

BinderCracker: Assessing the Robustness of Android System Services

Huan Feng, Kang G. Shin

*Department of Electrical Engineering and Computer Science
The University of Michigan
{huanfeng, kgshin}@umich.edu*

Abstract

In Android, communications between apps and system services are supported by a transaction-based Inter-Process Communication (IPC) mechanism. *Binder*, as the cornerstone of this IPC mechanism, separates two communicating parties as client and server. As with any client-server model, the server should not make any assumption on the validity (sanity) of client-side transaction. To our surprise, we find this principle has frequently been overlooked in the implementation of Android system services. In this paper, we demonstrate the prevalence and severity of this vulnerability surface and try to answer why developers keep making this seemingly simple mistake. Specifically, we design and implement *BinderCracker*, an automatic testing framework that supports parameter-aware fuzzing and has identified more than 100 vulnerabilities in six major versions of Android, including the latest version Android 6.0, Marshmallow. Some of the vulnerabilities have severe security implications, causing privileged code execution or permanent Denial-of-Service (DoS). We analyzed the root causes of these vulnerabilities to find that most of them exist because system service developers only considered exploitations via public APIs. We thus highlight the deficiency of testing only on client-side public APIs and argue for the necessity of testing and protection on the *Binder* interface — the actual security boundary. Specifically, we discuss the effectiveness and practicality of potential countermeasures, such as precautionary testing and runtime diagnostic.

1 Introduction

Android is the most popular smartphone OS and dominates the global market with a share of more than 82% [34]. By the end of 2015, the total number of Android devices surpassed 1.4 billion, and there were more than 1.6 million mobile apps available in Google Play for download [15, 26]. The developers of these apps are not always trustworthy; many of them might be inexpe-

rienced, careless or even malicious. Therefore, proper isolation between apps and the system is essential for robustness and security.

To meet this requirement, apps in Android execute in application sandboxes. If an app wants to interact with system services, it must perform Inter-Process Communications (IPC). *Binder*, as the cornerstone of this IPC mechanism, separates two communicating parties as client and server. Each server (i.e., system service in this case) exports a list of public APIs that clients (i.e., mobile apps) can invoke. The input parameters of each API go through extensive sanity checks on the client-side before being packed into a transaction and then sent to the server. An AIDL (Android Interface Description Language) file further enforces the schema of each transaction, serving as an explicit contract between the client and the server. These client-side public APIs get examined and tested by thousands of participating vendors of AOSP (Android Open Source Project) and are believed to be robust.

Robust public APIs and AIDL enforcement can protect system services against erroneous input from careless or inexperienced app developers, but render ineffective against developers with an adversarial mindset. This is because both enforcements reside in the process space of the client (attacker) and thus can eventually be circumvented. In other words, any system service that hinges on the validity (sanity) of client-side transaction is fundamentally vulnerable — the server should be robust on its own. This is probably a best engineering practice for any system that adopts a client-server model, but has surprisingly been overlooked in the implementation of many Android system services. In this paper, we conduct a comprehensive study of this attack surface in Android. Specifically, we answer two important questions:

- Q1. *How prevalent is this issue in Android, a major open source project?*
- Q2. *Why are so many developers making this seemingly simple mistake?*

To answer these questions, we designed and implemented *BinderCracker*, an automatic testing framework for Android system services. *BinderCracker* automatically crawls the RPC interfaces of both Java and native system services, and injects fuzzing transactions via the Binder surface. This directly challenges the robustness of the error-handling mechanisms in the system services. To increase the scale and depth of our testing, *BinderCracker* supports parameter-aware fuzzing with semi-valid inputs. This requires us to record and mutate existing transactions and understand their semantics. Doing this in the context of Android is challenging because a transaction may contain remote object handles that cannot be recorded in the form of raw bytes. Moreover, transactions frequently contain dynamic and non-primitive data types that are difficult to mutate in a sensible way. *BinderCracker* overcomes these challenges by implementing a replay engine that has in-depth understanding of Binder transactions. It utilizes the dependencies between transactions to reconstruct remote object handles during runtime, and tracks the hierarchy of non-primitive data types to unmarshal them into primitive types.

We examined more than 2400 service APIs in 6 major versions of Android, including the latest Android 6.0 (Marshmallow). In total, we identified more than 100 vulnerabilities, most of which are unfixed to date. Many of the vulnerabilities we identified are found to be able to crash the entire Android Runtime, while others can cause specific system services or system apps to fail. Some vulnerabilities have severe security implications, and may result in system memory corruption, privileged code execution, targeted or permanent Denial-of-Service (DoS). This extensive testing also demonstrates the effectiveness of our parameter-aware fuzzing — it identified 7x more vulnerabilities than simple black-box fuzzing with the same amount of time. Furthermore, since *BinderCracker* is parameter-aware, we can now unearth different vulnerabilities in the same RPC method by fuzzing different parameters in the same transaction.

We further analyzed the root causes of the identified vulnerabilities and studied 117 of them in Android source codes. Most of them exist because the developers only considered exploitations of public APIs. Therefore, many risky scenarios are assumed to be ‘unlikely’ or even ‘impossible’ in their mindset. Here, we list three most common mistakes made by system service developers that contribute to a vast majority of the vulnerabilities. First, private APIs are assumed to be unknown to others, thus no sanity-check is made. Second, client-side enforcements are assumed to be secure, so there is no double check of them at the server-side. Third, the de-serialization process is assumed to be always undisturbed, hence no sanity check is conducted during this

process. All of these are unreliable assumptions because they either directly or indirectly depend on the validity of client-side transactions. Moreover, we demonstrate most of the vulnerabilities can never be found by testing only public APIs. We, therefore, highlight the deficiency of testing only client-side public APIs and argue for the necessity of testing and protection at the Binder surface — the actual security boundary.

As our findings indicate, new vulnerabilities keep on emerging, especially with releases of new Android versions. To address this emerging attack surface, we need to eliminate potential vulnerabilities as early as possible in the development cycle. Specifically, we suggest the use of various precautionary testing techniques (including *BinderCracker*) before each product release. This can stop a large number of vulnerabilities from reaching the end-users. In fact, many severe vulnerabilities [6, 7, 27] could have been avoided, had *BinderCracker* been deployed. We also describe how to enhance the visibility of Binder transaction during runtime to support more informative runtime diagnosis, in case some vulnerabilities leak through *BinderCracker* and eventually reach the end-users.

This paper makes five main contributions by:

- Conducting the first comprehensive study that assesses the robustness of Android system services and unveiling an alarming perspective.
- Designing and implementing *BinderCracker*, an automatic testing framework that supports parameter-aware fuzzing on the Binder surface;
- Conducting an extensive test on 6 major Android versions and identifying 100+ vulnerabilities, some of which have severe security implications;
- Unearthing the root causes of these vulnerabilities and highlighting the necessity of protection on Binder — the actual security boundary;
- Discussing the effectiveness and practicality of potential countermeasures, such as precautionary testing and runtime diagnostic techniques.

The rest of the paper is organized as follows. Section 2 summarizes related work in the field of software testing and Android security. Section 3 introduces Binder and AIDL in Android, and describes how Android uses these to build system services. Section 4 examines the attack surface under our investigation. Section 5 details the design and implementation of an automatic testing framework, *BinderCracker*, which exposes vulnerable system services. Section 6 describes our tests on stock Android firmware and analyzes the root causes and security implications of the discovered vulnerabilities. Section 7 gives

a comprehensive discussion on how to effectively eliminate these vulnerabilities in the development cycle. Section 8 discusses the long-term value and other potential uses of our work, and finally, the paper concludes with Section 9.

2 Related Work

Discussed below is related work in the field of software testing and Android security.

Software Testing. In the software community, robustness testing falls into two categories: *functional* and *exceptional* testing. Functional testing focuses on verifying the functionality of software using expected input, while exceptional testing tries to apply unexpected and faulty inputs to crash the system. Numerous efforts have been made in the software testing community to test the robustness of Android [1, 2, 17, 21, 24, 36]. Most of them focus on the functional testing of GUI elements [1, 2, 17, 21]. Some have conducted exceptional testing on the evolving public APIs [24]. In this paper, we highlight the deficiency of testing only on public APIs and conduct an exceptional testing on lower-level Binder-based RPC interfaces.

Android Security. Android has received significant attention from the research community as an open source operating system [4, 10, 11, 16, 23, 28, 32]. Existing Android security studies largely focus on the imperfection of high-level permission model [12, 13, 25], and the resulting issues, such as information leakage [10], privilege escalation [4, 28] and collusion [23]. Our work highlights the insufficient protection of Android’s lower-level Binder-based RPC mechanism and how it affects the robustness of system services.

There also exist a few studies focusing on the IPC mechanism of Android [5, 9, 19, 22, 29]. However, they largely focus on one specific instance of Android IPC — Intent. Since the senders and recipients of Intents are both apps, manipulating Intents will not serve the purpose of exposing vulnerabilities in system services. Some researchers also provide recommendations for hardening Android IPCs [19, 22] and point out that the key issue in Intent communication is the lack of formal schema. We demonstrate that even for mechanisms enforcing a formal schema, such as AIDL, robustness remains as a critical issue. Gong *et al.* [14] also conducted experiments on the fuzzing of Binder interface. However, they focused on implementing Proof-of-Concept (PoC) exploits using identified vulnerabilities, instead of comprehensively assessing and understanding the attack surface. In fact, they only tested simple black-box

fuzzing on one Android version, which is a small subset of our work. We regard their work as a parallel and independent effort from an industry perspective.

3 Android IPC and Binder

Android executes apps and system services as different processes and enforces isolation between them. To enable different processes to exchange information with each other, Android provides, Binder, a secure and extensible IPC mechanism. Described below are the basic concepts in the Binder framework and an explanation of how a typical system service is built using these basic primitives.

3.1 Binder

In Android, Binder provides a message-based communication channel between two processes. It consists of (i) a kernel-level driver that achieves communication across process boundaries, (ii) a Binder library that uses `ioctl` syscall to talk with the kernel-level driver, and (iii) upper-level abstracts that utilize the Binder library. Conceptually, Binder takes a classical client-server architecture. A client can send a transaction to the remote server via the Binder framework and then retrieves its response. The parameters of the transaction are marshalled into a `Parcel` object which is a serializable data container. The `Parcel` object is sent through the Binder driver and then gets delivered to the server. The server de-serializes the parameters of the `Parcel` object, processes the transaction, and returns a response in a similar way back to the client. This allows a client to achieve *Remote Procedure Call* (RPC) and invoke methods on remote servers as if they were local. This Binder-based RPC is one of the most frequent forms of IPC in Android, and underpins the implementation of most system services.

3.2 AIDL

Many RPC systems use IDL (*Interface Description Language*) to define and restrict the format of a remote invocation [22], and so does Android. The AIDL (*Android Interface Description Language*) file allows the developer to define the RPC interface both the client and the server agree upon [3]. Android can automatically generate Stub and Proxy classes from an AIDL file and relieve the developers from (re-)implementing the low-level details to cope with native Binder libraries. The auto-generated Stub and Proxy classes will ensure that the declared list of parameters will be properly serialized, sent, received, and de-serialized. The developer only needs to provide a `.aidl` file and implement the corresponding interface. In other words, the AIDL file serves as an

```

interface IQueueService {
    boolean add(String name);
    String peek();
    String poll();
    String remove();
}

```

Figure 1: An example AIDL file which defines the interface of a service that implements a queue.

explicit contract between client and server. This enforcement makes the Binder framework extensible, usable, and robust. Fig. 1 shows an example AIDL file that defines the interface of a service that implements a queue.

3.3 System Service

We now describe how the low-level concepts in the Binder framework are structured to deliver a system service, using Wi-Fi service as an example. To implement the Wi-Fi service, system developers only need to define its interfaces as an AIDL description, and then implement the corresponding server-side logic (`WifiService`) and client-side wrapper (`WifiManager`) (see Fig. 2). The serialization, transmission, and deserialization of the interface parameters are handled by the codes automatically generated from the AIDL file. Specifically, when the client invokes some RPC method in the client-side wrapper `WifiManager`, the Proxy class `IWifiManager.Stub.Proxy` will marshal the input parameters in a `Parcel` object and send it across the process boundary via the Binder driver. The Binder library at the server-side will then unmarshal the parameters and invoke the `onTransact` function in the Stub class `IWifiManager.Stub`. This eventually invokes the service logic programmed in `WifiService`. Fig. 2 provides a clear illustration of the entire process.

4 The Attack Surface

The Binder framework separates two communicating parties as client and server. In the most common scenario, the client is a user-level app and the server is a system service. As with any client-server model, the server should never trust the client. In the particular case of Android which uses the Binder framework to build a light-weight RPC mechanism, the following two important properties have to be guaranteed.

- The RPC interfaces in both client and server sides should be consistent: they should expect the same list of input and return parameters.




Figure 2: How does an app communicate with a system service using Binder-based RPC (using Wi-Fi service as an example)? The red shaded region represents the codes that need to be provided/implemented by the service developer.

- Each parameter of the RPC interface should be properly checked: the server should only accept a transaction if all of its parameters are valid.

To guarantee these properties, Android adopts an AIDL enforcement and conducts extensive testings on public APIs. AIDL serves as an explicit contract between the server and the client. It defines the RPC interfaces a service trying to provide. A system service developer can work on top of an AIDL interface and leave the serialization, transmission and de-serialization to be handled automatically by the codes generated from the AIDL file. Bugs in the codes that are manually written by the developer are eliminated further by testing on public APIs and feedbacks from thousands of vendors and hundreds of millions of users. Thanks to these mechanisms, Android services, especially the widely-used and well-maintained system services, could be made robust.

However, when the client is malicious and trying to sabotage system services intentionally, these mechanisms cannot be counted on as a secure measure. This is because both enforcements reside in the process space of the client (attacker), thus can eventually be circumvented. Specifically, although AIDL enforcement and public API testing consolidate the upper and middle layers of the program stack, the service APIs (RPC interfaces) are still directly exposed in the low-layer of the stack — the Binder driver. By directly injecting faulty transactions into the Binder driver, an attacker can circumvent these existing protections in the upper layers and directly confront the server-side exception handling mechanisms. Fig. 3 gives an illustrative view of this attack surface.

Ideally, this should not affect the robustness of system services. The server-side codes should be robust on its own without making any assumption on the client-side

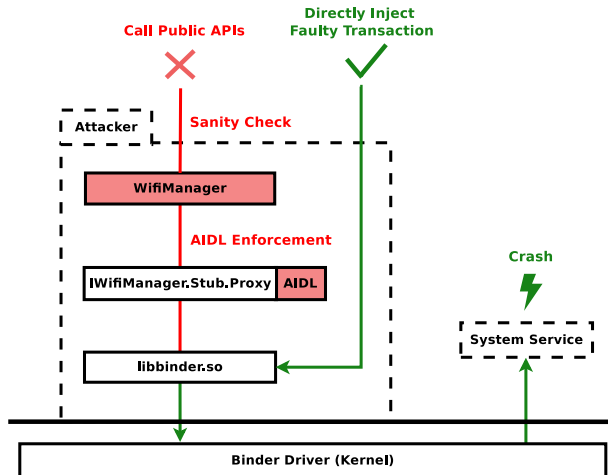


Figure 3: By injecting faulty transactions via the Binder driver, an attacker can bypass the protections in the upper layers of the program stack, and directly confront the server-side exception handling mechanisms.

transactions. This is probably a best engineering practice for any system that adopts a client-server model. However, we found this principle frequently overlooked in the implementation of Android system services, thus motivating us to perform a comprehensive study of this attack surface. Specifically, we want to answer two questions: (1) *how prevalent is this issue in Android, a major open source project?* and (2) *why are so many developers making this seemingly simple mistake?*

4.1 Attack Model

In this paper, we assume the adversary is a malicious app developer trying to sabotage system services. The adversary may mount this attack for various malicious purposes, such as launching a Denial-of-Service (DoS) attack, achieving privileged code execution, etc. A system service can be generic, existing in Android framework base, or vendor-specific, introduced by device manufacturers. The attacker has no root permission and cannot penetrate the security of OS kernel.

5 Automatic Vulnerability Discovery

To answer the questions raised above, we design and implement an automatic testing framework, *BinderCracker*, that can effectively unearth vulnerabilities in system services. To the best of our knowledge, *BinderCracker* is the first tool that supports parameter-aware fuzzing on the Binder surface.

```

struct binder_transaction_data {
    union {
        size_t handle; // (1).target service
        void *ptr;
    }target;
    void *cookie;
    unsigned int code; // (2).RPC method
    unsigned int flags;
    pid_t sender_pid;
    uid_t sender_euid;
    size_t data_size;
    size_t offsets_size;
    union {
        struct {
            binder_uintptr_t buffer;
            binder_uintptr_t offsets;
        } ptr;
        __u8 buf[8];
    } data; // (3).transactional data
};

```

Figure 4: The data struct sent through the Binder driver via the `ioctl` libc call. This struct contains three important pieces of information we need to modify to send a fuzzing transaction.

5.1 BinderCracker: An Overview

To fuzz a system service, *BinderCracker* must be able to send mal-formatted transactions to it. The attack surface we use is fundamental — the client can send any arbitrary transaction to the server because the client has complete control over its own process space, and can thus bypass any client-side enforcement. This can be achieved by either taking advantage of hidden Android APIs or hijacking the libc call that underpins the Binder communication library. Both of these techniques can be achieved in user-level without extending the Android system [35]. Basically, *BinderCracker* is manipulating (either directly or indirectly) a `binder_transaction_data` struct sent to the Binder driver. This data struct contains three important pieces of information we need to modify to send a fuzzing transaction and has the format as shown in Fig. 4.

The `target.handle` field specifies the service this transaction is sent to. The `code` field represents a specific RPC method we want to fuzz. The data struct contains the serialized bytes of the list of parameters for the RPC method, which is inherently a `Parcel` object. `Parcel` is a container class that provides a convenient set of serialization and de-serialization methods for different data types. Both the client and the server work directly with this `Parcel` object to send and receive the input parameters. Later in this section, we will elaborate on how to modify the `handle` and `code` variables to redirect the transaction to a specific RPC method of a specified ser-

vice, and how to fuzz the `Parcel` object to facilitate testing with different policies.

5.2 Transaction Redirection

There is a one-to-one mapping from the `handle` variable in the `binder_transaction_data` object to system service. This mapping is created during runtime and maintained by the Binder driver. Since the client has no control over the Binder driver, it cannot get this mapping directly. For system services that are statically cached, we can get them indirectly by querying a static service manager which has a fixed `handle` of 0. This service manager is a centralized controller for service registry and will be started before any other services. By sending a service interface descriptor (such as `android.os.IWindowManager`) to the service manager, it will return an `IBinder` object which contains the `handle` for the specified service. For system services that are dynamically allocated, we can retrieve them by recursively replaying the supporting transactions that generate these services. We will elaborate this later when discussing transaction fuzzing.

After getting the `handle` of a system service, we need to further specify the `code` variable in the `binder_transaction_data` object. Each code represents a different RPC method defined in the AIDL file. This mapping can be found in the Stub files which are automatically generated from the AIDL file. The code variable typically ranges from 1 to the total number of methods declared in the AIDL file. For native system services that are not implemented in Java, this mapping is directly coded in either the source files or the header files. Therefore, we scan both the AIDL files and the native source codes of Android to construct the mapping between transaction codes and RPC methods.

5.3 Transaction Fuzzing

After being able to redirect a Binder transaction to a chosen RPC method of a chosen system service, the next step is to manipulate the transaction data and create faulty transactions that are unlikely to occur in normal circumstances. Here, we take three widely-used fuzzing policies: sending empty transaction, random transactions, and semi-valid transactions. The first two policies are easy to implement because it is agnostic of the RPC method we target — we only need to fill the transaction with either `NULL` values or randomly generated bytes. The last policy, however, requires us to understand the semantics of a transaction to fuzz each parameter individually. Next, we explain how *BinderCracker* supports parameter-aware fuzzing with semi-valid transactions and why it is challenging even when we already




Figure 5: When fuzzing a transaction, we need to replay the supporting transactions according to their relative order in the dependency graph. This way, all the remote objects this transaction requires will be reconstructed during runtime.

know the RPC interfaces.

Parameter-Aware Fuzzing. To increase the scale and depth of testing, *BinderCracker* supports fuzzing with semi-valid transactions. A transaction is said to be *semi-valid* if all of the parameters it contains are valid except for one. Semi-valid transactions can dive deeper into the program structure without being early rejected, thus is able to reveal more in-depth vulnerabilities. To test with semi-valid transactions, we need to first record valid (seed) transactions, and then mutate the parameters in each transaction. This requires *BinderCracker* to be parameter-aware and is challenging for two reasons. First, recording a transaction is challenging when the transaction involves remote objects that cannot be recorded as raw bytes. In this scenario, values in the raw bytes are merely handles to the remote objects and become meaningless once out of the current execution context. Second, mutating a transaction is challenging when the transaction contains dynamic or non-primitive data types. Since the internal structure of this data type is unknown, we do not know how to mutate it in a sensible way. For example, many RPC interfaces take `Intent` as an input parameter. As a non-primitive data type, an `Intent` may contain arbitrary types of primitive types (i.e., `Int`, `String`, `Double`), depending on what has been put into it during runtime. Next, we will detail how we overcome these technical challenges.

Remote Object. In Android, more than 14% of RPC methods and 37% of user-level RPC calls involve a remote object. The most frequent form of a remote object is an `IBinder` object, which is widely-used for registering and invoking remote callbacks. Recording the raw bytes of these objects won't work since they are merely object handles. We overcome this challenge by maintaining a dependency graph among transactions. When recording each transaction, we iterate through the list of remote objects it takes as input and generates as output. Then, we construct a dependency graph that records how dynamic `IBinders` are produced and consumed. Before trying to replay a transaction, we need to execute

Version	API	Market	Device	Build #
4.1.1	16	9.0%	Galaxy Note 2	JRO03C
4.2.2	17	12.2%	Galaxy S4	JDQ39
4.4.2	19	36.1%	Galaxy S4	KOT49H
5.0.1	21	16.9%	Nexus 5	LRX22C
5.1.0	22	15.7%	Nexus 5	LMY47I
6.0.0	23	0.7%	Nexus 5	MRA58K

Figure 8: The list of Android versions we tested using *BinderCracker*.

tions, and analyze how they survived a major open-source project like Android.

6.1 Setups

We tested 6 major versions of Android: 4.1 (Jelly Bean), 4.2 (JellyBean), 4.4 (KitKat), 5.0 (Lollipop), 5.1 (Lollipop) and 6.0 (Marshmallow). All our experiments are conducted by running *BinderCracker* on official firmwares from major device manufacturers. An official firmware went through extensive testing by the vendors and is believed to be ready for a public release. Each firmware is tested in the initial state, right after it is installed. We didn’t install any third-party app or change any configuration except for turning on the adb debugging option, ruling out the influence of external factors. Fig. 8 lists the detailed information of all the firmwares we tested.

An RPC method is found to be *vulnerable* if testing it resulted in a fatal exception, crashing part of, or the entire Android Runtime. Each unique crash report (stack traces) under an RPC interface is further referred to as an *individual vulnerability*. For each vulnerability reported here, we followed the process of: 1) identify it on an official ROM, 2) manually confirm that it can be reproduced, and 3) inspect the source codes for a root cause analysis. For vendor-specific vulnerabilities of which source codes are not available, such as many of the customized system services provided by Samsung, we only record the stack trace.

6.2 Black-box Fuzzing Results

We conducted a comprehensive black-box fuzzing test on 6 major versions of Android. Specifically, we examined more than 98 generic system services (by Google) and 72 vendor-specific services (by Samsung), which covers more than 2400 low-level RPC methods. For each method, we sent either an empty transaction or a transaction filled with random bytes. In total, we identified 54 vulnerabilities, 39 of which are found in generic system services, and 15 are found in vendor-specific services.

Level	Exception Type	Count
Java	NullPointerException	13
	StackOverflowError	4
	UnsatisfiedLinkError	2
	ArrayIndexOutOfBoundsException	2
	OutOfResourcesException	1
	OutOfMemoryError	1
	StringIndexOutOfBoundsException	1
	IOException	1
	SEGV_MAPPER	26
	Native	SLTKILL
	SEGV_ACCERR	1

Figure 9: List of exceptions occurred in a comprehensive black-box fuzzing test.

On average, each version of Android we tested contains 15 vulnerabilities. The latest version of Android (6.0) still contains 5 vulnerabilities, 2 of which are new. 8 out of the 54 vulnerabilities can crash the entire Android Runtime (system servers), 13 can crash media servers, and 13 can cause crash of other system services and apps. Most of the identified bugs are due to accessing invalid memory addresses. We also found other causes of a crash such as `StackOverflow`. Fig. 9 list all the exception types and the number of their occurrences discovered in our test.

Note that new vulnerabilities have been kept emerging on this attack surface whenever there is a major upgrade of Android version. We also noticed almost all of the vulnerabilities are found within the first few fuzzing transactions, which means that a longer fuzzing time did not lead to the discovery of new bugs. This suggests the inefficiency of black-box fuzzing, probably due to the extremely large fuzzing space. As we will show later, more vulnerabilities are expected if more semantically-rich fuzzing techniques are used.

6.3 Parameter-aware Fuzzing Results

To increase the effectiveness of our testing, *BinderCracker* supports parameter-aware fuzzing with semi-valid transactions. Generating semi-valid transactions requires recording and mutating of existing valid transactions. Here, we collected more than one million valid transactions by running 30 popular apps in two latest Android versions (Android 5.1 and Android 6.0). Based on this seed dataset, we performed a semi-valid fuzzing test on more than 445 RPC methods of 78 system services. Note that we only tested the RPC methods that appeared in our seed dataset, which is a subset of all available RPC methods. To increase the coverage of the seed dataset, one can increase the duration of data collection or incorporate other data sources, such as the unit test cases for

each system service.

We found that semi-valid fuzzing can significantly increase the scale and the depth of our testing. In total, we identified 89 vulnerabilities in Android 5.1 and Android 6.0 which is 7x more than simple fuzzing. Compared to the vulnerabilities identified using simple fuzzing, the vulnerabilities exposed by semi-valid fuzzing are more interesting and have severer security implications (to be discussed later). Moreover, since semi-valid fuzzing is parameter-aware, we can expose different vulnerabilities in the same RPC method by fuzzing different parameters in the same API. For example, by fuzzing different variables contained in the Intent parameter, *BinderCracker* identified more than 20 vulnerabilities in a single RPC method `startActivity` in `ActivityManagerService`. Semi-valid testing also facilitates the process of identifying the corresponding bug in the source codes since we now know which input parameter results in the crash. Later, we will summarize the root causes of all identified vulnerabilities (in both simple and semi-valid fuzzing).

6.4 Root Cause Analysis

The direct causes of crashes are uncaught exceptions such as `NullPointerException` or `SEGV_MAPPER`, but the fundamental cause behind them is deeper. For each crashed system service of which source codes are available, we looked into the source codes and analyzed the root causes of the vulnerabilities. Specifically, we are interested in why these vulnerabilities survived in a major open source project like Android. In summary, we found that most of the vulnerabilities we identified are very likely to have been overlooked by system service developers. A likely explanation is many system developers only considered exploitation of public APIs, thus directly injecting faulty transactions to the Binder driver creates many scenarios that are believed to be ‘unlikely’ or ‘impossible’ in their mindset. Here, we highlight some of the new attack vectors identified by our approach which contribute to most of the vulnerabilities we identified.

First, an attacker can manipulate RPC parameters even if they are not directly exposed via public APIs. For example, `IAudioFlinger` provides an RPC method `REGISTER_CLIENT`. This method is only implicitly called in the Android middleware and is never exposed via public interfaces. Therefore, the developers of this system service may not expect an arbitrary input from this RPC method and didn’t perform a proper check of the input parameters. In our test, sending a list of null parameters via the Binder driver can easily crash this service. This suggests that developers should not overlook RPC interfaces that are private or hidden.

Second, an attacker can bypass sanity checks around

```
android.widget.RemoteViews
private RemoteViews(Parcel parcel,
    BitmapCache bitmapCache) {
    int mode = parcel.readInt();
    ...
    if (mode == MODE_NORMAL) {
        ...
    } else {
        // recursively calls itself
        mL = new RemoteViews(parcel,
            mBitmapCache);
        // recursively calls itself
        mP = new RemoteViews(parcel,
            mBitmapCache);
        ...
    }
    ...
}
```

Figure 10: The constructor of the `RemoteView` class contains a loophole which can cause a `StackOverflow` exception. Specifically, a bad recursion will occur if the input `Parcel` object follows a certain pattern.

the public API, no matter how comprehensive they are. For example, the `IBluetooth` service provides a method called `registerAppConfiguration`. All of the parameters of this RPC method are directly exposed via a public API and there are multiple layers of sanity check around this interface. Therefore, if there is an erroneous input from the public API, the client will throw an exception and crash without even sending the transaction to the server side. However, using our approach, an attack transaction is directly injected to the Binder driver without even going through these client-side checks. This suggests that the server should always double-check input parameters on its own.

Third, an attacker can exploit the serialization process of certain data types and create inputs that are hazardous at the server side. For example, `RemoteView` is a `Parcelable` object that represents a group of hierarchical views. It contains a loophole in its de-serialization module which can cause a `StackOverflow` exception. As shown in Fig. 10, a bad recursion will occur if the input `Parcel` object follows a certain pattern. By directly manipulating the serialized bytes of the `Parcel` sent via the Binder driver, this loophole can be triggered and crash the server. This suggests that RPC methods with serializable inputs require special attention and sanity check is also essential in the de-serializaiton process.




Figure 11: Many of the vulnerabilities we identified are found to be able to crash the entire Android Runtime (`system_server`), while others can cause specific system services (`mediaserver`) or system apps (`nfc`, `contacts`, etc) to fail.

In summary, most of the vulnerabilities originate from unreliable assumptions on the transactions from the client-side. This suggests that there is a misconception of where the security boundary is for Android system services — while many may assume the security boundary is at the client-side public APIs, the actual security boundary can only stand behind the Binder driver. Therefore, we advocate the importance of pure server-side error handling, and highlight the necessity of direct Binder-level testing on RPC interfaces of remote services.

6.5 Security Implications

Most of the vulnerabilities we discovered can be used to launch a Denial-Of-Service (DoS) attack. Some of them are found to be able to crash the entire Android Runtime, while others can cause specific system services or system apps to fail. Fig. 11 shows the distribution of the affected services (apps). When launching a DoS attack, the attacker can trigger a crash either consistently or only under certain conditions, for example, when a competitor’s app is running. This can create the impression that the competitor’s app is buggy and unusable. We even identified multiple vulnerabilities (in the de-serialization process of `Intent`) that can cause targeted crash of almost any system/user-level apps, without crashing the entire system. Specifically, an attacker can craft an `Intent` that contains a mal-formatted `Bundle` object and send to the target app. This will cause a crash during the de-serialization process of the `Intent` object before the target app can conduct any sanity check. Moreover, it can be very challenging to identify the attacker app under these scenarios because the OS only knows which ser-

vice/app is broken, but cannot tell who crashed it. We will discuss more about the attack attribution process in the next section.

Some of the vulnerabilities we discovered can cause more serious security problems. We found that in several RPC methods, the server-side fails to check potential Integer overflows. This may lead to disastrous consequences when exploited by an experience attacker. For example, in `IGraphicBufferProducer` an Integer overflow exists such that when a new `NativeHandle` is created, the server will malloc smaller memory than it actually requested (see Fig. 12). Subsequent writes to this data struct will corrupt the heap on the server-side. This vulnerability has been demonstrated to be able to achieve privileged code execution, and insert any arbitrary code into `system_server` [6]. We also found a vulnerability in `IContentService` that can lead to a infinite boot-loop, which can only be resolved by factory recovery or flushing a new ROM. This is also classified as High risk according to the official specification of Android severity levels [30].

Besides RPC methods that are not well-implemented, we also discovered RPC methods that are not properly protected by existing Permission models. In official ROMs of Samsung Galaxy 4 (Android 4.2.2 and Android 4.4.2), an attacker can reboot the device by directly sending a transaction to `PackageManagerService` via the Binder driver without requiring the `REBOOT` permission. This is critical since `REBOOT` is a sensitive permission only granted to system apps. The other service is `ICoverManager`, a customized service from Samsung. An attacker can invoke a certain RPC method of `ICoverManager` and block the entire screen with a pop-up blank Activity. The blank Activity cannot be revoked

```

native_handle_t* native_handle_create(int
numFds, int numInts)
{
    // numFds & numInts are not checked!
    native_handle_t* h = malloc( ...
        + sizeof(int)*(numFds+numInts));

    h->version = sizeof(native_handle_t);
    h->numFds = numFds;
    h->numInts = numInts;

    return h;
}

```

Figure 12: The constructor of the `native_handle` has an Integer Overflow vulnerability that can cause a heap corruption on the server-side. This can lead to privileged code execution in `system_server`.

using any virtual or physical button and the only exit is restarting the device.

6.6 Vulnerabilities: Fixed and Unfixed

We examine how many of the vulnerabilities remain unfixed and are potentially zero-day when they are found. Our analysis is based on the public changes of the source codes across different Android versions and revisions. We skip the 15 vulnerabilities in vendor-specific system services and 7 in generic system services due to the unavailability of source codes. Note that not all generic system services are open source, especially when it is related to decryption/encryption or interactions with OEM hardware.

Of the 117 analyzed vulnerabilities in Android code bases, only 18 have been fixed by adding additional sanity checks of input parameters. Another 12 vulnerabilities ‘disappeared’ during several major Android version upgrades either because 1) the corresponding source codes (or API) have been deleted; or 2) new updates in other parts of the source codes accidentally bypass the vulnerable source codes. For example, some crashes are caused by a recursive call in the `RemoteView` class (see Fig. 10). Similar crashes disappeared after Android 5.0. We looked in the source codes and found this is not because the bug has been fixed, but because in new versions of Android a faulty transaction will create an additional `Exception` before it reaches the vulnerable codes. The additional `Exception` is properly caught and accidentally avoids the fatal crash caused by the real vulnerability. We do not consider this as a ‘fix’ since an attacker can still recreate the crash by manually crafting a transaction which bypass the new code updates. As of this writing,




Figure 13: Number of the vulnerabilities that are fixed, disappeared and unfixed.

there are still 87 vulnerabilities left unfixed. Fig. 13 illustrates the proportion of vulnerabilities that are fixed, disappeared and unfixed. We have already submitted all unfixed vulnerabilities to AOSP by the time of this submission.

7 Defenses

As our testing results demonstrated, new vulnerabilities emerge whenever there is a major upgrade of the Android code base. This is because, considering the code size of Android, it is almost impossible to prevent the developers from writing buggy codes. Therefore, the only solution is to eliminate potential bugs as early as possible in the development cycle. To this end, we discuss potential defense mechanisms and summarize them into two categories: (1) precautionary testing, exposing vulnerabilities before releasing the new ROM; (2) runtime defense, defending against potential bugs after the ROM has been deployed. Next, we will go through several potential defense mechanisms in each category and discuss whether it is applicable or practical for our problem. Then, we demonstrate how to enhance the visibility of ongoing attacks by enabling runtime diagnostic of `Binder` transactions.

7.1 Precautionary Testing

Before releasing a new ROM, developers can conduct precautionary testing. The defense can be done early, in the development phase of each system service, or later, after the entire ROM gets built.

Android has already adopted a static code analysis tool, `lint`, to check potential bugs and optimizations for correctness, security, performance, usability, accessibility and internationalization [18]. Specifically, `lint` provides a feature that supports inspection with annotations. This allows the developer to add metadata tags to variables, parameters and return values. For example, the developer can mark an input parameter as

@NonNull, indicating that it cannot be Null, or mark it as @IntRange(from=0, to=255), enforcing that it can only be within a given range. Then, lint automatically analyzes the source codes and prompts potential violations. This can be extended to support inspections of RPC interfaces, allowing developers to explicitly declare the constraints for each RPC input parameter. This way, many potential bugs can be eliminated during the development phase. This defense is practical and comprehensive but requires system developers to specify the metadata tags for each RPC interface.

We can also conduct precautionary testing during runtime after the ROM has been built. Our system, *BinderCracker*, itself is effective in identifying vulnerabilities and can be used as an automatic testing tool. By fuzzing various system services with different policies, a large number of vulnerabilities can be eliminated before reaching the end-users. Actually, many severe vulnerabilities [6, 7, 27] could have been avoided if *BinderCracker* had been deployed. Note that the effectiveness of *BinderCracker* depends on the quality and coverage of the seed transactions. Besides collecting execution traces of a large number of apps, another potential way of generating a comprehensive seed dataset is to incorporate the functional unit tests of each system service.

7.2 Runtime Defense

It will be helpful if Android can provide some real-time defense against potential vulnerabilities even after the ROM has been deployed on end-users' devices. Here, we focus on specific defenses on the Binder layer, excluding generic defenses such as Address Space Layout Randomization (ASLR), SELinux, etc. They have been extensively discussed in other literature [20, 31, 33] and are not specific to our scenario. There are two potential defenses one can provide on the Binder surface during runtime: (i) intrusion detection/prevention, identifying and rejecting transactions that are malicious, and (ii) intrusion diagnostics, making an attack visible after the transaction has already caused some damage.

To provide runtime intrusion prevention, one needs to perform some type of abnormality detection on incoming transactions. This works by examining the input parameters of valid/invalid RPC invocations and characterizing the rules or boundaries. However, in our case, it is not practical for the following reasons. First, Binder transactions occur at a very high frequency but a mobile device is itself constrained in energy and computation power. Second, parameters in Binder transactions are very diverse, codependent, and evolving dynamically during runtime, and hence clear boundaries or rules may not exist. Third, end-users are not likely to endure even the smallest false-positive rate. One can,

of course, build a very conservative blacklist-based system and hard-coding rules of each potential vulnerability in the database. However, this seems unnecessary, especially when Android nowadays supports directly pushing security updates (patches) to devices of end-users.

An alternative solution is to diagnose, instead of prevent. It would be helpful if we can provide more visibility of how malicious transactions actually undermine a device. Even though this cannot stop the single device from being attacked, we can still utilize the collected statistics to develop in-time security patches, benefiting the vast majority of end-users. According to our experience in bug-hunting process (in Section 6), improvements on the following two aspects can effectively increase the visibility of Binder transactions.

Construct IPC Graph. In the case of an attack on the Binder interface, the victim is the recipient of the transaction while the attacker is the sender. Under most circumstances, there will only be visible consequences (i.e., crashes) on the server side while the attacker stays in the mist. This makes it difficult to conduct attack attribution, especially when an attack is mounted via a chain of transactions across multiple processes. To enhance the visibility of attacks in this process, we can instrument each system service to maintain the senders of recent transactions. Each transaction represents an edge in the IPC graph, linking apps and system services. Any user-level app that is linkable to the victim system service is a potential initiator of the attack. Similar techniques have been proven to be effective in improving the visibility of remote systems [8].

Maintain Transaction Schema. Even if we know which transaction causes a crash, it is often challenging to identify the corresponding vulnerability, especially when the bug is non-trivial. The most frequent obstacle is lack of visibility of the transaction schema. A transaction contains only raw bytes and may be unmarshalled in any arbitrary way. It is extremely tedious to anatomize a transaction and verify that it actually contains an invalid parameter causing a specific vulnerability. We propose to maintain transaction schema at runtime when the system service parses a transaction (similar to Fig. 6). If the server experiences some exception, the recorded schema will be attached with the crash report to provide more informative feedback.

8 Discussion

Our work has lasting values beyond the vulnerabilities we presented in this paper. First, we comprehensively assessed a risky attack surface that has long been over-

looked by the system developers of Android. As our experimental results demonstrated, new vulnerabilities are still emerging on this attack surface and *BinderCracker* can help eliminate potential vulnerabilities in future releases of Android.

Second, the lessons learned can transcend to other platforms facing similar issues, such as vehicular systems (CAN buses and ECUs), wearable devices, etc. We highlight that, although many systems adopt a client-server model in the design of their internal system components, they rarely follow the security standards of a real client-server model as in a networked environment. In many scenarios, a component may fall into the wrong hands and create serious security threats.

Third, our parameter-aware fuzzing is generic and not limited to system services. In fact, it also works for services exported by user-level apps. For example, apps like Facebook also host service in its own process space and export it to other apps. By performing fuzzing on this interface, more app-level vulnerabilities are expected to be unearthed. We didn't discuss it here mainly because source codes of app-level services are mostly unavailable. Therefore, it is difficult for us to analyze the root causes and security implications of the identified vulnerabilities.

For two reasons, we test each firmware in its initial state: right after it has been flushed, without installing third-party apps, inserting SIM cards, or connecting to WiFi. First, we want to exclude vulnerabilities caused by external factors that may not be reproducible. Second, we want to rule out the possibility that our fuzzing tests would negatively affect cellular providers or Internet services, for the sake of responsible research. All of the vulnerabilities we identified have been reported to AOSP. Part of the vulnerabilities have been accepted and will be patched in future versions of Android, while the rest are still under review. In this paper, we have only revealed details about the vulnerabilities that have been confirmed so far.

9 Conclusion

In this paper, we have conducted a field study accessing the robustness of Android system services. Specifically, we have designed and implemented *BinderCracker*, an automatic testing framework that can help expose vulnerable system services by fuzzing the Binder interface. *BinderCracker* supports parameter-aware fuzzing and identified more than 100 vulnerabilities in 6 major versions of Android. We summarized these vulnerabilities, explained their security implications and analyzed their root causes. Based on our observation, we highlighted the deficiency of testing only on client-side public APIs and advocated testing and protection at the Binder

interface — the actual security boundary. Several potential defenses as well as their practicality have been discussed to help eliminate vulnerabilities as early as possible in the development cycle.

References

- [1] AMALFITANO, D., FASOLINO, A. R., AND TRAMONTANA, P. A gui crawling-based technique for android mobile application testing. In *Software Testing, Verification and Validation Workshops (ICSTW), 2011 IEEE Fourth International Conference on* (2011), IEEE, pp. 252–261.
- [2] AMALFITANO, D., FASOLINO, A. R., TRAMONTANA, P., DE CARMINE, S., AND MEMON, A. M. Using gui ripping for automated testing of android applications. In *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering* (2012), ACM, pp. 258–261.
- [3] Android interface definition language (aidl). <http://developer.android.com/guide/components/aidl.html>.
- [4] BUGIEL, S., DAVI, L., DMITRIENKO, A., FISCHER, T., AND SADEGHI, A.-R. Xmandroid: A new android evolution to mitigate privilege escalation attacks.
- [5] CHIN, E., FELT, A. P., GREENWOOD, K., AND WAGNER, D. Analyzing inter-application communication in android. In *Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2011), MobiSys '11, ACM, pp. 239–252.
- [6] Cve-2015-1528. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-1528>.
- [7] Cve-2015-1528. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-1474>.
- [8] DIETZ, M., SHEKHAR, S., PISETSKY, Y., SHU, A., AND WALLACH, D. S. Quire: Lightweight provenance for smart phone operating systems. In *USENIX Security Symposium* (2011), vol. 31.
- [9] ELISH, K. O., YAO, D., AND RYDER, B. G. On the need of precise inter-app icc classification for detecting android malware collusions. In *Proceedings of IEEE Mobile Security Technologies (MoST), in conjunction with the IEEE Symposium on Security and Privacy* (2015).
- [10] ENCK, W., GILBERT, P., CHUN, B.-G., COX, L. P., JUNG, J., MCDANIEL, P., AND SHETH, A. N. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation* (Berkeley, CA, USA, 2010), OSDI'10, USENIX Association, pp. 1–6.
- [11] ENCK, W., OCTEAU, D., MCDANIEL, P., AND CHAUDHURI, S. A study of android application security. In *Proceedings of the 20th USENIX Conference on Security* (Berkeley, CA, USA, 2011), SEC'11, USENIX Association, pp. 21–21.
- [12] FELT, A. P., CHIN, E., HANNA, S., SONG, D., AND WAGNER, D. Android permissions demystified. In *Proceedings of the 18th ACM Conference on Computer and Communications Security* (New York, NY, USA, 2011), CCS '11, ACM, pp. 627–638.
- [13] FELT, A. P., HA, E., EGELMAN, S., HANEY, A., CHIN, E., AND WAGNER, D. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security* (New York, NY, USA, 2012), SOUPS '12, ACM, pp. 3:1–3:14.
- [14] Fuzzing android system services by binder call. <https://www.blackhat.com/docs/us-15/materials/us-15-Gong-Fuzzing-Android-System-Services-By-Binder-Call-To-Escalate-Privilege.pdf>.

- [15] Google says there are now 1.4 billion active android devices worldwide. <http://www.androidcentral.com/google-says-there-are-now-14-billion-active-android-devices-worldwide>.
- [16] HORNYACK, P., HAN, S., JUNG, J., SCHECHTER, S., AND WETHERALL, D. These aren't the droids you're looking for: Retrofitting android to protect data from imperious applications. In *Proceedings of the 18th ACM Conference on Computer and Communications Security* (New York, NY, USA, 2011), CCS '11, ACM, pp. 639–652.
- [17] HU, C., AND NEAMTIU, I. Automating gui testing for android applications. In *Proceedings of the 6th International Workshop on Automation of Software Test* (2011), ACM, pp. 77–83.
- [18] Improving your code with lint. <http://developer.android.com/tools/debugging/improving-w-lint.html>.
- [19] KANTOLA, D., CHIN, E., HE, W., AND WAGNER, D. Reducing attack surfaces for intra-application communication in android. In *Proceedings of the Second ACM Workshop on Security and Privacy in Smartphones and Mobile Devices* (New York, NY, USA, 2012), SPSM '12, ACM, pp. 69–80.
- [20] LEE, B., LU, L., WANG, T., KIM, T., AND LEE, W. From zygote to morula: Fortifying weakened aslr on android. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy* (Washington, DC, USA, 2014), SP '14, IEEE Computer Society, pp. 424–439.
- [21] MACHIRY, A., TAHILIANI, R., AND NAIK, M. Dynodroid: An input generation system for android apps. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering* (2013), ACM, pp. 224–234.
- [22] MAJI, A. K., ARSHAD, F. A., BAGCHI, S., AND RELLERMEYER, J. S. An empirical study of the robustness of inter-component communication in android. In *Proceedings of the 2012 42Nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (Washington, DC, USA, 2012), DSN '12, IEEE Computer Society, pp. 1–12.
- [23] MARFORIO, C., FRANCILLON, A., CAPKUN, S., CAPKUN, S., AND CAPKUN, S. *Application collusion attack on the permission-based security model and its implications for modern smartphone systems*. Department of Computer Science, ETH Zurich, 2011.
- [24] MCDONNELL, T., RAY, B., AND KIM, M. An empirical study of api stability and adoption in the android ecosystem. In *Software Maintenance (ICSM), 2013 29th IEEE International Conference on* (2013), IEEE, pp. 70–79.
- [25] NAUMAN, M., KHAN, S., AND ZHANG, X. Apex: Extending android permission model and enforcement with user-defined runtime constraints. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security* (New York, NY, USA, 2010), ASIACCS '10, ACM, pp. 328–332.
- [26] Number of apps available in leading app stores as of july 2015. <http://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/>.
- [27] PELES, O., AND HAY, R. One class to rule them all: 0-day deserialization vulnerabilities in android. In *9th USENIX Workshop on Offensive Technologies (WOOT 15)* (2015).
- [28] RANGWALA, M., ZHANG, P., ZOU, X., AND LI, F. A taxonomy of privilege escalation attacks in android applications. *Int. J. Secur. Netw.* 9, 1 (Feb. 2014), 40–55.
- [29] SASNAUSKAS, R., AND REGEHR, J. Intent fuzzer: Crafting intents of death. In *Proceedings of the 2014 Joint International Workshop on Dynamic Analysis (WODA) and Software and System Performance Testing, Debugging, and Analytics (PERTEA)* (New York, NY, USA, 2014), WODA+PERTEA 2014, ACM, pp. 1–5.
- [30] Security updates and resources. <https://source.android.com/security/overview/updates-resources.html>.
- [31] SHABTAI, A., FLEDEL, Y., AND ELOVICI, Y. Securing android-powered mobile devices using selinux. *IEEE Security & Privacy*, 3 (2009), 36–44.
- [32] SHABTAI, A., FLEDEL, Y., KANONOV, U., ELOVICI, Y., DOLEV, S., AND GLEZER, C. Google android: A comprehensive security assessment. *IEEE Security and Privacy* 8, 2 (Mar. 2010), 35–44.
- [33] SMALLEY, S., AND CRAIG, R. Security enhanced (se) android: Bringing flexible mac to android. In *NDSS* (2013), vol. 310, pp. 20–38.
- [34] Smartphone os market share, q4 2014. <http://www.idc.com/prodserv/smartphone-os-market-share.jsp>.
- [35] XU, R., SAÏDI, H., AND ANDERSON, R. Aurasium: Practical policy enforcement for android applications. In *Proceedings of the 21st USENIX Conference on Security Symposium* (Berkeley, CA, USA, 2012), Security'12, USENIX Association, pp. 27–27.
- [36] YE, H., CHENG, S., ZHANG, L., AND JIANG, F. Droidfuzzer: Fuzzing the android apps with intent-filter tag. In *Proceedings of International Conference on Advances in Mobile Computing & Multimedia* (New York, NY, USA, 2013), MoMM '13, ACM, pp. 68:68–68:74.