. 2017. Directed greybox fuzzing. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2329–2344.
- [9] ICP DAS. [n.d.]. GTM-203-3GWA AT Commands Manual. Retrieved from http://ftp.icpdas.com/pub/cd/usbcd/napdos/3g_modem/gtm-203m-3gwa/manual/gtm-203m-3gwa_atcommands_manual.pdf. [Online; accessed May 1, 2020].
- [10] S. Gan, C. Zhang, X. Qin, X. Tu, K. Li, Z. Pei, and Z. Chen. 2018. CollAFL: Path sensitive fuzzing. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'18)*, Vol. 00. 660–677. DOI : <https://doi.org/10.1109/SP.2018.00040>
- [11] Roei Hay. 2017. fastboot oem vuln: Android bootloader vulnerabilities in vendor customizations. In *Proceedings of the 11th {USENIX} Workshop on Offensive Technologies ({WOOT}'17)*.
- [12] Roei Hay and Michael Goberman. 2017. Attacking nexus 6 & 6P custom bootmodes. Retrieved from <https://www.docdroid.net/dxKUj5c/attacking-nexus-6-6p-custom-bootmodes.pdf>.
- [13] Aki Helin. [n.d.]. Radamsa. Retrieved from <https://gitlab.com/akihe/radamsa>. [Online; accessed May 1, 2020].
- [14] Christian Holler, Kim Herzig, and Andreas Zeller. 2012. Fuzzing with code fragments. In *Proceedings of the 21st USENIX Conference on Security Symposium (Security'12)*. USENIX Association, 38.
- [15] Syed Rafiul Hussain, Omar Chowdhury, Shagufta Mehnaz, and Elisa Bertino. 2018. LTEInspector: A systematic approach for adversarial testing of 4G LTE. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS'18)*.
- [16] Syed Rafiul Hussain, Mitziu Echeverria, Omar Chowdhury, Ninghui Li, and Elisa Bertino. 2019. Privacy attacks to the 4G and 5G cellular paging protocols using side channel information. (2019).
- [17] Motorola Inc.[n.d.]. Motorola AT Command Set. Retrieved from https://ipfs.io/ipfs/QmXoypizjW3WknFjInKLwHCnL72vedxjQkDDP1mXWo6uco/wiki/Motorola_phone_AT_commands.html. [Online; accessed May 1, 2020].
- [18] Multi-Tech Systems Inc.[n.d.]. AT Commands For CDMA Wireless Modems. Retrieved from https://canarysystems.com/downloads/documentation/third_party_documentation/CDMA_AT_Commands.pdf. [Online; accessed May 1, 2020].
- [19] Multi-Tech Systems Inc.[n.d.]. EVDO and CDMA AT Commands Reference Guide. Retrieved from <https://www.multitech.com/documents/publications/manuals/s000546.pdf>. [Online; accessed May 1, 2020].
- [20] Sony Mobile Communications Inc.[n.d.]. Sony Ericsson AT Command set. Retrieved from <https://www.activexperts.com/sms-component/at/sonyericsson/>. [Online; accessed May 1, 2020].
- [21] ETSI (European Telecommunications Standards Institute). [n.d.]. Digital cellular telecommunications system (Phase 2+); AT Command set for GSM Mobile Equipment (ME) (3GPP TS 07.07 version 7.8.0 Release 1998). Retrieved from https://www.etsi.org/deliver/etsi_ts/100900_100999/100916/07.08.00_60/ts_100916v070800p.pdf. [Online; accessed May 1, 2020].
- [22] ETSI (European Telecommunications Standards Institute). [n.d.]. Digital cellular telecommunications system (Phase 2+) (GSM); Universal Mobile Telecommunications System (UMTS); LTE; AT command set for User Equipment (UE) (3GPP TS 27.007 version 13.6.0 Release 13). Retrieved from https://www.etsi.org/deliver/etsi_ts/127000_127099/127007/13.06.00_60/ts_127007v130600p.pdf. [Online; accessed May 1, 2020].
- [23] ETSI (European Telecommunications Standards Institute). [n.d.]. Digital cellular telecommunications system (Phase 2+); Specification of the Subscriber Identity Module -Mobile Equipment (SIM-ME) interface (3GPP TS 51.011 version 4.15.0 Release 4). Retrieved from https://www.etsi.org/deliver/etsi_TS/151000_151099/151011/04.15.00_60/ts_151011v041500p.pdf. [Online; accessed May 1, 2020].

- [24] ETSI (European Telecommunications Standards Institute). [n.d.]. Digital cellular telecommunications system (Phase 2+), Universal Mobile Telecommunications System UMTS, LTE, AT command set for User Equipment UE. Retrieved from https://www.etsi.org/deliver/etsi_ts/127000_127099/127007/10.03.00_60/ts_127007v100300p.pdf. [Online; accessed May 1, 2020].
- [25] ETSI (European Telecommunications Standards Institute). [n.d.]. Digital cellular telecommunications system (Phase 2+); Use of Data Terminal Equipment - Data Circuit terminating; Equipment (DTE - DCE) interface for Short Message Service (SMS) and Cell Broadcast Service (CBS) (GSM 07.05 version 5.3.0). Retrieved from https://www.etsi.org/deliver/etsi_gts/07/0705/05.03.00_60/gsmts_0705v050300p.pdf. [Online; accessed May 1, 2020].
- [26] Imtiaz Karim, Fabrizio Cicala, Syed Rafiul Hussain, Omar Chowdhury, and Elisa Bertino. 2019. Opening Pandora's box through AT-Fuzzer: Dynamic analysis of AT interface for Android smartphones. In *Proceedings of the 35th Computer Security Applications Conference (ACSAC'19)*. Association for Computing Machinery, New York, NY, 529–543. DOI: <https://doi.org/10.1145/3359789.3359833>
- [27] George Klees, Andrew Ruef, Benji Cooper, Shiyi Wei, and Michael Hicks. 2018. Evaluating fuzz testing. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS'18)*. ACM, New York, NY, 2123–2138. DOI: <https://doi.org/10.1145/3243734.3243804>
- [28] Caroline Lemieux, Rohan Padhye, Koushik Sen, and Dawn Song. 2018. PerfFuzz: Automatically generating pathological inputs. In *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA'18)*. ACM, New York, NY, 254–265. DOI: <https://doi.org/10.1145/3213846.3213874>
- [29] Caroline Lemieux and Koushik Sen. 2018. FairFuzz: A targeted mutation strategy for increasing greybox fuzz testing coverage. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE'18)*. ACM, New York, NY, 475–485. DOI: <https://doi.org/10.1145/3238147.3238176>
- [30] Yuekang Li, Bihuan Chen, Mahinthan Chandramohan, Shang-Wei Lin, Yang Liu, and Alwen Tiu. 2017. Steelix: Program-state based binary fuzzing. In *Proceedings of the 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2017)*. ACM, New York, NY, 627–637. DOI: <https://doi.org/10.1145/3106237.3106295>
- [31] Yuwei Li, Shouling Ji, Chenyang Lv, Yuan Chen, Jianhai Chen, Qinchun Gu, and Chunming Wu. 2019. V-Fuzz: Vulnerability-oriented evolutionary fuzzing. *CoRR* abs/1901.01142 (2019).
- [32] Angela Lonozetta, Peter Cope, Joseph Campbell, Bassam Mohd, and Thamer Hayajneh. 2018. Security vulnerabilities in Bluetooth technology as used in IoT. *J. Sensor Actuat. Netw.* 7, 3 (2018), 28.
- [33] Huawei Technologies Co. Ltd. [n.d.]. HUAWEI MU609 HSPA LGA Module Application Guide. Retrieved from [https://www.paoli.cz/out/media/HUAWEI_MU609_HSPA_LGA_Module_Application_Guide_V100R002_02\(1\).pdf](https://www.paoli.cz/out/media/HUAWEI_MU609_HSPA_LGA_Module_Application_Guide_V100R002_02(1).pdf). [Online; accessed May 1, 2020].
- [34] Ulrike Meyer and Susanne Wetzel. 2004. A man-in-the-middle attack on UMTS. In *Proceedings of the 3rd ACM Workshop on Wireless Security*. ACM, 90–97.
- [35] Barton P. Miller, Louis Fredriksen, and Bryan So. 1990. An empirical study of the reliability of UNIX utilities. *Commun. ACM* 33, 12 (Dec. 1990), 32–44. DOI: <https://doi.org/10.1145/96267.96279>
- [36] Collin Mulliner and Charlie Miller. 2009. Fuzzing the phone in your phone. In *Proceedings of the Black Hat USA Conference*.
- [37] NIST (National Institute of Standards and Technology). [n.d.]. CVE-2016-4030. Retrieved from <https://nvd.nist.gov/vuln/detail/CVE-2016-4030>. [Online; accessed May 1, 2020].
- [38] NIST (National Institute of Standards and Technology). [n.d.]. CVE-2016-4031. Retrieved from <https://nvd.nist.gov/vuln/detail/CVE-2016-4031>. [Online; accessed May 1, 2020].
- [39] NIST (National Institute of Standards and Technology). [n.d.]. CVE-2016-4032. Retrieved from <https://nvd.nist.gov/vuln/detail/CVE-2016-4032>. [Online; accessed May 1, 2020].
- [40] NIST (National Institute of Standards and Technology). [n.d.]. CVE-2019-16400. Retrieved from <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-16400>. [Online; accessed May 1, 2020].
- [41] NIST (National Institute of Standards and Technology). [n.d.]. CVE-2019-16401. Retrieved from <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-16401>. [Online; accessed May 1, 2020].
- [42] André Pereira, Manuel Correia, and Pedro Brandão. 2014. Charge your device with the latest malware. In *Proceedings of the Black Hat Europe Conference*.
- [43] André Pereira, Manuel Correia, and Pedro Brandão. 2014. USB connection vulnerabilities on Android smartphones: Default and vendors' customizations. In *Communications and Multimedia Security*, Bart De Decker and André Zúquete (Eds.). Springer Berlin, 19–32.
- [44] André Pereira, Manuel Correia, and Pedro Brandão. 2014. USB connection vulnerabilities on Android smartphones: Default and vendors' customizations. In *Proceedings of the IFIP International Conference on Communications and Multimedia Security*. Springer, 19–32.
- [45] Theofilos Petsios, Jason Zhao, Angelos D. Keromytis, and Suman Jana. 2017. SlowFuzz: Automated domain-independent detection of algorithmic complexity vulnerabilities. *CoRR* abs/1708.08437 (2017).
- [46] Van-Thuan Pham, Marcel Böhme, Andrew E. Santosa, Alexandru Razvan Caciulescu, and Abhik Roychoudhury. 2018. Smart greybox fuzzing. *CoRR* abs/1811.09447 (2018).
- [47] Sanjay Rawat, Vivek Jain, Ashish Kumar, Lucian Cojocar, Cristiano Giuffrida, and Herbert Bos. 2017. Vuzzer: Application-aware evolutionary fuzzing. In *Proceedings of the Network and Distributed System Security Symposium (NDSS'17)*.

- [48] P. Roberto and F. Aristide. 2014. Modem interface exposed via USB. In *Proceedings of the Black Hat Europe Conference*. Retrieved from <https://github.com/ud2/advisories/tree/master/android/samsung/nocve-2016-0004>.
- [49] J. Ruderman, G. Kwong, C. Holler, J. de Mooij, D. Keeler, J. Schwartzentruber, and The SpiderMonkey. [n.d.]. jsfunfuzz. Retrieved from <https://github.com/MozillaSecurity/funfuzz/tree/master/src/funfuzz/js/jsfunfuzz>. [Online; accessed May 1, 2020].
- [50] David Rupprecht, Katharina Kohls, Thorsten Holz, and Christina Pöpper. 2019. Breaking LTE on layer two. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA*. 1121–1136. DOI: [10.1109/SP.2019.00006](https://doi.org/10.1109/SP.2019.00006)
- [51] Mike Ryan. 2013. Bluetooth: With low energy comes low security. In *Proceedings of the 7th {USENIX} Workshop on Offensive Technologies*.
- [52] SIMCom. [n.d.]. AT Command Set SIMCOM-SIM5320-ATCEN-V2.02. Retrieved from https://cdn-shop.adafruit.com/datasheets/SIMCOM_SIM5320_ATC_EN_V2.02.pdf. [Online; accessed May 1, 2020].
- [53] PEACH TECH. [n.d.]. Peach fuzzer platform. Retrieved from <https://www.peach.tech/>. [Online; accessed May 1, 2020].
- [54] Wireless Solutions Telit. [n.d.]. AT commands reference guide. Retrieved from https://www.telit.com/wp-content/uploads/2017/09/Telit_AT_Commands_Reference_Guide_r24_B.pdf. [Online; accessed May 1, 2020].
- [55] Dave (Jing) Tian, Grant Hernandez, Joseph I. Choi, Vanessa Frost, Christie Raules, Patrick Traynor, Hayawardh Vijayakumar, Lee Harrison, Amir Rahmati, Michael Grace, and Kevin R. B. Butler. 2018. ATtention spanned: Comprehensive vulnerability analysis of AT commands within the Android ecosystem. In *Proceedings of the 27th USENIX Security Symposium (USENIX Security'18)*. 273–290. Retrieved from <https://www.usenix.org/conference/usenixsecurity18/presentation/tian>.
- [56] Spandan Veggalam, Sanjay Rawat, Istvan Haller, and Herbert Bos. 2016. IFuzzer: An evolutionary interpreter fuzzer using genetic programming. In *Computer Security – ESORICS 2016*, Ioannis Askoxylakis, Sotiris Ioannidis, Sokratis Katsikas, and Catherine Meadows (Eds.). Springer International Publishing, Cham, 581–601.
- [57] J. Wang, B. Chen, L. Wei, and Y. Liu. 2017. Skyfire: Data-driven seed generation for fuzzing. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'17)*. 579–594. DOI: <https://doi.org/10.1109/SP.2017.23>
- [58] Junjie Wang, Bihuan Chen, Lei Wei, and Yang Liu. 2018. Superion: Grammar-aware greybox fuzzing. *CoRR* abs/1812.01197 (2018).
- [59] Wikipedia. [n.d.]. Wikipedia. Retrieved from https://en.wikipedia.org/wiki/Hayes_command_set. [Online; accessed May 1, 2020].
- [60] Christos Xenakis and Christoforos Ntantogian. 2015. Attacking the baseband modem of mobile phones to breach the users' privacy and network security. In *Proceedings of the 7th International Conference on Cyber Conflict: Architectures in Cyberspace (CyCon'15)*. IEEE, 231–244.
- [61] Christos Xenakis, Christoforos Ntantogian, and Orestis Panos. 2016. (U) SimMonitor: A mobile application for security evaluation of cellular networks. *Comput. Secur.* 60 (2016), 62–78.
- [62] M. Zalewski. [n.d.]. American fuzzy lop. Retrieved from <http://lcamtuf.coredump.cx/afl/>. [Online; accessed May 1, 2020].
- [63] Michal Zalewski. [n.d.]. Mangleme. Retrieved from <https://github.com/WebKit/webkit/tree/master/Tools/mangleme>. [Online; accessed May 1, 2020].

Received May 2020; accepted August 2020