

DRAFT - The Woes of Android Reverse Engineering: from Large Scale Analysis to Dynamic Deobfuscation

Current version: commit 5497988

List of Notes

<input type="checkbox"/>	operation	18
<input type="checkbox"/>	I would have said <code>operating</code> but <code>grammarly</code> disagrees	18
<input type="checkbox"/>	to have their abuser install	18
<input type="checkbox"/>	<code>jfl</code> says <code>install ing</code> , <code>jm</code> says no, <code>grammarly</code> is on the side of <code>jm</code>	18
<input type="checkbox"/>	of	19
<input type="checkbox"/>	<code>jfl</code> asks <code>in ?</code> <code>grammarly</code> says <code>of</code>	19
<input type="checkbox"/>	Reflection is another common obfuscation technique against static analysis. Instead of directly invoking methods, the generic <code>Method.invoke()</code> API is used, and the method is retrieved from its name in the form of a character string. Finding the value of this string can be quite difficult to determine statically, so it is once again an issue more suitable for dynamic analysis. When encountering a complex case of reflection (i.e., using ciphered strings) or code loading, a reverse engineer will switch to dynamic analysis to collect the relevant data (the name of the methods called or the code that was loaded), then switch back to static analysis. This is doable for a manual analysis; unfortunately, the more complex tools that would require that runtime information to perform an accurate analysis may not have a way to access this new data.	20
<input type="checkbox"/>	Peu développé. Expliquer qu'un reverser, s'il trouve de la reflection ou du dyn load peut éventuellement capturer les données en analyse dynamique. Mais ensuite ces données deviennent inutiles s'il retourne à de l'analyse static. En effet, il fait souvent les deux en alternances. Il avait besoin que les data issues de l'analyse dyn soient prises en compte par l'analyse statique, par exemple ... TODO : trouver un exemple simple à formuler	20
<input type="checkbox"/>	TODO: Ça serait bien de souligner Dyn Code Load et Reflection	33
<input type="checkbox"/>	We believe that the effort of reviewing the literature for making a comprehensive overview of available approaches should be pushed further: an existing published approach with a software that cannot be used for technical reasons endangers both the reproducibility and reusability of research.	35
<input type="checkbox"/>	A mettre en avant?	36
<input type="checkbox"/>	to see if they can be used as an indication as to which tools can still be used today.	36
<input type="checkbox"/>	Mettre en avant	36
<input type="checkbox"/>	TODO: <code>cl-shadow</code>	81
<input type="checkbox"/>	TODO: small intro	87
<input type="checkbox"/>	Maybe talk about v41 in RASTA? this will break a lot of things	91
<input type="checkbox"/>	TODO: Ça serait bien de faire un PR ou deux à <code>Jadx/Androguard/Soot</code> quand même	91
<input type="checkbox"/>	Ici je pensais lire comment on transforme le code qui load du code, mais on me parle de multi dex	99
<input type="checkbox"/>	Un exemple aiderait à comprendre <code>jm</code> : <code>j</code> en ai pas qui prennent pas 3 pages de listing	100

- TODO: Comment on dit proprement que cest tout pété? 106
- TODO: Flowdroid results are inconclusive: some apks have more leak after and as many apks have less? also, runing flowdroid on the same apk can return a different number of leak??? 111



THÈSE DE DOCTORAT DE

CENTRALESUPÉLEC

ÉCOLE DOCTORALE N° 601

*Mathématiques, Télécommunications, Informatique,
Signal, Systèmes, Électronique*
Spécialité : *Informatique*

Par

Jean-Marie MINEAU

The Woes of Android Reverse Engineering: from Large Scale Analysis to Dynamic Deobfuscation

Les difficultés de la rétro-ingénierie Android: de l'analyse large échelle au dé-brouillage dynamique

Thèse présentée et soutenue à Rennes, le 2025-12-09

Unité de recherche : IRISA

Thèse N°: 2025CSUP0006

Composition du jury :

Président :	Olivier Barais	Professeur des Universités	Université de Rennes
Rapporteurs :	Vincent Nicomette	Professeur des Universités	INSA de Toulouse
	Julien Signoles	Directeur de Recherche	CEA LIST
Examineurs :	Simone Aonzo	Maître de Conférences	Eurecom
Dir. de thèse :	Jean-François Lalande	Professeur des Universités	CentraleSupélec
	Valérie Viet Triem Tong	Professeure	CentraleSupélec

ACKNOWLEDGEMENTS

1

2 First of all, I would like to express my gratitude to Vincent Nicomette and Julien Signoles, who
3 agreed to be the *rapporteurs* for this thesis, as well as to Olivier Barais and Simone Aonzo for
4 being part of the Jury.

5 I'm also grateful for the support of my thesis advisors, Jean-François Lalande and Valérie Viet
6 Triem Tong. They trusted me to explore my strange ideas, reminded me to take things one at
7 a time when I was overwhelmed by the number of issues one encounters when working with
8 Android, and overall gave me great advice during those three years. Now that I have finished
9 my PhD, I can confirm that the curve of the evolution of the morale of the PhD student Jean-
10 François drew me is pretty accurate.

11 I want to thank Marie Babel and Jacque Klein for being part of my *Comité de Suivi Individuel*,
12 and my colleagues of the 5th floor, from both the PIRAT\('); and the SUSHI team. Those 3
13 years were fun, the croissants were good, and the conversations were weird. What more could
14 you ask for? And I wish courage and good luck to those on the tail-end of their PhD, you're
15 almost there!

16 As promised, I thank Amelie for proofreading my French summary. I forgive you for the typos
17 you missed. Not as promised, I also thank you for the moral support and all the cute animal
18 pictures exchanged during the last 3 years.

19 I also want to thank the *Pains-Perdus*, some of you may or may not have pushed me into doing
20 a PhD, and you managed to endure my many rants about Android.

21 Last but not least, I thank my family. Antoine for the good food every time I'm in the area.
22 Kpu for all the sword fighting when you were in Rennes, now that you have the HEMA virus,
23 my job is done. My mother, for giving me my first taste of infosec, even if it was definitely not
24 on purpose. Bypassing parental control on a wifi router is a great way to start. And my father,
25 for showing a 12-year-old me how to read and modify the source code of Battle for Wesnoth,
26 who knew this first step in the world of reverse engineering would lead to a PhD?

27 Thank you.

TABLE OF CONTENTS

29	1	Introduction	1
30	2	Background and Motivation	7
31	2.1	Introduction	7
32	2.2	Android Background	8
33	2.3	Problems of the Reverse Engineer	18
34	2.4	State of the Art	19
35	2.5	Conclusion	28
36	3	Evaluating the Reusability of Android Static Analysis Tools	29
37	3.1	Introduction	29
38	3.2	Methodology	30
39	3.3	Experiments	36
40	3.4	Failures Analysis	42
41	3.5	State-of-the-Art Comparison	46
42	3.6	Recommendations	47
43	3.7	Limitations	48
44	3.8	Future Works	49
45	3.9	Conclusion	50
46	4	Class Loaders in the Middle: Confusing Android Static Analysers	53
47	4.1	Introduction	53
48	4.2	Analysing the Class Loading Process	54
49	4.3	Obfuscation Techniques	60
50	4.4	Shadow Attacks in the Wild	65
51	4.5	Discussion	71
52	4.6	Conclusion	75
53	5	The Application of Theseus: After Adding Runtime Data, it is Still Your Application ..	77
54	5.1	Introduction	78
55	5.2	Overview	78
56	5.3	Code Transformation	79
57	5.4	Collecting Runtime Information	87
58	5.5	Results	90
59	5.6	Limitations and Future Works	99
60	5.7	Conclusion	102

61	6	Conclusion	103
62	6.1	Contributions of this Thesis	103
63	6.2	Perspectives for Future Work	104

INDEX OF FIGURES

65	Figure 1: Source code for a simple Java method and its Call and Control Flow Graphs ...	15
66	Figure 2: Tool selection methodology overview	34
67	Figure 3: Experiment methodology overview	34
68	Figure 4: Exit status for the Drebin dataset	36
69	Figure 5: Exit status for the RASTA dataset	36
70	Figure 6: Exit status evolution for the RASTA dataset	38
71	Figure 7: Finishing rate by bytecode size for APK detected in 2022	39
72	Figure 8: Finishing rate by discovery year with a bytecode size $\in [4.08, 5.2]$ MB	39
73	Figure 9: Finishing rate by min SDK with a bytecode size $\in [4.08, 5.2]$ MB	40
74	Figure 10: Exit status comparing goodware (left bars) and malware (right bars) for the RASTA	
75	dataset	41
76	Figure 11: Heatmap of the ratio of error reasons for all tools for the RASTA dataset	43
77	Figure 12: The class loading hierarchy of Android	55
78	Figure 13: Location of SDK classes during development and at runtime	57
79	Figure 14: Call Graphs of an application calling <code>Main.bad()</code> from a shadowed <code>Obfuscation</code>	
80	class	64
81	Figure 15: Redefined SDK classes, sorted by the first SDK they appeared in	66
82	Figure 16: Process to add runtime information to an APK	79
83	Figure 17: Inserting DEX files inside an APK	83
84	Figure 18: Exit status of static analysis tools on original APKs (left) and patched APKs	
85	(right)	94
86	Figure 19: Call Graph of <code>Main.main()</code> generated by Androguard before patching	97
87	Figure 20: Call Graph of <code>Main.main()</code> generated by Androguard after patching	97

INDEX OF TABLES

89	Table 1: Considered tools [33]: availability and usage reliability	30
90	Table 2: Selected tools, forks, selected commits and running environment	32
91	Table 3: Average size and date of goodware/malware parts of the RASTA dataset	40
92	Table 4: DEX size and Finishing Rate (FR) per decile	42
93	Table 5: Average number of errors, analysis time, memory per unitary analysis – compared by	
94	exit status	42
95	Table 6: Commonly found dependencies	50
96	Table 7: Comparison of API methods between documentation and emulators	59
97	Table 8: Working attacks against static analysis tools	62
98	Table 9: Shadow classes compared to SDK 34 for a dataset of 49 975 applications	65
99	Table 10: Shadow classes compared to SDK 34 for a dataset of 49 975 applications	69
100	Table 11: Summary of the dynamic exploration of the applications from the RASTA dataset	
101	collected by Androzoo in 2023	91
102	Table 12: Most common dynamically loaded files	92
103	Table 13: Edges added to the call graphs computed by Androguard by instrumenting the	
104	applications	93
105	Table 14: Average time and memory consumption of Soot, Apktool and Androscalpel	98

INDEX OF LISTINGS

107	Listing 1: Class instantiation	55
108	Listing 2: Default Class Loading Algorithm for Android Applications	56
109	Listing 3: The method generating the .dex filenames from the AOSP	60
110	Listing 4: Main body of test apps	61
111	Listing 5: Implementation of Reflection found in <code>classes11.dex</code> (shadows Listing 6)	69
112	Listing 6: Implementation of Reflection executed by ART (shadowed by Listing 5)	69
113	Listing 7: Instantiating a class using <code>Class.newInstance()</code>	80
114	Listing 8: Instantiating a class using <code>Constructor.newInstance(...)</code>	80
115	Listing 9: Calling a method using reflection	80
116	Listing 10: A reflection call that can call any method	80
117	Listing 11: Listing 9 after the de-reflection transformation	81
118	Listing 12: Pseudo-code of the renaming algorithm	85
119	Listing 13: Code of the main class of the application, as shown by Jadx, before patching .	95
120	Listing 14: Code of <code>Main.main()</code> , as shown by Jadx, after patching	96

LIST OF ACRONYMS AND NOTATIONS

122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150

Acronyms	Meanings
ADB	Android Debug Bridge, a tool to connect to an Android emulator of smartphone to run commands, start applications, send events and perform other operations for testing and debugging purpose
AOSP	Android Open Source Project, the project hosting the most of the Android operating system source code
API	Application Programming Interface, in the Android ecosystem, it is a set of classes with known method signatures that can be called by an application to interact with the Android framework
APK	Android Package, the file format used to install application on Android. The APK format is an extension of the JAR format
ART	Android RunTime, the runtime environment that execute an Android application. The ART is the successor of the older Dalvik Virtual Machine
AXML	Android XML. The specific flavor of XML used by Android. The main specificity of AXML is that it can be compile in a binary version inside an APK
CFG	Control-Flow Graph, a graph representing the control structures (<i>e.g.</i> “if” blocks) of the code of a method/application
DEX	Dalvik Executable, the file format for the bytecode used for applications by Android
DFG	Data-Flow Graph, a graph representing the flow of information in an application
FR	Finishing Rate, the number of runs that finished over the number of total runs of an analysis
HPC	High-Performance Computing, the use of supercomputers and computer clusters
MWE	Minimum Working Example, in this context, a small example that can be used to check if a tool is working
IDE	Integrated Development Environment, a software providing tools for software development
JAR	Java ARchive file, the file format used to store several java class files. Sometimes used by Android to store DEX files instead of java classes

151
152
153
154
155
156
157
158
159
160
161
162

Acronyms	Meanings
JNI	Java Native Interface, the native library used to interact with Java classes of the application and Android API
OAT	Of Ahead Time, an ahead of time compiled format for DEX files
NDK	Native Development Kit, the set of tools used to build C and C++ code for Android
SDK	Software Development Kit, a set of tools for developing software targeting a specific platform. In the context of Android, the version of the SDK can be associated to a version of Android, and application compatibility is defined in term of compatible SDK version
XML	eXtensible Markup Language, a language to store data
ZIP	ZIP is an archive format. A ZIP file contains other files, that may be compressed

188 in the store for malicious software¹. Although its **operation** is kept secret, it seems that
189 the Bouncer is both comparing the applications with known malware code and running the
190 applications in Google's cloud infrastructure to detect hidden behaviour. Despite Google's
191 efforts, malicious applications are still found in the Play Store [2]. Also, it is not uncommon for
192 people in abusive situations **to have their abuser install** on their phone a stalkerware (spying
193 application) found outside of the Play Store [20].

I would have said "operating" but grammatically disagrees

194 For these reasons, it is important to be able to analyse an application and understand what
195 it does. This process is called reverse engineering. A lot of work has been done to reverse
196 engineer computer software, but Android applications come with specific challenges that need
197 to be addressed. For instance, Android applications are distributed in a specific file format,
198 the APK format, and the code of the application is mainly compiled into an Android-specific
199 bytecode: Dalvik. An Android reverse engineer will need tools that can read those Android-
200 specific formats. A first test in the process of reverse engineering an application would be to
201 simply read the content of the application and the code in it. Tools like Apktool can be used to
202 convert the binary files of an application into a human-readable format. Other tools like Jadx can
203 go further and try to generate Java code from the bytecode in the application. Because Android
204 applications tend to be quite large, it can be quite tedious to understand what it does just from
205 reading their bytecode. To address this issue, many tools/approaches have been developed [33,
206 61] to extract higher-level information about the behaviour of the application without having to
207 manually analyse the application. For example, Flowdroid [8] aims to detect information leaks:
208 given a set of methods that can generate private information, and a set of methods that send
209 information to the outside, Flowdroid will detect if private information is sent to the outside.
210 Once again, those kinds of tools need to target Android specifically. Android runs its applications
211 code differently than a computer would run software. One example would be the handling of
212 entry points: computer software usually has one entry point, whereas Android applications have
213 many, and Android will choose depending on context. Unfortunately, those tools are hard to
214 use, and even when they work on small example applications, it is not uncommon for them to
215 fail to run on real-life applications [53]. This is worrying. Android applications are becoming
216 more complex every year, and tools that cannot handle this complexity will fail more often.
217 This leads us to our first problem statement:
218

jfl says "installing", jm says no, grammatically is on the side of jm

219 1. <https://googlemobile.blogspot.com/2012/02/android-and-security.html>

220 **Pb1:** *To what extent are previously published Android analysis tools still usable today, and*
 221 *what factors impact their reusability?*

222 Many tools have been published to analyse Android applications, but the Android ecosystem
 223 is evolving rapidly. Tools developed 5 years ago might not be usable anymore. We will
 224 endeavour to identify which tools are still usable today, and for the others, what causes
 225 them to no longer be an option.

226 Another issue is that Android application developers sometimes use various techniques to slow
 227 down reverse engineering. This process is called obfuscation. Malware developers do that to
 228 hide malicious behaviour and avoid detection, but the use of obfuscation is no proof that an
 229 application is malicious. Indeed, legitimate application developers can also use obfuscation to
 230 protect their intellectual property. Thus, developers and reverse engineers are playing a game
 231 of cat and mouse, constantly inventing new techniques to hide or reveal the behaviour of an
 232 application.

233 There are two types of reverse engineering techniques: static and dynamic. Static analysis
 234 consists of examining the application without running it, while dynamic analysis studies

the action of the application while it is running. Both methods have their drawbacks, and
 techniques will often capitalise on the drawbacks of one of those methods. For instance, an
 application can try to detect if it is running in a sandbox environment and not act maliciously
 if it is the case. Similarly, an application can dynamically load bytecode at runtime, and this
 bytecode will not be available during a static analysis. Dynamic code loading relies on Java
 classes called `ClassLoader` that are central components of the Android runtime environment.
 Because dynamic code loading is such a difficult problem for static analysis, dynamic class
 loading is often ignored when doing static analysis. However, class loading is not limited to
 dynamic code loading. As a matter of fact, the Android Runtime is constantly performing class
 loading to load classes from the application or from the Android platform itself. This blind spot
 in static analysis tools raises our second problem statement:

246 **Pb2:** *What is the default Android class loading algorithm, and does it impact static analysis?*

247 Class loading is an operation often ignored by static analysis tools. The exact algorithm
 248 used is not well known and might not be accurately modelled by static analysis tools. If it
 249 is the case, discrepancies between the model of the tools and the one used by Android could
 250 be used as a base for new obfuscation techniques.

jfl asks
 “in”? 236
 grammatically
 says “of” 238

Reflection is another common obfuscation technique against static analysis. Instead of directly invoking methods, the generic `Method.invoke()` API is used, and the method is retrieved from its name in the form of a character string. Finding the value of this string can be quite difficult to determine statically, so it is once again an issue more suitable for dynamic analysis. When encountering a complex case of reflection (*i.e.*, using ciphered strings) or code loading, a reverse engineer will switch to dynamic analysis to collect the relevant data (the name of the methods called or the code that was loaded), then switch back to static analysis. This is doable for a manual analysis; unfortunately, the more complex tools that would require that runtime information to perform an accurate analysis may not have a way to access this new data.

Some contributions made the results they computed available to other tools by modifying the application (instrumenting) in a way that reflects those results. This led us to our last problem statement:

Pb3: *Can we use instrumentation to provide dynamic code loading and reflection data collected dynamically to static analysis tools and improve their results?*

Dynamic code loading and reflection are problems most suited for dynamic analysis. However, static analysis tools do not have access to collected data. Encoding this information inside valid applications could be a way to make it universally available to any static analysis tool. Ideally, this encoding should not degrade the quality of the static analysis compared to the original application.

Contributions

The contributions of this thesis are the following:

1. We evaluate the reusability of Android static analysis tools published by the community: we rebuild static analysis tools in their original environment as container images. With those containers, those tools are now readily available on any environment capable of running either Docker or Singularity. We tested those tools on a dataset of real-life applications balanced in order to have a significant number of applications with different characteristics to assess which characteristics impact the success of a tool. This work was presented at the 21st International Conference on Software and Systems Reuse (ICSR 2024) conference [44].
2. We model the default class loading behaviour of Android. Based on this model, we define a class of obfuscation techniques that we call *shadow attacks* where a class definition in an APK shadows the actual class definition. We show that common state-of-the-art tools like Jadx or Flowdroid do not implement this model correctly and thus can fall for those shadow attacks. We analysed a large number of recent Android applications and found that

- 303 applications with class shadowing do exist, though they are the result of quirks in the APK
304 compilation process and not deliberate obfuscation attempts. This work was published in
305 the Digital Threats journal [45].
- 306 3. We propose an approach to allow static analysis tools to analyse applications that perform
307 dynamic code loading: We collect at runtime the bytecode dynamically loaded and the
308 reflection calls information, and patch the APK file to perform those operations statically.
309 Finally, we evaluate the impact this transformation has on the tools we containerised
310 previously.
- 311 4. We released under the GPL licence the software we used in the experiments presented in
312 this thesis. For Chapter 3, this includes the code used to test the output of each tool and the
313 code to analyse the results of the experiment, in addition to the containers to run the tested
314 tools. We also released Androscalpel, a Rust crate to manipulate Dalvik bytecode, that we
315 used to create Theseus, a set of scripts that implement the approach presented in Chapter 5.
316 The complete list and location of the software we release are available in Appendix A.

317 **Outline**

318 This dissertation is composed of 6 chapters. This introduction is the first chapter. It is followed
319 by Chapter 2, which gives background information about Android and the various analysis
320 techniques targeting Android applications.

321 The next three chapters are dedicated to the contributions of this thesis. First Chapter 3
322 studies the reusability of static analysis tools. Next, in Chapter 4, we model the default class
323 loading algorithm used by Android and show the consequences for reverse engineering tools
324 that implement the wrong model. Then Chapter 5 presents an approach that allows for static
325 analysis tools to analyse applications that load bytecode at runtime.

326 Finally, Chapter 6 summarises the contributions of this thesis and opens perspectives for future
327 work.

BACKGROUND AND MOTIVATION

This is a Unix system. I know this.

— Alexis "Lex" Murphy, Jurassic Park

332 2.1 Introduction

In order to understand the challenges of reverse engineering Android applications, we first need to understand some key concepts and specificities of Android. In particular, the format in which applications are distributed, as well as the runtime environment that runs those applications, are very specific to Android. To handle those specificities, a reverse engineer must have appropriate tools. Some of those tools are used recurrently, either by the reverse engineer themselves, or as a basis for other more complex tools that implement more advanced analysis techniques.

Among those techniques, the ones that do not require running the application are called static analysis. Over time, many of those tools have been released. To compare those different tools, different benchmarks have been proposed, highlighting different strengths and weaknesses of each tool.

Unfortunately, static analysis has its limits. One such limit is that it cannot analyse what is not inside the application. Platform classes are classes that are present directly on the smartphone, and not in the application. Some of those classes are well known and taken into account by analysis tools, but the rest of those classes, often called *hidden API*, are not. In addition to platform classes, classes that are loaded dynamically (*i.e.*, at runtime) are also not always available to static analysis. This led static analysis tools to disregard the class loading process altogether, leaving the subject relatively unexplored.

When static analysis fails, for instance, because of dynamic class loading, the reverse engineer will fall back on dynamic analysis. Dynamic analysis is the counterpart of static analysis: it is based on the analysis of the execution of the application. Depending on the context, the reverse engineer will then alternate between different techniques, using previous results to improve the

354 next iteration. Regrettably, analysis tools mostly return results in an ad hoc format, making it
355 difficult to make other tools aware of the retrieved information. Some tools, however, encode
356 their result in the form of a new augmented Android application. The idea being that any
357 Android analysis tools must be able to handle an Android application in the first place, so they
358 will have access to that new information.

359 We will begin this chapter with a presentation of the bases of the Android ecosystem. The
360 reader already familiar with Android reverse engineering might want to skip Section 2.2 and
361 directly read Section 2.3, where we put our problem statements in perspective. We will then
362 examine the state of the art related to those problem statements in Section 2.4, and conclude
363 this chapter in Section 2.5.

364 2.2 Android Background

365 We begin this chapter with background information about Android and the reverse engineering
366 of Android applications. We start with a description of Android applications and their execution
367 environment, then list some useful basic tools for reverse engineering, and finish with the basics
368 of static analysis for Android.

369 2.2.1 Android

370 Android is the smartphone operating system developed by Google. It is based on a Long Term
371 Support Linux Kernel, to which patches developed by the Android community are added. On
372 top of the kernel, Android redeveloped many of the usual components used by Linux-based
373 operating systems, like the init system or the standard C library, and added new ones, like the
374 ART that executes the applications. Those changes make Android a unique operating system.

375 2.2.1.1 Android Applications

376 Applications in the Android ecosystem are distributed in the APK format. APK files are JAR
377 files with additional features, which are themselves ZIP files with additional features.

378 A minimal APK file contains a file `AndroidManifest.xml`, the `META-INF/` folder containing the
379 JAR manifest and signature files, and an APK Signing Block at the end of the ZIP file. The
380 code of the application is then stored in a custom format, the Dalvik bytecode, or in the binary
381 ELF format, called native code in the Android ecosystem, or both. Dalvik bytecode is stored
382 in the `classes.dex`, `classes2.dex`, `classes3.dex`, ... while native code is stored in `lib/<arch>/`
383 `*.so`. The `res/` folder contains the resources required for the user interface. When resources
384 are present in `res/`, the file `resources.arsc` is also present at the root of the archive. The
385 `assets/` folder contains the files that are used directly by the code application. Depending on
386 the application and compilation process, any kind of other files and folders can be added to the
387 application.

388 **Signature** Android applications are cryptographically signed to prove the authorship. Appli-
389 cations signed with the same key are considered developed by the same entity. This allows
390 updating the applications securely, and applications can declare security permissions to restrict
391 access to some features to only applications with the same author.

392 Android has several signature schemes coexisting:

- 393 • The v1 signature scheme is the JAR signing scheme, where the signature data is stored in
394 the `META-INF/` folder.
- 395 • The v2, v3 and v3.1 signature scheme are store in the ‘APK Signing Block’ of the APK.
396 The v2 signature scheme was introduced in Android 7.0, and to keep retro-compatibility
397 with older versions, the v1 scheme is still used in addition to the APK Signing Block. The
398 Signing block is an unindexed binary section added to the ZIP file, between the ZIP entries
399 and the Central Directory. The signature was added in an unindexed section of the ZIP to
400 avoid interfering with the v1 signature scheme that signed the files inside the archive, and
401 not the archive itself.
- 402 • The v4 signature scheme is complementary to the v2/v3 signature scheme. Signature data
403 are stored in an external, `.apk.idsig` file.

404 **Android Manifest** The Android Manifest is stored in the `AndroidManifest.xml`, encoded in
405 the binary AXML format. The manifest declares important information about the application:

- 406 • Generic information like the application name, ID and icon.
- 407 • The Android compatibility of the applications, in the form of 3 values: the Android `min-`
408 `sdk`, `target-sdk` and `max-sdk`. Those are the minimum, targeted and maximum versions of
409 the Android SDK supported by the application.
- 410 • The application components (Activity, Service, Receiver and Provider) of the application
411 and their associated classes.
- 412 • Intent filters to list the intents that can start or be sent to the application components.
- 413 • Security permissions required by the application.

414 **Code** An application usually contains at least a `classes.dex` file containing Dalvik bytecode.
415 This is the format executed by the Android ART. It is common for an application to have
416 more than one DEX file when an application needs to reference more methods than the format
417 allows in one file (each method referenced inside a DEX is associated with a 16-bits number,
418 limiting their number to 65 536). Support for multiple DEX files was added in the SDK 21
419 version of Android, and applications that have multiple DEX files are sometimes referred to as
420 ‘multi-dex’.

421 In addition to DEX files, and sometimes instead of DEX files, applications can contain `.so`
422 ELF (Executable and Linkable Format) files in the `lib/` folder. In the Android ecosystem,
423 binary code is called native code. Because native code is compiled for a specific architecture, `.so`

424 files are present in different versions, stored in different subfolders, depending on the targeted
425 architecture. For example, `lib/arm64-v8a/libexample.so` is the version of the `example` library
426 compiled for an ARM 64 architecture. Because smartphones mostly use ARM processors, it is
427 not rare to see applications that only have the ARM version of their native code.

428 **Resources** Developing graphical interfaces for applications requires many kinds of specific
429 assets, which are stored in `lib/`. Those resources include bitmap images, text, layout, etc. Data
430 like layout, colour or text are stored in binary AXML. An additional file, `resources.arsc`, in a
431 custom binary format, contains a list of the resource names, ids, and their properties.

432 **Compilation Process** For the developer, the compilation process is handled by Android
433 Studio and is mostly transparent. Behind the scenes, Android Studio relies on Gradle to
434 orchestrate the different compilation steps:

435 The sources XML files like `AndroidManifest.xml` and the one in `res/` are compiled to binary
436 AXML by `aapt`, which also generates the resource table `resources.arsc` and a `R.java` file that
437 defines for each resource variables named after the resource, set to the ID of the resource.
438 The `R.java` file allows the developer to refer to resources with readable names and avoid using
439 the often automatically generated resource IDs, which can change from one version of the
440 application to another.

441 The source code is then compiled. The most common programming languages used for Android
442 applications are Java and Kotlin. Both are first compiled to Java bytecode in `.class` files using
443 the language compiler. To allow access to the Android API, the `.class` are linked during the
444 compilation to an `android.jar` file that contains classes with the same signatures as the ones
445 in the Android API for the targeted SDK. The `.class` files are then converted into the DEX
446 format using `d8`. During those steps, both the original language compiler and `d8` can perform
447 optimisations on the classes, like code shrinking, inlining, etc.

448 If the application contains native code, the original C or C++ code is compiled using Android
449 tools from the NDK to target the different possible architectures.

450 `aapt` is then used once again to package all the generated AXML, DEX, `.so` files, as well as the
451 other resource files, assets, `resources.arsc`, and any additional files deemed necessary to form
452 the final ZIP file. `aapt` ensures that the generated ZIP is compatible with the requirements of
453 Android. For instance, the `resources.arsc` will be mapped directly in memory at runtime, so
454 it must not be compressed inside the ZIP file.

455 If necessary, the ZIP file is then aligned using `zipalign`. Again, this is to ensure compatibility
456 with Android optimisations: some files like `resources.arsc` need to be 4-bits aligned to be
457 mapped in memory.

458 The last step is to sign the application using the `apksigner` utility.

459 Since 2021, Google has required that new applications in the Google Play app store be uploaded
460 in a new format called Android App Bundles. The main difference is that Google will perform
461 the last packaging steps and generate (and sign) the application itself. This allows Google to
462 generate different applications for different targets and to avoid including unnecessary files in
463 the application, like native code targeting the wrong architecture.

464 **2.2.1.2 Android Runtime**

465 Android runtime environment has many specificities that set it apart from other platforms.
466 A heavy emphasis is put on isolating the applications from one another as well as from the
467 system's critical capabilities. The code execution itself can be confusing at first. Instead of the
468 usual linear model with a single entry point, applications have many entry points that are called
469 by the Android framework in accordance with external events.

470 ***Application Architecture*** Android applications expose their components to the Android
471 Runtime (ART) via classes inheriting specific classes from the Android SDK. Four classes
472 represent application components that can be used as entry points:

- 473 • *Activities*: An activity represents a single screen with a user interface. This is the component
474 used to interact with a user.
- 475 • *Services*: A service serves as an entry point to run the application in the background.
- 476 • *Broadcast receivers*: A broadcast receiver is an entry point used when a matching event is
477 broadcast by the system.
- 478 • *Content providers*: A content provider is a component that manages data accessible by other
479 applications through the content provider.

480 Components must be listed in the `AndroidManifest.xml` of the application so that the system
481 knows of them. In the life cycle of a component, the system will call specific methods defined
482 by the classes associated with each component type. Those methods are to be overridden by
483 the classes defined in the application if they are specific actions to be performed. For instance,
484 an activity might compute some values in `onCreate()`, called when the activity is created, save
485 the value of those variables to the file system in `onStop()`, called when the activity stop being
486 visible to the user, and recover the saved values in `onRestart()`, called when the user navigate
487 back to the activity.

488 In addition to the components declared in the manifest that act as entry points, the Android
489 API heavily relies on callbacks. The most obvious cases are for the user interface; for example,
490 a button will call a callback method defined by the application when clicked. Other parts of the
491 API also rely on non-linear execution; for example, when an application sends an intent (see
492 next paragraph), the intent sent in response is transmitted back to the application by calling
493 another method.

494 **Application Isolation and Interprocess Communication** On Android, each application
495 has its own storage folders and the application processes are isolated from each other, and
496 from the hardware interfaces. This sandboxing is done using Linux security features like
497 group and user permissions, SELinux, and seccomp. The sandboxing is adjusted according
498 to the permissions requested in the `AndroidManifest.xml` file of the applications. In addition,
499 most features of the Android system can only be accessed through Binder, Android's main
500 interprocess communication channel.

501 Binder is a component of the Android framework, external to the application, that all applica-
502 tions can communicate with. Applications can send messages to Binder, called **intents**. Binder
503 will check if the application is allowed to send it, and then forward it to the appropriate
504 component. This component can then respond with another intent. Applications must declare
505 intent filters to indicate which intent can be sent to the application, and which classes receive
506 the intents. Intents are central to Android applications and are not just used to access Android
507 capabilities. For instance, activities and services are started by receiving intents, and it is not
508 uncommon for an application to self-send intents to switch between activities. Intents can also
509 be sent directly from Android to the application: when a user starts an application by tapping
510 the app icon, Android will send an intent to the class of the application that defined the intent
511 filter for the `android.intent.action.MAIN` intent. One interesting feature of the Binder is that
512 intents do not need to explicitly name the targeted application and class: intents can be implicit
513 and request an action without knowing the exact application that will perform it. An example
514 of this behaviour is when an application wants to open a file: an `android.intent.action.VIEW`
515 intent is sent with the file location and type, and Binder will find and start an application
516 capable of viewing this file.

517 **Platform Classes** In addition to the classes they include, Android applications have access
518 to classes provided by Android, stored on the phone. Those classes are called *platform classes*.
519 They are divided between SDK classes and hidden API. The SDK classes can be seen as the
520 Android standard library. They are documented by Google and have a certain stability from
521 version to version. In case of breaking changes, the changes are listed by Google as well. The
522 list of SDK classes is available at compile time in the form of an `android.jar` file to link against.
523 On the contrary, hidden API are undocumented methods used internally by the ART. Still,
524 they are loaded by the application and can be used by it.

525 **Class Loading and Reflection** Class loading is the mechanism used by Android to find and
526 select the class implementation when encountering a reference to a class. Android developers
527 mainly use it to load bytecode dynamically from a source other than the application itself (*e.g.*,
528 a file downloaded at runtime), using `ClassLoader` objects. `Class` objects are retrieved from those
529 class loaders using their name in the form of strings to identify them. Those `Class` can then

530 be instantiated into an object, and `Methods` objects can be used to call the methods of the
531 instantiated object. The process of manipulating `Class` and `Methods` objects instead of using
532 bytecode instructions is called reflection. Reflection is not limited to bytecode that has been
533 dynamically loaded: it can be used for any class or method available to the application.

534 Because the `ClassLoader` objects are only used when loading bytecode dynamically or when
535 using reflection, it is often forgotten that the ART uses class loaders constantly behind the
536 scene, allowing classes from the application and platform classes to cohabit seamlessly.

537 ∴

538 In this subsection, we presented the most notable specificities of the Android ecosystem. In the
539 next section, we will continue with the various tools available for an Android reverse engineer.

540 2.2.2 Reverse Engineering Tools

541 Due to the specificities of Android, reverse engineers need tools adapted to Android. The
542 development tools provided by Google can be used for basic operations, but a reverse engineer
543 will quickly need more specialised tools. Usually, the first step while analysing an application
544 is to look at its content. `Apktool` and `Jadx` are common tools used to convert the content of
545 an application into a readable format. Analysing an application this way, without running it,
546 is called static analysis. For more advanced forms of static analysis, `Androguard` and `Soot` can
547 be used as libraries to automate analyses. When static analysis becomes too complicated (*e.g.*,
548 if the application uses obfuscation techniques), a reverse engineer might switch to dynamic
549 analysis. This time, the application is executed, and the analyst will scrutinise the behaviour
550 of the application. `Frida` is a good option to help with this dynamic analysis. It is a toolkit
551 that can be used to intercept method calls and execute custom scripts while an application is
552 running.

553 **Android Studio** The whole Android development ecosystem is packaged by Google in the
554 IDE `Android Studio`¹. In practice, `Android Studio` is a source-code editor that wraps around
555 the different tools of the Android SDK. The SDK tools and packages can be installed manually
556 with the `sdkmanager` tool. Among the notable tools in the SDK are:

- 557 • `emulator`: this tool allows running an emulated Android phone on a computer. Although
558 very useful, Android emulator has several limitations. For once, it cannot emulate another
559 architecture. An `x86_64` computer cannot emulate an ARM smartphone. This can be an
560 issue because a majority of smartphones run on ARM processors. Also, for certain versions of
561 Android, the proprietary `GooglePlay` libraries are not available on rooted emulators. Lastly,
562 emulators are not designed to be stealthy and can easily be detected by an application.
563 Malware will avoid detection by not running its payload on emulators.

564 1. <https://developer.android.com/studio>

- 565 • ADB: a tool to send commands to an Android smartphone or emulator. It can be used to
566 install applications, send instructions, events, and generally perform debugging operations.
- 567 • Platform Packages: Those packages contain data associated with a version of Android needed
568 to compile an application. Especially, they contain the so-called `android.jar` files, which
569 contain the list of API for a version of Android.
- 570 • `d8`: The main use of `d8` is to convert Java bytecode files (`.class`) to Android DEX format.
571 It can also be used to perform different levels of optimisation of the bytecode generated.
- 572 • `aapt/aapt2` (Android Asset Packaging Tool): This tool is used to build the APK file. It is
573 commonly used by other tools that repackage applications like `Apktool`. Behind the scenes,
574 it converts XML to binary AXML and ensures that each file has the right compression and
575 alignment. (*e.g.*, some resource files are mapped in memory by the ART, and thus need to
576 be aligned and not compressed).
- 577 • `apksigner`: the tool used to sign an APK file. When repackaging an application, for example,
578 with `Apktool`, the new application needs to be signed.

579 **Apktool** `Apktool`¹ is a *reengineering tool* for Android APK files. It can be used to disassemble
580 an application: it will extract the files from the APK file, convert the binary AXML to text
581 XML, and use `smali/backsmali`² to convert the DEX files to `smali`, an assembler-like language
582 that matches the Dalvik bytecode instructions. The main strength of `Apktool` is that after
583 disassembling an application, its content can be edited and reassembled into a new APK.

584 **Androguard** `Androguard`³ [15] is a Python library for parsing and disassembling APK files.
585 It can be used to automatically read Android manifests, resources, and bytecode. Contrary to
586 `Apktool`, which generates text files, it can be used as a library to programmatically analyse the
587 application. It can also perform additional analysis, like computing a call graph or control flow
588 graph of the application (we will explain what those graphs are later in Section 2.2.3). However,
589 contrary to `Apktool`, it cannot repackage a modified application.

590 **Jadx** `Jadx`⁴ is an application decompiler. It converts DEX files to Java source code. It is not
591 always capable of decompiling all classes of an application, so it cannot be used to recompile
592 a new application, but the code generated can be very helpful to reverse an application. In
593 addition to decompiling DEX files, `Jadx` can also decode Android manifests and application
594 resources.

595 **Soot** `Soot`⁵ [7] was originally a Java optimisation framework. It could lift Java bytecode to
596 other intermediate representations that can be optimised, then converted back to bytecode.

597 1. <https://apktool.org/>

598 2. <https://github.com/JesusFreke/smali>

599 3. <https://github.com/androguard/androguard>

600 4. <https://github.com/skylot/jadx>

601 5. <https://github.com/soot-oss/soot>

602 Because Dalvik bytecode and Java bytecode are equivalent, support for Android was added to
603 Soot, and Soot features are now leveraged to analyse and modify Android applications. One
604 of the best-known examples of Soot usage for Android analysis is Flowdroid [8], a tool that
605 computes data flow in an application.

606 A new version of Soot, SootUp¹, is currently being worked on. Compared to Soot, it has a
607 modernised interface and architecture, but it is not yet feature-complete, and some tools like
608 Flowdroid are still using Soot.

609 **Frida** Frida² is a dynamic instrumentation toolkit. It allows the reverse engineer to inject
610 and run JavaScript code inside a running application.

611 To instrument an application, the Frida server must be running as root on the phone, or the
612 Frida library must be injected inside the APK file before installing it. Frida defines a JavaScript
613 wrapper around the Java Native Interface (JNI) used by native code to interact with Java classes
614 and the Android API. In addition to allowing interaction with Java objects from the application
615 and the Android API, this wrapper provides the option to replace a method implementation
616 with a JavaScript function (that itself can call the original method implementation if needed).
617 This makes Frida a powerful tool capable of collecting runtime information or modifying the
618 behaviour of an application as needed.

619 The main drawback of using Frida is that it is a known tool, easily detected by applications.
620 Malware might implement countermeasures that avoid running malicious payloads if Frida is
621 detected.

622 ∴

623 Those tools are quite useful for manual operations. However, considering the complexity of
624 modern Android applications, it might take a lot of work for a reverse engineer to analyse one
625 application. Different techniques have been developed to streamline the analysis. Next, we will
626 see the most common of those techniques for static analysis.

627 **2.2.3 Static Analysis**

628 A static analysis program examines an APK file without executing it to extract information
629 from it. Basic static analysis can include extracting information from the `AndroidManifest.xml`
630 file or decompiling bytecode to Java code with tools like Apktool or Jadx. Unfortunately,
631 simply reading the bytecode does not scale. To do so, a human analyst is needed, making it
632 complicated to analyse a large number of applications, and even for single applications, the size
633 and complexity of some applications can quickly overwhelm the reverse engineer.

634 1. <https://github.com/soot-oss/SootUp>

635 2. <https://frida.re/>

```

1 public static void fizzBuzz(int n) {
2   for (int i = 1; i <= n; i++) {
3     if (i % 3 == 0 && i % 5 == 0) {
4       Buzzer.fizzBuzz();
5     } else if (i % 3 == 0) {
6       Buzzer.fizz();
7     } else if (i % 5 == 0) {
8       Buzzer.buzz();
9     } else {
10      Log.e("fizzbuzz", String.valueOf(i));
11    }
12  }
13 }

```

a) A Java program

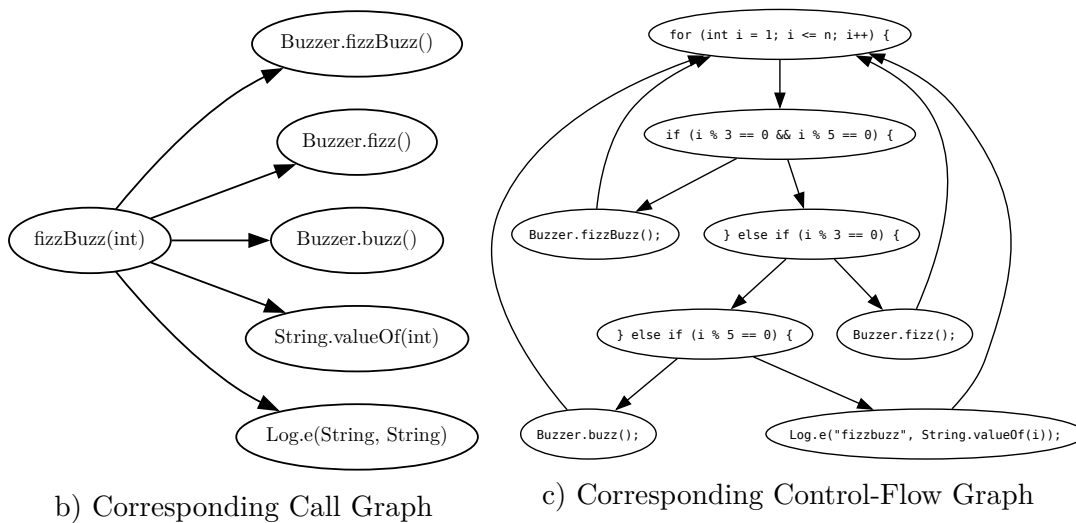


Figure 1: Source code for a simple Java method and its Call and Control Flow Graphs

636 Control flow analysis is often used to mitigate this issue. The idea is to extract the behaviour, the
637 flow, of the application from the bytecode, and to represent it as a graph. A graph representation
638 is easier to work with than a list of instructions and can be used for further analysis. Depending
639 on the level of precision required, different types of graphs can be computed. The most basic of
640 those graphs is the call graph. A call graph is a graph where the nodes represent the methods in

641 the application, and the edges represent calls from one method to another. Figure 1 b) show the
642 call graph of the code in Figure 1 a). A more advanced control-flow analysis consists of building
643 the control-flow graph. This time, instead of methods, the nodes represent instructions, and
644 the edges indicate which instruction can follow which instruction. Figure 1 c) represents the
645 control-flow graph of Figure 1 a), with code statements instead of bytecode instructions.

646 Once the control-flow graph is computed, it can be used to compute data-flows. Data-flow
647 analysis is used to follow the flow of information in the application. By defining a list of methods
648 and fields that can generate critical information (taint sources) and a list of methods that
649 can consume information (taint sinks), taint-tracking detects potential data leaks (if a data
650 flow links a taint source and a taint sink). For example, `TelephonyManager.getImei()` returns
651 a unique, persistent, device identifier. This can be used to identify the user, and it cannot be
652 changed if compromised. This makes `TelephonyManager.getImei()` a good candidate as a taint
653 source. On the other hand, `URLRequest.start()` sends a request to an external server, making it
654 a taint sink. If a data-flow is found linking `TelephonyManager.getImei()` to `URLRequest.start()`,
655 this means the application is potentially leaking critical information to an external entity, a
656 behaviour that is probably not wanted by the user.

657 Static analysis is powerful as it can detect unwanted behaviour in an application, even if the
658 behaviour does not manifest itself when running the application. However, static analysis tools
659 must overcome many challenges when analysing Android applications.

660 **the Java object-oriented paradigm** A call to a method can, in fact, correspond to a call
661 to any method overriding the original method in subclasses.

662 **the multiplicity of entry points** Each component of an application can be an entry point
663 for the application.

664 **the event-driven architecture** Methods in the applications can be called when events occur,
665 in an unknown order.

666 **the interleaving of native code and bytecode** Native code can be called from bytecode
667 and vice versa, but tools often only handle one of those formats.

668 **the potential dynamic code loading** An application can run code that was not originally
669 in the application.

670 **the use of reflection** Methods can be called from their name as a string object, which is
671 difficult to identify statically.

672 **the continual evolution of Android** each new version of Android brings new features
673 that analysis tools must be aware of. For instance, the multi-dex feature presented in
674 Section 2.2.1.1 was introduced in Android SDK 21. Tools unaware of this feature only
675 analyse the `classes.dex` file and will ignore all other `classes<n>.dex` files.

676 :TODO: Ca serait bien de souligner Dyn Code Load et Reflection :

677 ∴

678 With the bases of Android application analysis in mind, we can now examine our problem
679 statements further.

680 2.3 Problems of the Reverse Engineer

681 In this section, we will develop on some issues encountered by reverse engineers, and link them
682 to our problem statements.

683 In the previous section, we listed some limitations to static analysis. Some limitations have
684 been known for some time now, and many contributions have been made to overcome them.
685 Those contributions often introduce new tools that implement solutions to those different issues.
686 Depending on the situation, a reverse engineer might want to use those tools or build another
687 tool on top of one. Unfortunately, they can be hard to use. And like we said previously, the fast
688 evolution of Android can be a significant obstacle. The combination of those two points can
689 lead a reverse engineer to spend a lot of time trying to use a tool without realising that the tool
690 does not work anymore. Our first problem statement **Pb1** focuses on this issue: *To what extent
691 are previously published Android analysis tools still usable today, and what factors impact their
692 reusability?* Determining which tools are still usable today is a first step, but finding out what
693 reasons make a tool stop working might help write more resilient tools in the future.

694 We also presented dynamic code loading as an obstacle for static analysis. Code loading is
695 achieved using class loader objects, causing class loaders to be generally associated with dynamic
696 code loading. However, class loading plays a much more important role in the ART. Class
697 loading originates from the Java ecosystem and was ported to Android so that developers could
698 keep writing applications in Java. Despite that, Android made a lot of changes to the original
699 Java classes and did not document those changes. Between static analysis, general oversight of
700 class loading, relegating it to dynamic analysis, and the lack of documentation of the actual
701 behaviour of the ART, the question of the impact of the class loading algorithm on static
702 analysis can be asked. Our second problem statement **Pb2** aims to answer this question: *What
703 is the default Android class loading algorithm, and does it impact static analysis?*

704 Circling back to known limitations of static analysis, dynamic code loading and reflection are
705 often used to obfuscate applications. Dynamic code loading allows hiding bytecode from static
706 analysis with relatively low effort. The bytecode can be downloaded at runtime, stored in the
707 application encrypted, hidden inside other files, generated at runtime, etc. In a way, reflection
708 can do the same thing, but for specific method calls: instead of the actual call, static analysis
709 will see a call to the generic `Method.invoke()` method. By contrast, it is relatively easy to find
710 the name of the method called or to intercept dynamically loaded bytecode using dynamic tools
711 like Frida. The issue that arises then is what to do with the collected data. Simply having it
712 greatly helps a manual analysis, but it cannot be used directly by tools that perform static

713 analyses. There is no standard representation for runtime information, and there is simply no
714 way to give a list of reflection sites and the associated method calls as a new input for most static
715 analysis tools. This means that in most cases, when a reverse engineer wants to improve static
716 analysis with dynamic analysis, they need to modify the static tools to receive the additional
717 runtime data. Doing so requires both time and knowledge of the internals of the tools used. Our
718 third problem statement, **Pb3**, explores an alternative approach that modifies the application
719 instead of the tool: *Can we use instrumentation to provide dynamic code loading and reflection*
720 *data collected dynamically to static analysis tools and improve their results?*

721 We will now explore the current state of the art for relevant contributions related to our problem
722 statements.

723 **2.4 State of the Art**

724 This section focuses on the state of the art related to our three problem statements: the
725 reusability of Android static analysis tools, the class loading mechanism of Android, and the
726 use of instrumentation to encode information collected dynamically.

727 **2.4.1 Reusability of Static Analysis Tools**

728 *To what extent are previously published Android analysis tools still usable today, and what*
729 *factors impact their reusability?*

730 In the past fifteen years, the research community has released many tools to detect or analyse
731 malicious behaviours in applications. The first step to answer this question is to list those
732 previously published tools. The number of publications related to static analysis can make it
733 difficult to find the right tool for the right task. Li *et al.* [33] published a systematic literature
734 review for Android static analysis before May 2015. They analysed 92 publications and classified
735 them by goal, method used to solve the problem and underlying technical solution for handling
736 the bytecode when performing the static analysis. In particular, they listed 27 approaches with
737 an open-source implementation available.

738 Interestingly, a lot of the tools listed rely on common tools to interact with Android
739 applications/DEX bytecode. Recurring examples of such support tools are Apktool (*e.g.*,
740 Amandroid [66], Bluseal [58], SAAF [26]), Androguard (*e.g.*, Adagio [21], Appareciumn [63],
741 Mallodroid [19]) or Soot (*e.g.*, Bluseal [58], DroidSafe [24], Flowdroid [8]): those tools are built
742 incrementally, on top of each other. This strengthens our idea that being able to reuse previous
743 tools is important.

744 Nevertheless, Li *et al.* focus more on the techniques and features described in the reviewed
745 publications, and experiments to evaluate whether the pointed out software are still usable were
746 not performed. We believe that the effort of reviewing the literature for making a comprehensive

overview of available approaches should be pushed further: an existing published approach with a software that cannot be used for technical reasons endangers both the reproducibility and reusability of research.

We will now explore this direction further by looking at other works that have been done to evaluate different analysis tools. Those evaluations often take the form of benchmarks and follow a similar method (we will look at the different contributions in more details in Section 2.4.1.2). They start by selecting a set of tools with similar goals to compare. Usually, those contributions are comparing existing tools to their own, but some contributions do not introduce a new tool and focus on surveying the state of the art for some technique. They then selected a dataset of applications to analyse. We will see in Section 2.4.1.1 that those datasets are often hand-crafted, except for some studies that select a few real-world applications that they manually reverse-engineered to get a ground truth to compare to the tool's result. Once the tools and test dataset are selected, the tools are run on the application dataset, and the results of the tools are compared to the expected results (ground truth) to determine the accuracy of each tool. Additional factors are sometimes compared as well: the number of false positives, false negatives, or even the time it took to finish the analysis. Occasionally, the number of applications a tool simply failed to analyse is also compared.

In Section 2.4.1.1, we will look at the dataset used in the community to compare analysis tools. Then, in Section 2.4.1.2, we will go through the contributions that benchmarked those tools to see if they can be used as an indication as to which tools can still be used today.

2.4.1.1 Application Datasets

Research contributions often rely on existing datasets or provide new ones in order to evaluate the developed software. Raw datasets such as Drebin[6] contain little information about the provided applications. As a consequence, dataset suites have been developed to provide, in addition to the applications, meta information about the expected results. For example, taint analysis datasets should provide the source and expected sink of a taint. In some cases, the datasets are provided with additional software for automating part of the analysis. One such dataset is DroidBench, which was released with the tool Flowdroid [8]. Later, the dataset ICC-Bench was introduced with the tool Amandroid [66] to complement DroidBench by introducing applications using Inter-Component data flows. These datasets contain carefully crafted applications containing flows that the tools should be able to detect. These hand-crafted applications can also be used for testing purposes or to detect any regression when the software code evolves. The drawback to using hand-crafted applications is that these datasets are not representative of real-world applications [50] and the obtained results can be misleading.

A mettre en avant?

Mettre en avant

782 Contrary to DroidBench and ICC-Bench, some approaches use real-world applications. Bosu
783 *et al.* [12] use DIALDroid to perform a threat analysis of Inter-Application communication
784 and published DIALDroid-Bench, an associated dataset. Similarly, Luo *et al.* released Taint-
785 Bench [40], a real-world dataset, and the associated recommendations to build such a dataset.
786 These datasets are useful for carefully spotting missing taint flows, but contain only a few dozen
787 applications.

788 In addition to those datasets, AndroZoo [3] collect applications from several application
789 marketplaces, including the Google Play store (the official Google application store), Anzhi
790 and AppChina (two Chinese stores), or FDroid (a store dedicated to free and open source
791 applications). Currently, AndroZoo contains more than 25 million applications that can be
792 downloaded by researchers from the SHA256 hash of the application. AndroZoo also provides
793 additional information about the applications, like the date the application was detected for the
794 first time by AndroZoo or the number of antiviruses from VirusTotal that flagged the application
795 as malicious. This will allow us to sample a dataset of applications evenly distributed over the
796 years. In addition to providing researchers with easy access to real-world applications, AndroZoo
797 makes it a lot easier to share datasets for reproducibility: instead of sharing hundreds of APK
798 files, the list of SHA256 is enough.

799 **2.4.1.2 Benchmarking**

800 We will now go through the different contributions that evaluated different static analysis tools
801 to see if they can give us some insights into the current usability of the tools.

802 The few experiments with datasets composed of real-world applications confirmed that some
803 tools, such as Amandroid [66] and Flowdroid [8], are less efficient on real-world applications [12,
804 40]. Unfortunately, those real-world applications datasets are rather small, and a larger number
805 of applications would be more suitable for our goal, *i.e.*, evaluating the reusability of a variety
806 of static analysis tools.

807 Pauck *et al.* [49] used DroidBench [8], ICC-Bench [66] and DIALDroid-Bench [12] to compare
808 Amandroid [66], DIAL-Droid [12], DidFail [28], DroidSafe [24], FlowDroid [8] and IccTA [30]. To
809 perform their comparison, they introduced the AQL (Android App Analysis Query Language)
810 format. AQL can be used as a common language to describe the computed taint flow as well as
811 the expected result for the datasets. It is interesting to notice that all the tested tools timed out
812 at least once on real-world applications, and that Amandroid [66], DidFail [28], DroidSafe [24],
813 IccTA [30] and ApkCombiner [31] (a tool used to combine applications) all failed to run on
814 applications built for Android API 26. These results suggest that a more thorough study of the
815 link between application characteristics (*e.g.*, date, size) should be conducted. Luo *et al.* [40]
816 used the framework introduced by Pauck *et al.* to compare Amandroid [66] and Flowdroid [8] on
817 DroidBench and their own dataset TaintBench, composed of real-world Android malware. They

818 found out that those tools have a low recall on real-world malware, and are thus over-adapted
819 to micro-datasets. Unfortunately, because AQL is only focused on taint flows, we cannot use it
820 to evaluate tools performing more generic analysis.

821 A first work about quantifying the reusability of static analysis tools was proposed by Reaves *et*
822 *al.* [53]. Seven Android analysis tools (Amandroid [66], AppAudit [69], DroidSafe [24], Epicc [48],
823 FlowDroid [8], MalloDroid [19] and TaintDroid [17]) were selected to check if they were still
824 readily usable. For each tool, both the usability and results of the tool were evaluated by
825 asking auditors to install and use it on DroidBench and 16 real-world applications. The auditors
826 reported that most of the tools require a significant amount of time to set up, often due to
827 dependency issues and operating system incompatibilities. Reaves *et al.* propose to solve these
828 issues by distributing a Virtual Machine with a functional build of the tool in addition to the
829 source code. Regrettably, these Virtual Machines were not made available, preventing future
830 researchers from taking advantage of the work done by the auditors. Reaves *et al.* also report
831 that real-world applications are more challenging to analyse, with tools having lower results,
832 taking more time and memory to run, sometimes to the point of not being able to run the
833 analysis. Considering it was noticed on a dataset of only 16 real-world applications, this result
834 is worrying. A more diverse dataset would be needed to better assess the extent of the issue
835 and give more insight into the factors impacting the performance of the tools.

836 Mauthe *et al.* present an interesting methodology to assess the robustness of Android decompilers
837 [42]. They used 4 decompilers on a dataset of 40 000 applications. The error messages of
838 the decompilers were parsed to list the methods that failed to decompile, and this information
839 was used to estimate the main causes of failure. It was found that the failure rate is correlated to
840 the size of the method, and that a consequent amount of failures are from third-party libraries
841 rather than the core code of the application. They also concluded that malware are easier to
842 entirely decompile, but have a higher failure rate, meaning that the ones that are hard to
843 decompile are substantially harder to decompile than goodware.

844 \therefore

845 To summarise, Li *et al.* made a systematic literature review of static analysis for Android that
846 listed 27 open-sourced tools. However, they did not test those tools. Reaves *et al.* did so for some
847 of them and analysed the difficulty of using them. They raised two major concerns about the
848 use of Android static analysis tools. First, they can be quite difficult to set up, and second, they
849 appear to have difficulties analysing real-world applications. This is problematic for a reverse
850 engineer, not only do they need to invest a significant amount of work to set up a tool properly,
851 but they also do not have any guarantees that the tool will actually manage to analyse the
852 application they are investigating.

853 In Chapter 3, we will try to set up the tools listed by Li *et al.* and test them on a large number
854 of real-world applications to see which can be used today. We will also aim at identifying what
855 characteristics of real-world applications make them harder to analyse.

856 2.4.2 Android Class Loading

857 *What is the default Android class loading algorithm, and does it impact static analysis?*

858 This subsection is mainly dedicated to class loading in Java and Android. Because we focus on
859 the *default* class loading algorithm, we will not focus on dynamic code loading (*i.e.*, loading of
860 additional bytecode while the application is already running). However, class loading is used,
861 without dynamic code loading, to load classes other than the one in the application: the platform
862 classes. In the second part of this subsection, we will look at the work that has been done
863 related to those platform classes.

864 2.4.2.1 Class Loading

865 Class loading is a fundamental element of Java; it defines which classes are loaded from where.
866 In Android, this is often associated with dynamic code loading, as the `ClassLoader` objects are
867 used to load code at runtime. However, class loading also intervenes to load platform classes
868 or classes from the application itself, and thus requires some attention when analysing an
869 application.

870 Class loading mechanisms have been studied in the general context of the Java language.
871 Gong [23] describes the JDK 1.2 class loading architecture and capabilities. One of the main
872 advantages of class loading is the type safety property that prevents type spoofing. As explained
873 by Liang and Bracha [35], by capturing events at runtime (new loaders, new class) and
874 maintaining constraints on the multiple loaders and their delegation hierarchy, authors can
875 avoid confusion when loading a spoofed class. This behaviour is now implemented in modern
876 Java virtual machines. Later, Tazawa and Hagiya [64] proposed a formalisation of the Java
877 Virtual Machine supporting dynamic class loading in order to ensure type safety. Those works
878 ensure strong safety for the Java Virtual Machine, in particular when linking new classes at
879 runtime. Although Android has a similar mechanism, the implementation is not shared with
880 the JVM of Oracle. Additionally, our problem statement does not focus on spoofing classes at
881 runtime, but on confusions that occur when using a static analyser used by a reverser that tries
882 to understand the code loading process offline.

883 Contributions about Android class loading focus on using the capabilities of class loading to
884 extend Android features or to prevent reverse engineering of Android applications. For instance,
885 Zhou *et al.* [76] extend the class loading mechanism of Android to support regular Java
886 bytecode, and Kritz and Maly [29] propose a new class loader to automatically load modules
887 of an application without user interactions.

888 Regarding reverse engineering, class loading mechanisms are frequently used by packers,
889 applications that load their actual code at runtime, for hiding all or parts of the code of
890 an application [16]. For example, packers exploits the class loading capability of Android to
891 load new code. They also combine the loading with code generation from ciphered assets or
892 code modification from native code calls [37] to increase the difficulty of recovery of the code.
893 Because parts of the original code will be only available at runtime, deobfuscation approaches
894 propose techniques that track DEX structures when manipulated by the application [68, 70,
895 74]. Those contributions interact with the class loading mechanism of Android to collect the
896 DEX structures at the right moment.

897 Some classes, however, are not loaded from the application, nor dynamically loaded by the
898 application. Those classes are platform classes, and apart from dynamic code loaded, they are
899 the main reason class loading is needed by Android. We will now look at the literature related
900 to them.

901 2.4.2.2 Platform Classes

902 Platform classes are divided between SDK classes that are documented, and the other classes,
903 often referred to as hidden APIs. SDK classes are clearly listed and documented by Google,
904 so they do not require as much attention as hidden APIs. As we said earlier, hidden API are
905 undocumented methods that can be used by an application, thus making them a potential
906 blind spot when analysing an application. However, not a lot of research has been done on the
907 subject.

908 Li *et al.* did an empirical study of the usage and evolution of hidden API [34]. They found that
909 hidden API are added and removed in every release of Android, and that they are used both
910 by benign and malicious applications.

911 More recently, He *et al.* [25] did a systematic study of hidden service API related to security.
912 They studied how the hidden API can be used to bypass Android security restrictions and
913 found that although Google countermeasures are effective, they need to be implemented inside
914 the system services and not the hidden API due to the lack of in-app privilege isolation: the
915 framework code is in the same process as the user code, meaning any restriction in the framework
916 can be bypassed by the user. Unfortunately, those two contributions do not explore further the
917 consequences of the use of hidden APIs for a reverse engineer.

918 \therefore

919 In conclusion, class loading mechanisms have been studied carefully in the context of the Java
920 language. However, the same cannot be said about Android, whose implementation diverges
921 significantly from classic Java Virtual Machines. Most work done on Android focuses on

922 extending Android capabilities using class loading, or on analysing dynamically the code loading
923 operations of an application.

924 In Chapter 4, we will model the behaviour of Android when loaded classes used by an application
925 that do not use dynamic code loading, and check if this behaviour matches the behaviour of
926 common analysis tools. We will also take some time to check if the state of the art related to
927 hidden API is up to date with the current Android versions.

928 **2.4.3 Allowing Static Analysis Tools to Analyse Obfuscated Application**

929 *Can we use instrumentation to provide dynamic code loading and reflection data collected*
930 *dynamically to static analysis tools and improve their results?*

931 Dynamic analysis of Android applications has been researched for a long time. Like static
932 analysis, it has its own challenges, which we will explore in this subsection. After that, we
933 will also look at contributions that sought to encode results inside the APK format or used
934 instrumentation to improve analyses in some way.

935 **2.4.3.1 Dynamic Analysis**

936 Some situations, like reflection of dynamic code loading, are difficult to solve with static analysis
937 and require a different approach: dynamic analysis. With dynamic analysis, the application is
938 actually executed, and the reverse engineer observes its behaviour. Monitoring the behaviour
939 can be achieved by various strategies: observing the filesystem, the display screen, the process
940 memory, the kernel, etc. Depending on the chosen level of observation, dynamic analysis can
941 become a serious technical challenge.

942 A basic example of dynamic analysis is presented by Bernardi *et al.* [10]: the logs generated by
943 `strace` are used to list the system calls generated in response to an event to determine if an
944 application is malicious or not.

945 More advanced methods are more intrusive and require modifying either the APK, the Android
946 framework, runtime, or kernel. TaintDroid [17], for example, modifies the Dalvik Virtual
947 Machine (the predecessor of the ART) to track the data flow of an application at runtime,
948 while AndroBlare [4, 5] try to compute the taint flow by hooking system calls using a Linux
949 Security Module. DexHunter [74] and AppSpear [73] also patch the Dalvik Virtual Machine/
950 ART, this time to collect bytecode loaded dynamically. Modifying the Android framework,
951 runtime, or kernel is possible thanks to the Android project being open-source, but this is
952 a delicate operation that requires revising a patch for each new version of Android. Thus, a
953 common issue faced by tools that took this approach is that they are stuck with a specific
954 version of Android. Some sandboxes limit this issue by using dynamic binary instrumentation,
955 like DroidHook [14], based on the Xposed framework, or CamoDroid [18], based on Frida. This

956 approach is a lot less stealthy than patching Android, but it is generally easier to set up and is
957 easier to port to new Android versions.

958 Another known challenge when analysing an application dynamically is the code coverage: if
959 some part of the application is not executed, it cannot be analysed. Considering that Android
960 applications are meant to interact with a user, this can become problematic for automatic
961 analysis. The Monkey tool developed by Google is one of the most used solution [61]. It sends
962 a random stream of events to the phone without tracking the state of the application. More
963 advanced tools statically analyse the application to model in order to improve the exploration.
964 Sapienz [41] and Stoa [60] use this technique to improve application testing. GroddDroid [1]
965 has the same approach but detects statically suspicious sections of code to target, and will
966 interact with the application to target those code sections.

967 Unfortunately, exploring the application entirely is not always possible, as some applications
968 will try to detect if they are in a sandbox environment (*e.g.*, if they are in an emulator, or if
969 Frida is present in memory) and will refuse to run some sections of code if this is the case.
970 Ruggia *et al.* [55] make a list of evasion techniques. They propose a new sandbox, DroidDungeon,
971 that, contrary to other sandboxes like DroidScope[71] or CopperDroid[62], strongly emphasises
972 resilience against evasion mechanisms.

973 A common objective of dynamic analysis is to collect bytecode loaded dynamically and reflection
974 information. Like we said earlier, DexHunter [74] and AppSpear [73] do that by instrumenting
975 the Android Runtime. Qu *et al.* [52] developed DyDroid, a hybrid framework using dynamic
976 analysis to intercept dynamic code loading and static analysis to determine the nature of the
977 loaded code. They used DyDroid to make an audit of the use of dynamic code loading in
978 applications from the Google Play store in 2016. It resulted that dynamic code loading was
979 mostly related to mobile advertisement, and that the code loading originated from a third-party
980 library included in the application, rather than the code of the application developer itself.
981 Similarly, StaDynA [75] is a framework that generates a call graph statically, then uses dynamic
982 analysis to analyse dynamic code loading and reflection calls to complete this call graph.

983 The issue with those approaches is that they are only compatible with their own subsequent
984 analysis. For instance, StaDynA only provides the call graph, and cannot be used as is to
985 improve the capacity of Flowdroid. This is unfortunate: the reverse engineer's next step will
986 depend on the context. Not being able to reuse the result of a previous analysis with any ad hoc
987 tools greatly limits their options. AppSpear has an interesting solution to this issue: the code
988 it intercepts is repackaged inside a new APK file that Android analysis tools should be able
989 to analyse. We will now explore further the contributions that take this approach to encode
990 results inside applications.

991 2.4.3.2 Improving Analysis with Instrumentation

992 Usually, instrumentation refers to the practice of modifying the behaviour of a program to collect
993 information during its execution. Frida is a good example of an instrumentation framework.
994 The term can also be used more generally to describe operations that modify the application
995 code. In this section, we will focus on the use of instrumentation that makes an application
996 easier to analyse by other tools, instead of just collecting additional information at runtime.

997 In the previous section, we gave the example of AppSpear [73], which reconstructs DEX files
998 intercepted at runtime and repackages the APK with the new code in it. DexLego [46] has a
999 similar but a lot more aggressive technique. It targets heavily obfuscated packers that decrypt
1000 then re-encrypt the method's instructions just in time. To get the bytecode, DexLego logs each
1001 instruction executed by the ART, and reconstructs the methods, then the DEX files, from this
1002 stream of instructions. The main limitation of this technique is that it carries over the limitation
1003 of dynamic analysis to static analysis: the bytecode injected in the application is limited to
1004 the instructions executed during the dynamic analysis. Nevertheless, it is an interesting way to
1005 encode the traces of a dynamic analysis in a way that can be used by any Android analysis tool.

1006 IccTa [30] technique is close to the idea of modifying the application to improve its analysis:
1007 it performs a first analysis to compute the potential inter-component communication of an
1008 application, then modifies the Jimple representation of this application before feeding it to
1009 Flowdroid to perform a taint analysis. Jimple is the intermediate language used by Soot, so
1010 even if IccTa does not generate a new application, this modified representation can probably be
1011 used by any tool based on the Soot framework or recompiled into a new application without too
1012 much effort. Samhi *et al.* [56] followed this direction to unify the analysis of bytecode and native
1013 code. Their tool, JuCify, uses Angr [59] to generate the call graph of the native code, and uses
1014 heuristics to encode this call graph into Jimple that can then be added to the Jimple generated
1015 by Soot from the bytecode of the application. Like IccTa, they use Flowdroid to analyse this
1016 new augmented representation of the application, but it should be usable by any analysis tools
1017 relying on Soot.

1018 Finally, DroidRA [32] uses the COAL [47] solver to statically compute the reflection information.
1019 The reflection calls are transformed into direct calls inside the application using Soot. Using
1020 COAL makes DroidRA quite good at solving the simpler cases, where the names of classes and
1021 methods targeted by reflection are already present in the application. Those cases are quite
1022 common; being able to solve those without resorting to dynamic analysis is quite useful. On
1023 the other hand, COAL will struggle to solve cases with complex string manipulation and is
1024 simply not able to handle cases that rely on external data (*e.g.*, downloaded from the internet
1025 at runtime). Likewise, this can only access code loaded dynamically if the code was present
1026 inside the application without any kind of obfuscation (*e.g.*, a DEX file in the assets of the
1027 application can be analysed, but not if it is ciphered).

1028

∴

1029 Instrumenting applications to encode the result of an analysis as a unified representation
1030 has been explored before. It has been used by tools like AppSpear and DexLego to expose
1031 heavily obfuscated bytecode collected dynamically. Similarly, DroidRA compute reflection
1032 information statically and injects the actual method calls inside the application it returns.
1033 However, AppSpear and DexLego focus primarily on specific obfuscation techniques, making
1034 their implementation difficult to port to more recent versions of Android, and DroidRA suffers
1035 from the limitation of static analysis. We believe that instrumentation is a promising approach
1036 to encoding that information. Especially, we think that it could be used to provide dynamic
1037 information that is not available to static analysis tools like DroidRA.

1038 In Chapter 5, we will try to use instrumentation to combine dynamic analysis (to collect code
1039 loaded dynamically and reflection information) with static analysis, regardless of the static
1040 analysis tool used.

1041 2.5 Conclusion

1042 This chapter presented the specificities of Android and the usual tools used as a basis for
1043 reverse engineering applications. Many contributions have been made to static analysis, and
1044 benchmarks have been proposed to compare the different tools that resulted from those
1045 contributions. Those benchmarks raised questions about the reusability of those tools and
1046 their capacity to handle real-world applications. We then looked at platform classes and class
1047 loading, a commonly recognised limitation of static analysis. Because of that, the issue is
1048 generally relegated to dynamic analysis, leaving the details of the class loading mechanisms
1049 of Android unexplored. To complement static analysis, we continued by looking at dynamic
1050 analysis. A variety of approaches have been proposed, balancing ease of use, maintainability
1051 and stealthiness. The results of those analyses are often in an ad hoc format, making it difficult
1052 to reuse with other tools. A few exceptions, as well as some static analysis tools, proposed an
1053 interesting solution to this issue: instrumenting the analysed application to encode the results
1054 of the analysis in the form of a valid APK, a format any Android analysis tools should be able
1055 to read. We liked this solution and believe it should be studied further. This process led us to
1056 explore three problem statements:

1057 **Pb1** *To what extent are previously published Android analysis tools still usable today, and what*
1058 *factors impact their reusability?*

1059 **Pb2** *What is the default Android class loading algorithm, and does it impact static analysis?*

1060 **Pb3** *Can we use instrumentation to provide dynamic code loading and reflection data collected*
1061 *dynamically to static analysis tools and improve their results?*

1062 In the next chapters, we will endeavour to contribute to the Android reverse engineering field
1063 by answering them.

1064

1065

1066

EVALUATING THE REUSABILITY OF ANDROID STATIC ANALYSIS TOOLS

1067

I keep going back in time.

1068

— Max Caulfield, Life is Strange "Out of Time"

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

This chapter intends to explore the robustness of past software dedicated to static analysis of Android applications. We pursue the community effort that identified software supporting publications that perform static analysis of mobile applications, and we propose a method for evaluating the reliability of this software. We extensively evaluate static analysis tools on a recent dataset of Android applications, including goodware and malware, that we designed to measure the influence of parameters such as the date and size of applications. Our results show that 54.55% of the evaluated tools are no longer usable and that the size of the bytecode and the min SDK version have the greatest influence on the reliability of the tested tools.

1081

3.1 Introduction

1082

1083

1084

1085

1086

In this chapter, we study the reusability of open source static analysis tools that appeared between 2011 and 2017, on a recent Android dataset. The scope of our study is **not** to quantify if the output results are accurate to ensure reproducibility, because all the studied static analysis tools have different goals in the end. On the contrary, we take the hypothesis that the provided tools compute the intended result, but may crash or fail to compute a result due to the evolution

1087 of the internals of an Android application, raising unexpected bugs during an analysis. This
 1088 chapter intends to show that sharing the software artefacts of a paper may not be sufficient to
 1089 ensure that the provided software will be reusable.

1090 Thus, our contributions are the following. We carefully retrieved static analysis tools for Android
 1091 applications that were selected by Li *et al.* [33] between 2011 and 2017. We contacted the
 1092 authors whenever possible to select the best candidate versions and to confirm the good usage
 1093 of the tools. We rebuild the tools in their original environment and share our Docker images.¹
 1094 We evaluated the reusability of the tools by measuring the number of successful analyses of
 1095 applications taken in the Drebin dataset [6] and in a custom dataset that contains more recent
 1096 applications (62 525 in total). The observation of the success or failure of these analyses enables
 1097 us to answer the following research questions:

1098 **RQ1** Which Android static analysis tools that are more than 5 years old are still available and
 1099 can be reused without crashing with a reasonable effort?

1100 **RQ2** How has the reusability of tools evolved over time, especially when analysing applications
 1101 that are more than 5 years away from the publication of the tool?

1102 **RQ3** Does the reusability of tools change when analysing goodware compared to malware?

1103 The chapter is structured as follows. Section 3.2 presents the methodology employed to build
 1104 our evaluation process, and Section 3.3 gives the associated experimental results. Section 3.4
 1105 investigates the reasons behind the observed failures of some of the tools. We then compare
 1106 in Section 3.5 our results with the contributions presented in Chapter 2. In Section 3.6, we
 1107 give recommendations for tool development that we drew from our experience running our
 1108 experiment. Finally, Section 3.7 lists the limit of our approach, Section 3.8 presents further
 1109 avenues that did not have time to pursue and Section 3.9 concludes the chapter.

1110 3.2 Methodology

1111 In this section, we describe our methodology to evaluate the reusability of Android static
 1112 analysis tools. Figure 2 and Figure 3 summarize our approach. We collected tools listed as open
 1113 source by Li *et al.*, checked if the tools were only using static analysis techniques, and selected
 1114 the most recent version of the tool. We then packaged the tools inside containers and checked
 1115 our choices with the authors. We then run those tools on a large dataset that we sampled, and
 1116 collected the exit status of the run (whether the tool completed the analysis or not).

1117 3.2.1 Collecting Tools

1118 We collected the static analysis tools from [33], plus one additional paper encountered during
 1119 our review of the state-of-the-art (DidFail [28]). They are listed in Table 1, with the original
 1120 release date and associated publication. We intentionally limited the collected tools to the ones

1121 1. on Docker Hub as `histausse/rasta-<toolname>:icsr2024`

Tool	Availability			Repo type	Decision	Comments
	Bin	Src	Doc			
A3E [9] (2013)	–	✓	✓	github	✗	Hybrid tool (static/dynamic)
A5 [65] (2014)	–	✓	✗	github	✗	Hybrid tool (static/dynamic)
Adagio [21] (2013)	–	✓	✓	github	✓	
Amandroid [66] (2014)	✓	✓	✓	github	✓	
Anadroid [36] (2013)	✗	✓	✓	github	✓	
Androguard [15] (2011)	–	✓	✓	github	✓	
Android-app-analysis [22] (2015)	✗	✓	✓	google	✗	Hybrid tool (static/dynamic)
Apparecium [63] (2015)	✓	✓	✗	github	✓	
BlueSeal [58] (2014)	✗	✓	○	github	✓	
Choi <i>et al.</i> [13] (2014)	✗	✓	○	github	✗	Works on source files only
DIALDroid [12] (2017)	✓	✓	✓	github	✓	
DidFail [28] (2014)	✓	✓	○	bitbucket	✓	
DroidSafe [24] (2015)	✗	✓	✓	github	✓	
Flowdroid [8] (2014)	✓	✓	✓	github	✓	
Gator [54, 72] (2014, 2015)	✗	✓	✓	edu	✓	
IC3 [47] (2015)	✓	✓	○	github	✓	
IccTA [30] (2015)	✓	✓	✓	github	✓	
Lotrack [38] (2014)	✗	✓	✗	github	○	Authors ack. a partial doc.
MalloDroid [19] (2012)	–	✓	✓	github	✓	
PerfChecker [39] (2014)	✗	✗	○	request	✓	Binary obtained from authors
Poeplau <i>et al.</i> [51] (2014)	✗	○	✗	github	✗	Related to Android hardening
Redexer [27] (2012)	✗	✓	✓	github	✓	
SAAF [26] (2013)	✓	✓	✓	github	✓	
StaDynA [75] (2015)	✗	✓	✓	request	✗	Hybrid tool (static/dynamic)
Thresher [11] (2013)	✗	✓	✓	github	○	Not built with author’s help
Wognsen <i>et al.</i> [67] (2014)	–	✓	✗	bitbucket	✓	

binaries, sources: –: not relevant, ✓: available, ○: partially available, ✗: not provided
documentation: ✓✓: excellent, MWE, ✓: few inconsistencies, ○: bad quality, ✗: not available
decision: ✓: considered; ○: considered but not built; ✗: out of scope of the study

Table 1: Considered tools [33]: availability and usage reliability

1122 selected by Li *et al.* [33] for several reasons. First, not using recent tools enables a gap of at
1123 least 5 years between the publication and the more recent APK files, which enables measuring
1124 the reusability of previous contributions with a reasonable gap of time. Second, collecting new
1125 tools would require inspecting these tools in depth, similarly to what has been performed by
1126 Li *et al.* [33], which is not the primary goal of this chapter. Additionally, selection criteria such
1127 as the publication venue or number of citations would be necessary to select a subset of tools,
1128 which would require an additional methodology.

1129 Some tools use hybrid analysis (both static and dynamic): A3E [9], A5 [65], Android-app-
1130 analysis [22], StaDynA [75]. They have been excluded from this study. We manually searched the

- 1131 tool repository when the website mentioned in the paper is no longer available (*e.g.*, when the
 1132 repository has been migrated from Google code to GitHub), and for each tool we searched for:
- 1133 • an optional binary version of the tool that would be usable as a fallback (if the sources
 1134 cannot be compiled for any reason).
 - 1135 • the source code of the tool.
 - 1136 • the documentation for building and using the tool with an MWE.

1137 In Table 1 we rated the quality of these artifacts with “✓” when available but may have
 1138 inconsistencies, a “○” when too much inconsistencies (inaccurate remarks about the sources,
 1139 dead links or missing parts) have been found, a “✗” when no documentation have been found,
 1140 and a double “✓” for the documentation when it covers all our expectations (building process,
 1141 usage, MWE). Results show that documentation is often missing or very poor (*e.g.*, Lotrack),
 1142 which makes the rebuild process very complex and the first analysis of an MWE.

1143 We finally excluded Choi *et al.* [13] as their tool works on the sources of Android applications,
 1144 and Poeplau *et al.* [51] that focus on Android hardening. As a summary, in the end, we have
 1145 20 tools to compare. Some specificities should be noted. The IC3 tool will be duplicated in
 1146 our experiments because two versions are available: the original version of the authors and a
 1147 fork used by other tools like IccTa. For Androguard, the default task consists of unpacking the
 1148 bytecode, the resources, and the Manifest. Cross-references are also built between methods and
 1149 classes. Because such a task is relatively simple to perform, we decided to duplicate this tool
 1150 and ask Androguard to decompile an APK and create a control flow graph of the code using
 1151 its decompiler: DAD. We refer to this variant of usage as `androguard_dad`. For Thresher and
 1152 Lotrack, because these tools cannot be built, we excluded them from experiments.

1153 Finally, starting with 26 tools of Table 1, with the two variations of IC3 and Androguard, we
 1154 have in total 22 static analysis tools to evaluate, of which two tools cannot be built and will be
 1155 considered as always failing.

1156 3.2.2 Source Code Selection and Building Process

1157 In a second step, we explored the best sources to be selected among the possible forks of a
 1158 tool. We reported some indicators about the explored forks and our decision about the selected
 1159 one in Table 2. For each source code repository called “Origin”, we reported in Table 2 the
 1160 number of GitHub stars attributed by users, and we mentioned if the project is still alive
 1161 (✓ in column Alive when a commit exists in the last two years). Then, we analysed the fork
 1162 tree of the project. We searched recursively for any forked repository that contains a more
 1163 recent commit than the last one of the branch mentioned in the documentation of the original
 1164 repository. If such a commit is found (the number of such commits is reported in column Alive
 1165 Forks Nb), we manually looked at the reasons behind this commit and considered whether we
 1166 should prefer this more up-to-date repository instead of the original one (column “Alive Forks

Tool	Origin		Alive Forks		Last Commit	Authors	Environment
	Stars	Alive	Nb	Usable	Date	Reached	Language – OS
Adagio [21]	74	✓	0	✗	2022-11-17	✓	Python – U20.04
Amandroid [66]	161	✗	2	✗	2021-11-10	✓	Scala – U22.04
Anadroid [36]	10	✗	0	✗	2014-06-18	✗	Scala/Java/Python – U22.04
Androguard [15]	4430	✓	3	✗	2023-02-01	✗	Python – Python 3.11 slim
Apparecium [63]	0	✗	1	✗	2014-11-07	✗	Python – U22.04
BlueSeal [58]	0	✗	0	✗	2018-07-04	✓	Java – U14.04
DIALDroid [12]	16	✗	1	✗	2018-04-17	✗	Java – U18.04
DidFail [28]	4	✗			2015-06-17	✓	Java/Python – U12.04
DroidSafe [24]	92	✗	3	✗	2017-04-17	✓	Java/Python – U14.04
Flowdroid [8]	868	✓	1	✗	2023-05-07	✓	Java – U22.04
Gator [54, 72]					2019-09-09	✓	Java/Python – U22.04
IC3 [47]	32	✗	3	✓	2022-12-06	✗	Java – U12.04 / 22.04
IccTA [30]	83	✗	0	✗	2016-02-21	✓	Java – U22.04
Lotrack [38]	5	✗	2	✗	2017-05-11	✓	Java – ?
MalloDroid [19]	64	✗	10	✗	2013-12-30	✗	Python – U16.04
PerfChecker [39]		✗			–	✓	Java – U14.04
Redexer [27]	153	✗	0	✗	2021-05-20	✓	Ocaml/Ruby – U22.04
SAAF [26]	35	✗	5	✗	2015-09-01	✓	Java – U14.04
Thresher [11]	31	✗	1	✗	2014-10-25	✓	Java – U14.04
Wognsen <i>et al.</i> [67]				✗	2022-06-27	✗	Python/Prolog – U22.04

✓: yes, ✗: no, UX.04: Ubuntu X.04

Table 2: Selected tools, forks, selected commits and running environment

1167 Usable”). As reported in Table 2, we excluded all forks, except IC3, for which we selected the
 1168 fork JordanSamhi/ic3, because they always contain experimental code with no guarantee of
 1169 stability. For example, a fork of Aparecium contains a port for Windows 7, which does not
 1170 suggest an improvement in the stability of the tool. For IC3, the fork seems promising: it has
 1171 been updated to be usable on a recent operating system (Ubuntu 22.04 instead of Ubuntu 12.04
 1172 for the original version) and is used as a dependency by IccTa. We decided to keep these two
 1173 versions of the tool (IC3 and IC3_fork) to compare their results.

1174 Then, we self-allocated a maximum of four days for each tool to successfully read and follow
 1175 the documentation, compile the tool and obtain the expected result when executing an analysis
 1176 of an MWE. We sent an email to the authors of each tool to confirm that we used the most
 1177 suitable version of the code, that the command line we used to analyse an application is the
 1178 most suitable one and, in some cases, requested some help to solve issues in the building process.
 1179 We reported in Table 2 the authors who answered our request and confirmed our decisions.

1180 From this building phase, several observations can be made. Using a recent operating system,
 1181 it is almost impossible in a reasonable amount of time to rebuild a tool released years ago. Too
 1182 many dependencies, even for Java-based programs, trigger compilation or execution problems.




Figure 2: Tool selection methodology overview

1183 Thus, if the documentation mentions a specific operating system, we use a Docker image of
 1184 this OS.

1185 Most of the time, tools require additional external components to be fully functional. It could
 1186 be resources such as the `android.jar` file for each version of the SDK, a database, additional
 1187 libraries or tools. Depending on the quality of the documentation, setting up those components
 1188 can take hours to days. This is why we automated in a Dockerfile the setup of the environment
 1189 in which the tool is built and run¹.

1190 Figure 2 summarises our tool selection process. We first looked for the tools listed as open
 1191 source by Li *et al.*. For the tools still available, we checked if they used dynamic analysis and
 1192 removed them. We then checked if there were more recent updates of the tools and selected the
 1193 most relevant version. Finally, we marked as non-reusable the tools that we could not set up
 1194 within a period of 4 days, even with the help of the authors.

1195 3.2.3 Runtime Conditions

1196 As shown in Figure 3, before benchmarking the tools, we built and installed them in Docker
 1197 containers to facilitate any reuse by other researchers. We converted them into Singularity
 1198 containers because we had access to such a cluster and because this technology is often used
 1199 by the HPC community for ensuring the reproducibility of experiments. We performed manual
 1200 tests using these Singularity images to check:

- 1201 • the location where the tool writes on the disk. For the best performances, we expect the
- 1202 tools to write on a mount point backed by an SSD. Some tools may write data at unexpected
- 1203 locations, which require small patches from us.

1204 1. To guarantee reproducibility, we published the results, datasets, Dockerfiles and containers:
 1205 • <https://github.com/histausse/rasta> .
 1206 • <https://zenodo.org/records/10144014> .
 1207 • <https://zenodo.org/records/10980349> .
 1208 • on Docker Hub as `histausse/rasta-<toolname>:icsr2024`.

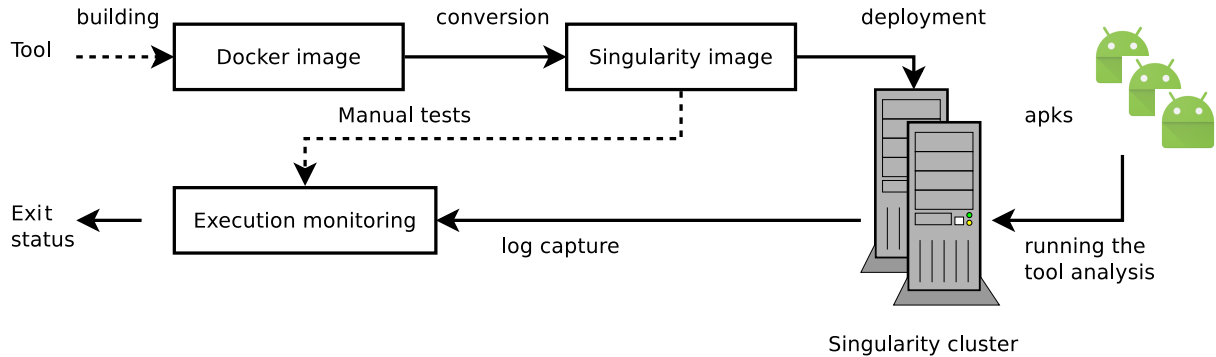


Figure 3: Experiment methodology overview

- 1209 • the amount of memory allocated to the tool. We checked that the tool could run an MWE
 1210 with a 64 GB limit of RAM.
- 1211 • the network connection opened by the tool, if any. We expect the tool not to perform any
 1212 network operation, such as the download of Android SDKs. Thus, we prepared the required
 1213 files and cached them in the images during the building phase. In a few cases, we patched
 1214 the tool to disable the download of resources.

1215 A campaign of tests consists of executing the 20 selected tools on all APKs of a dataset. The
 1216 constraints applied to the clusters are:

- 1217 • No network connection is authorised to limit any execution of malicious software.
 1218 • The allocated RAM for a task is 64 GB.
 1219 • The allocated maximum time is 1 hour.
 1220 • The allocated object space/stack space is 64 GB / 16 GB if the tool is a Java-based program.

1221 For the disk files, we use a mount point that is stored on an SSD disk, with no particular
 1222 size limit. Note that, because the allocation of 64 GB could be insufficient for some tools, we
 1223 evaluated the results of the tools on 20% of our dataset (described later in Section 3.2.4) with
 1224 128 GB of RAM and 64 GB of RAM and checked that the results were similar. With this
 1225 confirmation, we continued our evaluations with 64 GB of RAM only.

1226 3.2.4 Dataset

1227 Two datasets are used in the experiments of this section. The first one is **Drebin** [6], from
 1228 which we extracted the malware part (5479 samples that we could retrieve) for comparison
 1229 purposes only. It is a well-known and very old dataset that should not be used anymore because
 1230 it contains temporal and spatial biases [50]. We intend to compare the rate of success on this old
 1231 dataset with a more recent one. The second one, **RASTA** (Reusability of Android Static Tools
 1232 and Analysis), we built to cover all dates between 2010 and 2023. This dataset is a random
 1233 extract of Androzoo [3], for which we balanced applications between years and size. For each

1234 year and inter-decile range of size in Androzoo, 500 applications have been extracted with an
 1235 arbitrary proportion of 7% of malware. This ratio has been chosen because it is the ratio of
 1236 goodwill/malware that we observed when performing a raw extract of Androzoo. For checking
 1237 the maliciousness of an Android application, we rely on the VirusTotal detection indicators. If
 1238 more than 5 antiviruses have flagged the application as malicious, we consider it malware. If
 1239 no antivirus has reported the application as malicious, we consider it goodwill. Applications
 1240 in between are dropped.

1241 For computing the release date of an application, we contacted the authors of Androzoo to
 1242 compute the minimum date between the submission to Androzoo and the first upload to Virus-
 1243 Total. Such a computation is more reliable than using the DEX date that is often obfuscated
 1244 when packaging the application.

1245 3.3 Experiments

1246 3.3.1 RQ1: Re-Usability Evaluation

1247 Figure 4 and Figure 5 compare the Drebin and RASTA datasets. They represent the success/
 1248 failure rate (green/orange) of the tools. We distinguished failure to compute a result from
 1249 timeout (blue) and crashes of our evaluation framework (in grey, probably due to out-of-memory
 1250 kills of the container itself). Because it may be caused by a bug in our own analysis stack, exit
 1251 statuses represented in grey (Other) are considered as unknown errors and not as failures of
 1252 the tool. We discuss further errors for which we have information in the logs in Section 3.4.




Figure 4: Exit status for the Drebin dataset




Figure 5: Exit status for the RASTA dataset

1253 Results on the Drebin datasets show that 11 tools have a high success rate (greater than 85%).
 1254 The other tools have poor results. The worst, excluding Lotrack and Tresher, is Anadroid with
 1255 a ratio under 20% of success.

1256 On the RASTA dataset, we observe a global increase in the number of failed status: 12 tools
 1257 (54.55%) have a finishing rate below 50%. The tools that have bad results with Drebin are,
 1258 of course, bad results on RASTA. Three tools (androguard_dad, blueseal, saaf) that were
 1259 performing well (higher than 85%) on Drebin, surprisingly fall below the bar of 50% of failure.
 1260 7 tools keep a high success rate: Adagio, Amandroid, Androguard, Apparecium, Gator, Mallo-
 1261 droid, Redexer. Regarding IC3, the fork with a simpler build process and support for modern
 1262 OS has a lower success rate than the original tool.

1263 Two tools should be discussed in particular. Androguard has a high success rate, which is not
 1264 surprising: it is used by a lot of tools, including for analysing applications uploaded to the
 1265 Androzoo repository. Nevertheless, when using Androguard decompiler (DAD) to decompile
 1266 an APK, it fails more than 50% of the time. This example shows that even a tool that is
 1267 frequently used can still run into critical failures. Concerning Flowdroid, our results show a
 1268 very low timeout rate (0.06%), which was unexpected: in our exchanges, Flowdroid’s authors
 1269 were expecting a higher rate of timeout and fewer crashes.

1270 As a summary, the final ratio of successful analysis for the tools that we could run is 54.9%.
 1271 When including the two defective tools, this ratio drops to 49.9%.

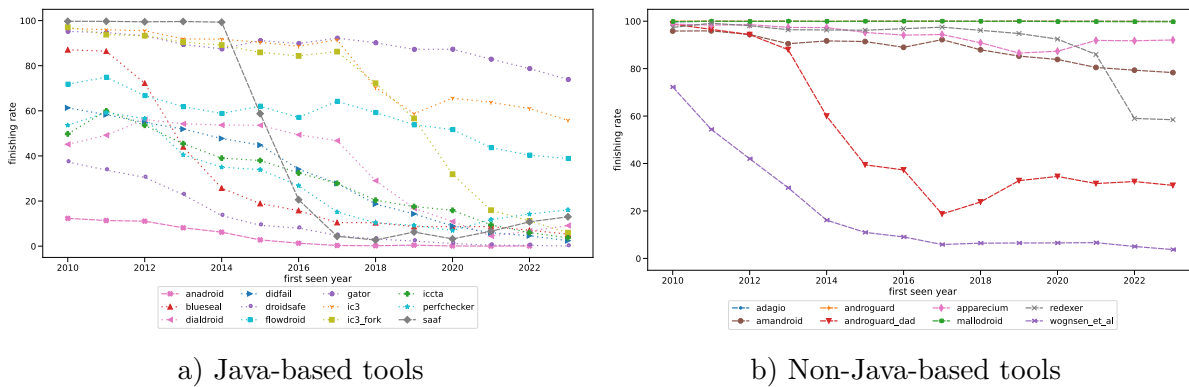


Figure 6: Exit status evolution for the RASTA dataset

1272 **RQ1 answer:** On a recent dataset, we consider that 54.55% of the tools are unusable. For
 1273 the tools that we could run, 54.9% of analyses are finishing successfully.

1274 3.3.2 RQ2: Size, SDK and Date Influence

1275 For investigating the effect of application dates on the tools, we computed the date of each
 1276 APK based on the minimum date between the first upload in AndroZoo and the first analysis
 1277 in VirusTotal. Such a computation is more reliable than using the DEX date, which is often
 1278 obfuscated when packaging the application. Then, for the sake of clarity of our results, we
 1279 separated the tools that have mainly Java source code from those that use other languages.
 1280 Among the ones that are Java-based programs, most of them use the Soot framework, which
 1281 may correlate the obtained results. Figure 6 a) (resp. Figure 6 b)) compares the success
 1282 rate of the tools between 2010 and 2023 for Java-based tools (resp. non Java-based tools).
 1283 For Java-based tools, a clear decrease in finishing rate can be observed globally for all tools.
 1284 For non-Java-based tools, 2 of them keep a high success rate (Androgard, Mallodroid). The
 1285 result is expected for Androgard, because the analysis is relatively simple and the tool is
 1286 largely adopted, as previously mentioned. Mallodroid, being a relatively simple script leveraging
 1287 Androgard, benefits from Androgard’s resilience. It should be noted that Saaf kept a high
 1288 success ratio until 2014 and then quickly decreased to less than 20% after 2014. This example
 1289 shows that, even with an identical source code and the same running platform, a tool can change
 1290 its behaviour over time because of the evolution of the structure of the input files.

1291 An interesting comparison is the specific case of IC3 and Ic3_fork. Until 2019, the success rate
 1292 was very similar. After 2020, ic3_fork is continuing to decrease, whereas IC3 keeps a success
 1293 rate of around 60%.

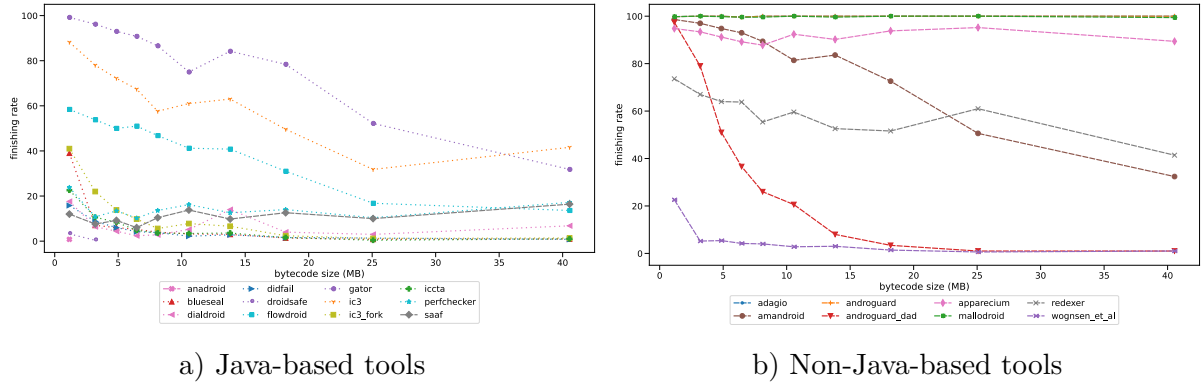


Figure 7: Finishing rate by bytecode size for APK detected in 2022

1294 To compare the influence of the date, SDK version and size of applications, we fixed one
 1295 parameter while varying another.

1296 *Fixed application year. (5000 APKs)* We selected the year 2022, which has a good amount of
 1297 representatives for each decile of size in our application dataset. Figure 7 a) (resp. Figure 7 b))
 1298 shows the finishing rate of the tools in function of the size of the bytecode for Java-based tools
 1299 (resp. non-Java-based tools) analysing applications of 2022. We can observe that all Java-based
 1300 tools have a finishing rate that decreases over the years. 50% of non-Java-based tools have the
 1301 same behaviour.

1302 *Fixed application bytecode size. (6252 APKs)* We selected the sixth decile (between 4.08 and
 1303 5.20 MB), which is well represented in a wide number of years. Figure 8 a) (resp. Figure 8 b))
 1304 represents the finishing rate depending on the year at a fixed bytecode size. We observe that 9

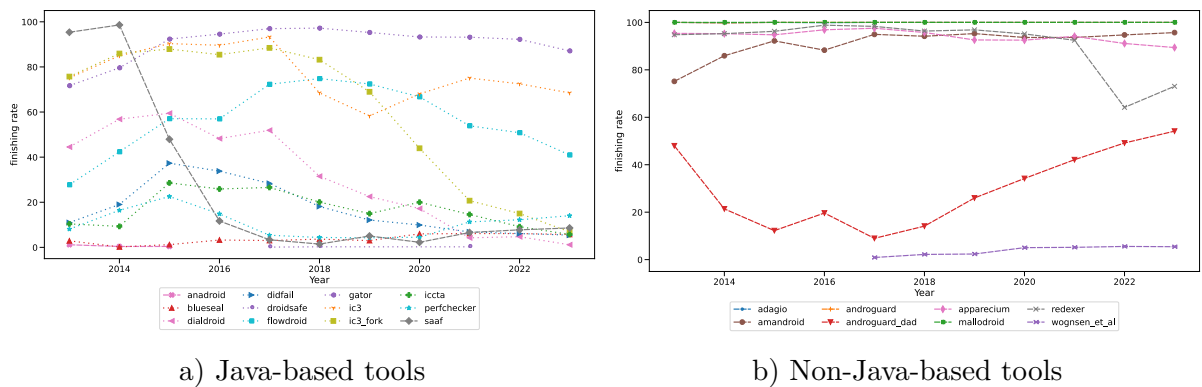


Figure 8: Finishing rate by discovery year with a bytecode size $\in [4.08, 5.2]$ MB

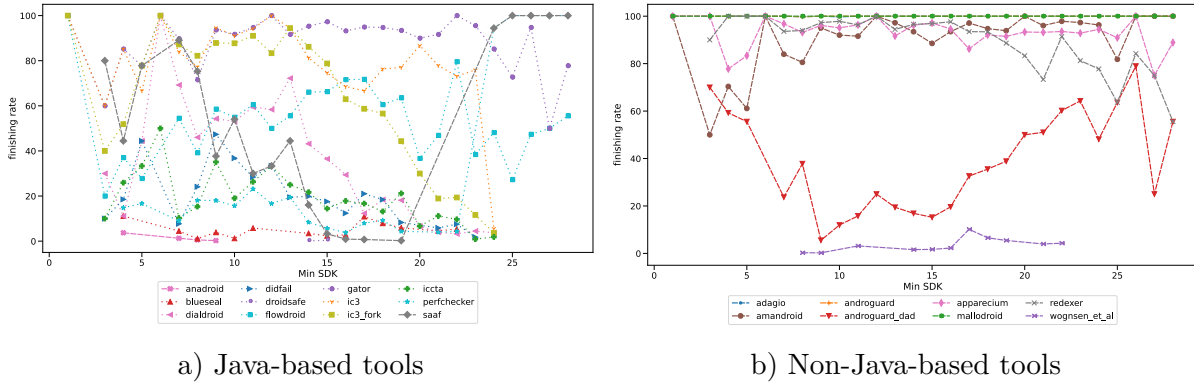


Figure 9: Finishing rate by min SDK with a bytecode size $\in [4.08, 5.2]$ MB

1305 tools out of 12 have a finishing rate dropping below 20% for Java-based tools, which is not the
 1306 case for non-Java-based tools.

1307 We performed similar experiments by varying the min SDK and target SDK versions, still with
 1308 a fixed bytecode size between 4.08 and 5.2 MB, as shown in Figure 9 a) and Figure 9 b). We
 1309 found that, contrary to the target SDK, the min SDK version has an impact on the finishing
 1310 rate of Java-based tools: 8 tools over 12 are below 50% after SDK 16. It is not surprising, as
 1311 the min SDK is highly correlated to the year.

1312 **RQ2 answer:** For the 20 tools that can be used partially, a global decrease in the success
 1313 rate of tools' analysis is observed over time. Starting at a 78% success rate, after five years,
 1314 tools have 61% success; after ten years, 45% success. The success rate varies based on the
 1315 size of the bytecode and SDK version. The date is also correlated with the success rate for
 1316 Java-based tools only.

1317 3.3.3 RQ3: Malware vs Goodware

RASTA part	Average size (MB)	
	APK	DEX
goodware	16.9	6.6
malware	17.2	4.3
total	16.9	6.5

Table 3: Average size and date of goodwill/malware parts of the RASTA dataset

1318 We sampled our dataset to have a variety of APK sizes, but the size of the application is not
 1319 entirely proportional to the bytecode size. Looking at Table 3, we can see that although malware
 1320 are, on average, bigger APKs, they contain less bytecode than goodwill. In the previous section,




Figure 10: Exit status comparing goodware (left bars) and malware (right bars) for the RASTA dataset

1321 we saw that the size of the bytecode has the most significant impact on the finishing rate of
 1322 analysis tools, and indeed, Figure 10 reflects that.

1323 In Figure 10, we compared the finishing rate of malware and goodware applications for the
 1324 evaluated tools. We can see that malware and goodware seem to generate a similar number of
 1325 timeouts. However, with the exception of two tools – apparecium and redexer, we can see a
 1326 trend of goodware being harder to analyse than malware. Some tools, like DAD or perfchecker,
 1327 show the finishing rate ratio augmented by more than 20 points.

Decile	Average DEX size (MB)		Finishing Rate: FR		Ratio Size Good/Mal	Ratio FR Good/Mal
	Good	Mal	Good	Mal		
1	0.13	0.11	0.85	0.82	1.17	1.04
2	0.54	0.55	0.74	0.72	0.97	1.03
3	1.37	1.25	0.63	0.66	1.09	0.97
4	2.41	2.34	0.57	0.62	1.03	0.92
5	3.56	3.55	0.53	0.59	1	0.9
6	4.61	4.56	0.5	0.61	1.01	0.82
7	5.87	5.91	0.47	0.57	0.99	0.83
8	7.64	7.63	0.43	0.56	1	0.76
9	11.39	11.26	0.39	0.58	1.01	0.67
10	24.24	21.36	0.33	0.46	1.13	0.73

Table 4: DEX size and Finishing Rate (FR) per decile

1328 We saw that the bytecode size may be an explanation for this increase. To investigate this
1329 further, Table 4 reports the bytecode size and the finishing rate of goodware and malware in
1330 each decile of bytecode size. We also computed the ratio of the bytecode size and finishing rate
1331 for the two populations. We observe that while the bytecode size ratio between goodware and
1332 malware stays close to one in each decile (excluding the two extremes), the goodware/malware
1333 finishing rate ratio decreases for each decile. It goes from 1.03 for the 2nd decile to 0.67 in the
1334 9th decile. We conclude from this table that, at equal size, analysing malware still triggers fewer
1335 errors than for goodware, and that the difference in errors generated between when analysing
1336 goodware and analysing malware increases with the bytecode size.

1337 **RQ3 answer:** Analysing malware applications triggers fewer errors for static analysis tools
1338 than analysing goodware for comparable bytecode size.

1339 3.4 Failures Analysis

1340 In this section, we investigate the reasons behind the high failure ratio presented in Section 3.3.
1341 Table 5 reports the average number of errors, the average time and memory consumption of
1342 the analysis of one APK file.

1343 3.4.1 Detected Errors

1344 During the running of our experiment, we parsed the standard output and error to capture:

- 1345 • Java errors and stack traces.
- 1346 • Python errors and stack traces.
- 1347 • Ruby errors and stack traces.
- 1348 • Log4j messages with a “ERROR” or “FATAL” level.
- 1349 • XSB error messages.
- 1350 • Ocaml errors.

1351 For example, Dialdroid reports an average of 55.9 errors for one successful analysis. On the
1352 contrary, some tools, such as Blueseal report very few errors at a time, making it easier to
1353 identify the cause of the failure.

1354 Because some tools send back a high number of errors in our logs (up to 46 698 for one
1355 execution), we tried to determine the error that is linked to the failed status. Unfortunately, our
1356 manual investigations confirmed that the last error of a log output is not always the one that
1357 should be attributed to the global failure of the analysis. The error that seems to generate the
1358 failure can occur in the middle of the execution, be caught by the code and then other subsequent
1359 parts of the code may generate new errors as consequences of the first one. Similarly, the first
1360 error in the logs is not always the cause of a failure. Sometimes errors successfully caught and
1361 handled are logged anyway. Thus, it is impossible to accurately extract the error responsible for

Exit status		adagio	amaandroid	anaandroid	androguard	androguard_dad	apparectum	bluesel	diandroid	didfail	droitsafe	flowdroid	gator	ic3	ic3_fork	iccta	malloandroid	perfchecker	redexer	ssaf	wogansen_et_al
Average number of errors (and standard deviation)																					
FINISHED	errors	0	0.9	0.02	0	0	3.33	0	55.88	1.13	7.17	1.04	0	1.22	0.32	14.3	4.37	0.28	1.29	0.22	2.13
	σ	0	3.23	0.19	0	0	261.92	0.05	63.73	3.94	37.87	26.32	0.04	23.1	2.73	71.74	277	1.77	1.28	0.83	66.5
FAILED	errors	0	2.79	2.34	1.35	1	21.63	1.02	33.79	6.6	12.53	14.64	0.32	3.66	1.29	17.34	1	1.15	3.45	6.35	4.11
	σ	0	8.7	0.94	0.48	0.02	466.97	0.21	108.56	31.56	74.01	49.07	0.78	18.06	0.71	42.81	0	4.7	4.52	22.97	48.81
TIMEOUT	errors	0	9.78	0.01	0	0	4.3	0.01	60.94	1.06	26.64	0.75	0	2.13	0.91	3.68	0	1.24	0	91.29	1.31
	σ	0	9.76	0.11	0	0	79.98	0.11	101.73	2.98	97.18	1.72	0	5.19	3.19	15.33	0	4.3	0	353.75	3.42
Average time (s)																					
FINISHED		17	405	149	16	27	98	158	768	270	676	29	33	156	159	90	28	4	16	56	696
FAILED	time	8	761	5	14	63	22	12	68	445	443	137	924	535	29	202	5	10	17	6	56
TIMEOUT		0	3601	3601	0	3604	3600	3601	3604	3600	3600	3601	3601	3601	3601	3600	0	3600	0	3602	3601
Average Memory (GB)																					
FINISHED		0.6	4.5	12.8	0.6	0.3	1.3	2.7	15.9	17.6	9.9	2	2	15.3	5	5	0	0	1	3	3
FAILED	memory	0.3	4.9	2.8	0.4	1	0.6	1.7	3.9	68.3	14.8	5	41.5	130.9	5	12	0	1	1	1	1
TIMEOUT		0	19	82.6	0	68.1	2.1	15.4	37.2	99.8	0.2	20.2	1.1	81.1	20	2	0	1	0	9	0

Table 5: Average number of errors, analysis time, memory per unitary analysis – compared by exit status

1362 a failed execution. Therefore, we investigated the nature of errors globally, without distinction
1363 between error messages in a log.

1364 Figure 11 draws the most frequent error objects for each of the tools. A black square is an error
1365 type that represents more than 80% of the errors raised by the considered tool. In between,
1366 grey squares show a ratio between 20% and 80% of the reported errors.

1367 First, the heatmap helps us to confirm that our experiment is running in adequate conditions.
1368 Regarding errors linked to memory, two errors should be investigated: `OutOfMemoryError` and
1369 `StackOverflowError`. The first one only appears for Gator with a low ratio. Several tools have
1370 a low ratio of errors concerning the stack. These results confirm that the allocated heap and
1371 stack are sufficient for running the tools with the RASTA dataset. Regarding errors linked to
1372 the disk space, we observe small ratios for the exception `IOException`, `FileNotFoundException` and
1373 `FileNotFoundException`. Manual inspections revealed that those errors are often a consequence
1374 of a failed Apktool execution.




Figure 11: Heatmap of the ratio of error reasons for all tools for the RASTA dataset

1375 Second, the black squares indicate frequent errors that need to be investigated separately. In
 1376 the next subsection, we manually analysed, when possible, the code that generates these high
 1377 ratios of errors, and we give feedback about the possible causes and difficulties in writing a
 1378 bug fix.

1379 3.4.2 Tool by Tool Investigation

1380 *Androguard and Androguard_dad* Surprisingly, while Androguard rarely fails to analyse an
 1381 APK, the internal decompiler of Androguard (DAD) fails more than half of the time. The

1382 analysis of the logs shows that the issue comes from the way the decompiled methods are stored:
1383 each method is stored in a file named after the method name and signature, and this file name
1384 can quickly exceed the size limit (255 characters on most file systems). It should be noticed that
1385 Androguard_dad rarely fails on the Drebin dataset. This illustrates the importance of testing
1386 tools on real and up-to-date APKs: even a bad handling of filenames can influence an analysis.

1387 *Malldroid and Apparecium* Malldroid and Apparecium stand out as the tools that raised the
1388 most errors in one run. They can raise more than 10 000 errors by analysis. However, it happened
1389 only for a few dozen APKs, and conspicuously, the same APKs raised the same high number
1390 of errors for both tools. The recurring error is a `KeyError` raised by Androguard when trying to
1391 find a string by its identifier. Although this error is logged, it seems successfully handled, and
1392 during a manual analysis of the execution, both tools seemingly perform their analysis without
1393 issue. This high number of occurrences may suggest that the output is not valid. Still, the tools
1394 claim to return a result, so, from our perspective, we consider those analyses as successful. For
1395 numerous other errors, we could not identify the reason why those specific applications raise
1396 so many exceptions. However, we noticed that Malldroid and Apparecium use an outdated
1397 version of Androguard (respectively version 3.0 and 2.0), and neither Androguard v3.3.5 nor
1398 DAD with Androguard v3.3.5 raise those exceptions. This suggests the issue has been fixed by
1399 Androguard and that Malldroid and Apparecium could benefit from a dependency upgrade.

1400 *Blueseal* Because Blue seal rarely logs more than one error when crashing, it is easy to identify
1401 the relevant error. The majority of crashes come from unsupported Android versions (due to
1402 the magic number of the DEX files not being supported by the version of back smali used by
1403 Blue seal) and methods whose implementation is not found (like native methods).

1404 *Droidsafe and SAAF* Our investigation of the most common errors raised by Droidsafe
1405 and SAAF showed that they are often preceded by an error from apktool. Indeed, 28
1406 654 runs of Droidsafe and 38 635 runs of SAAF failed after raising at least one of
1407 `brut.androlib.AndrolibException` or `brut.androlib.err.UndefinedResObject`, suggesting that
1408 those tools would benefit from an upgrade of apktool.

1409 *IC3 and IC3_fork* We compared the number of errors between IC3 and IC3_fork. IC3_fork
1410 reports fewer errors for all types of analysis, which suggests that the author of the fork has
1411 removed the outputted errors from the original code: the thrown errors are captured in a
1412 generic `RuntimeException`, which removes the semantics, making it harder for our investigations.
1413 Nevertheless, IC3_fork has more failures than IC3: the number of errors reported by a tool is
1414 not correlated to the final success of its analysis.

1415 *Flowdroid* Our exchanges with the authors of Flowdroid led us to expect more timeouts from
1416 executions taking too long than failed runs. Surprisingly, we only got 0.06% of timeout, and
1417 a high number of failures. We tried to detect recurring causes of failures, but the complexity

1418 of Flowdroid makes the investigation difficult. Most exceptions seem to be related to concur-
1419 rency. Other errors that came up regularly are `java.nio.channels.ClosedChannelException`,
1420 which is raised when Flowdroid fails to read from the APK, although we did not find the
1421 reason for failure, null pointer exceptions when trying to check if a null value is in a
1422 `ConcurrentHashMap` (in `LazySummaryProvider.getClassFlows()`) and `StackOverflowError` from
1423 `StronglyConnectedComponentsFast.recurse()`. We randomly selected 20 APKs that generated
1424 stack overflows in Flowdroid and retried the analysis with 500GB of RAM allocated to the
1425 JVM. 18 of those runs still failed with a stack overflow without using all the allocated memory,
1426 and the other two failed after raising null pointer exceptions from `getClassFlows`. This shows
1427 that the lack of memory is not the primary cause of those failures.

1428 As a conclusion, we observe that a lot of errors can be linked to bugs in dependencies. Our
1429 attempts to upgrade those dependencies led to new errors appearing: we conclude that this is
1430 not a trivial task that requires familiarity with the inner code of the tools.

1431 3.5 State-of-the-Art Comparison

1432 In this section, we will compare our results with the contributions presented in Chapter 2.

1433 Luo *et al.* released TaintBench [40], a real-world benchmark and the associated recommen-
1434 dations to build such a benchmark. These benchmarks confirmed that some tools, such as
1435 Amandroid and Flowdroid, are less efficient on real-world applications. We confirm the hypoth-
1436 esis of Luo *et al.* that real-world applications lead to less efficient analysis than using handcrafted
1437 test applications or old datasets [40]. In addition, even if Drebin is not hand-crafted, it is quite
1438 old and seems to present similar issues as handcrafted datasets when used to evaluate a tool:
1439 we obtained really good results compared to the RASTA dataset – which is more representative
1440 of real-world applications.

1441 Our findings are also consistent with the numerical results of Pauck *et al.* that showed that
1442 58.89% of DIALDroid-Bench [12] real-world applications are analysed successfully with the 6
1443 evaluated tools [49]. Six years after the release of DIALDroid-Bench, we obtain a lower ratio
1444 of 40.05% for the same set of 6 tools but using the RASTA dataset of 62 525 applications.
1445 We extended this result to a set of 20 tools and obtained a global success rate of 54.9%. We
1446 confirmed that most tools require a significant amount of work to get them running [53]. Our
1447 investigations of crashes also confirmed that dependencies on older versions of Apktool are
1448 impacting the performances of Anadroid, Saaf and Wognsen *et al.* in addition to DroidSafe and
1449 IccTa, already identified by Pauck *et al.*.

1450 Third, we extended to 20 different tools the work done by Reaves *et al.* on the usability of
1451 analysis tools (4 tools are in common, we added 16 new tools and two variations). We confirmed
1452 that most tools require a significant amount of work to get them running. We encounter similar
1453 issues with libraries and operating system incompatibilities, and noticed that, as time passes,

1454 dependency issues may impact the build process. For instance, we encountered cases where the
1455 repositories hosting the dependencies were closed, or cases where Maven failed to download
1456 dependencies because the OS version did not support SSL, which is now mandatory to access
1457 Maven Central.

1458 **3.6 Recommendations**

1459 In light of our findings in Section 3.4 and the issues we encountered while packaging the tools,
1460 we summarise some takeaways that we believe developers should follow to improve the success
1461 of reusing their software.

1462 We understand software developed for research purposes is not and should not be held to
1463 the same standards as production software. However, research is incremental, and it is not
1464 sustainable to start each tool from scratch. It is critical to be able to build upon tools already
1465 published, and efforts should be made to allow that when releasing a tool.

1466 During the packaging and testing of the tools we examined in our experiment, the most notable
1467 issues we encountered could have been avoided by following classical development best practices.

1468 To make a tool easy to reuse, it should have documentation with at least:

- 1469 • Instructions about how to install the dependencies.
- 1470 • Instructions about how to build the tool (if the tool needs to be built).
- 1471 • Instructions about how to use the tool (*e.g.*, command line arguments).
- 1472 • Instructions about how to interpret the results of the tools (we only checked for the existence
1473 of the results in our experiment, but we found that some results can be quite obscure).

1474 In addition to the documentation, a minimum working example with the expected result of the
1475 tools allows a potential user to check if everything is working as intended. This MWE have the
1476 additional benefit that it can serve as an example in the documentation.

1477 Another best practice to follow is to pin the version of dependencies of the tools. Many modern
1478 dependency management tools can handle that: for instance, for Python, Poetry or uv generate
1479 a lock file with the exact version of the libraries to use, Cargo does the same for Rust, in
1480 Java, this can be an option in Gradle, and dependencies in Maven `pom.xml` files are usually
1481 the exact version. For other dependencies that are not managed by a dependency manager
1482 – for instance, the Java virtual machine to use, the Python interpreter, resource files – the
1483 version to use should be clearly documented. Alternatively, tools like `nixpkg` can be used to
1484 pin every dependency. The worst case we encountered during our experiment was a tool whose
1485 documentation instructed us to install the `z3` dependencies with a simple `git clone`, without
1486 specifying the commit to use. The `z3` project is still actively maintained, so the dependency
1487 installed was not compatible, and finding a compatible version required checking releases one
1488 by one. Dependencies fetched with a version control system should always indicate the exact
1489 version to use (in the case of `git`, a commit, tag, or release should be used).

1490 We also found that interactions with the running environment can become very problematic
1491 when the environment changes. To minimise the issues, packaging the tool inside a Docker
1492 container or even a virtual machine can ensure that future users have at least access to a working
1493 version of the tool.

1494 Finally, when possible, continuous integration, tests and code reviews should be implemented
1495 to improve the reliability of the developed tool.

1496 Concerning the actual code of the tool, more attention should be paid to error reporting. When
1497 a tool failed to perform its analysis, it should be clear to the user, and the reason should be
1498 clearly reported. In some cases, this may imply *not* trying to recover from unrecoverable errors:
1499 this often leads to errors seemingly unrelated to the initial issue. This is often a problem in
1500 Java code, where the developers are strongly encouraged to catch all exceptions, and in bash
1501 scripts that run several programs in a row without checking the exit statuses.

1502 Good error reporting can allow future users to solve issues encountered using the tools: for
1503 instance, the log generated by Androguard's decompiler clearly shows that the issue is file names
1504 exceeding the size limit. This issue could easily be fixed by changing the filenames used to store
1505 the results. In contrast, the errors generated by Flowdroid are so opaque that we have no idea
1506 how we could solve them.

1507 At last, an important remark concerns the libraries used by a tool. We have seen two types of
1508 libraries:

- 1509 • internal libraries manipulating internal data of the tool.
- 1510 • external libraries that are used to manipulate the input data (APKs, bytecode, resources).

1511 We observed during our manual investigations that external libraries are the ones leading to
1512 crashes because of variations in recent APKs (file format, unknown bytecode instructions, multi-
1513 DEX files). We believe that the developer should provide enough documentation to make a later
1514 upgrade of these external libraries possible. For example, old versions of Apktool are the top-
1515 most libraries raising errors, but breaking changes introduced by upgrading from v1.X versions
1516 to v2.X versions prevent us from upgrading Apktool.

1517 **3.7 Limitations**

1518 Some limitations of our approach should be kept in mind.

1519 Our application dataset is biased in favour of Androguard, because Androzoo has already used
1520 Androguard internally when collecting applications and discarded any application that cannot
1521 be processed with this tool.

1522 Despite our best efforts, it is possible that we made mistakes when building or using the tools.
1523 It is also possible that we wrongly classified a result as a failure. To mitigate this possible

1524 problem, we contacted the authors of the tools to confirm that we used the right parameters
1525 and chose a valid failure criterion. Before running the final experiment, we also ran the tools
1526 on a subset of our dataset and manually investigated the most common errors to ensure that
1527 they are not trivial errors that can be solved.

1528 The timeout value and memory limits are arbitrarily fixed. To mitigate this issue, a small
1529 extract of our dataset has been analysed with more memory/time, and we checked that there
1530 was no significant difference in the results.

1531 Finally, the use of VirusTotal for determining if an application is malware or not may be wrong.
1532 To limit the impact of errors, we used a threshold of at most 5 antiviruses (resp. no more
1533 than 0) reporting an application as being malware (resp. goodware) for taking a decision about
1534 maliciousness (resp. benignness).

1535 **3.8 Future Works**

1536 A first extension to this work would obviously be to study more tools. We restricted ourselves
1537 to the tools listed by Li *et al.*, but it would be interesting to compare our result to the finishing
1538 rate of recently released tools. It would be interesting to see if they are better at handling large
1539 APKs, but also to see if older applications are more challenging for them due to discontinued
1540 features.

1541 Another avenue would be to define a benchmark to check the ability of tools to handle real-
1542 world applications. Our dataset is too large for a simple benchmark and is sampled to have a
1543 variety of application sizes and years of publication. Hence, the first step would be to sample
1544 a dataset for this benchmark. Current benchmark datasets focus on the accuracy of the tested
1545 tools, with difficult-to-analyse applications. It could be interesting to extract from our results
1546 some of the applications that the most tools failed to analyse, and either use them directly or
1547 study them to craft simpler applications reproducing the same challenges as those applications.
1548 Such datasets would need to be updated regularly: we saw that there is a trend for newer
1549 applications to be harder to analyse, a frozen dataset would ignore this factor.

1550 In addition to the finishing rate, it would be both interesting and useful to have reference values.
1551 Table 6 list common Android-related dependencies we encountered when packaging the tools.
1552 We can see that each tool uses at least one of those dependencies. It would be reasonable to
1553 consider the best finishing ratio a tool can have to be the finishing ratio of a tool that would
1554 perform an “empty analysis” using the same dependencies. Considering the prevalence of those
1555 dependencies, having those theoretical minimums could also guide future tool developers when
1556 choosing their dependencies.

Tool	Soot	Androguard	Apktool
adagio		✓	
amandroid			✓
anandroid			✓
androguard		✓	
androguard_dad		✓	
apparecium		✓	
blueseal	✓		✓
dialdroid	✓		
didfail	✓		
droidsafe	✓		✓
flowdroid	✓		
gator	✓		✓
ic3	✓		
ic3_fork	✓		
iccta	✓		✓
malloandroid		✓	
perfchecker	✓		
redexer			✓
saaf			✓
wognsen_et_al			✓

Table 6: Commonly found dependencies

1557 3.9 Conclusion

1558 Since the release of Android, many tools have been published in order to analyse Android
 1559 applications. In Chapter 2, we went through contributions that benchmark and compare some
 1560 of those tools. Those contributions suggested that analysing real-world applications might be
 1561 more challenging than expected. This led us to question the reusability of those tools (**Pb1**).

1562 This chapter has assessed the suggested results of the literature [40, 49, 53] about the reliability
 1563 of static analysis tools for Android applications. With a dataset of 62 525 applications, we
 1564 established that 54.55% of 22 tools are not reusable. 2 of those were due to the fact that we did
 1565 not manage to use the tools, even with the help of the author. We consider the 10 other tools
 1566 to be unusable due to the fact that they fail to finish their analysis more than 50% of the time..
 1567 In total, the analysis success rate of the tools that we could run for the entire dataset is 54.9%.
 1568 The characteristics that have the most influence on the success rate are the bytecode size and
 1569 the min SDK version. Finally, we showed that malware APKs generate fewer fatal errors than
 1570 goodware when analysed.

1571 Following Reaves *et al.* recommendations [53], we publish the Docker and Singularity images we
 1572 built to run our experiments alongside the Docker files. This will allow the research community
 1573 to use the tools directly without the build and installation penalty.

1574
1575

Pb1: *To what extent are previously published Android analysis tools still usable today, and what factors impact their reusability?*

1576
1577
1578
1579
1580
1581
1582

More than half the tools we selected were not usable. In some cases, it was due to our inability to set up the tool correctly. Mostly, it was due to the high failure rate when analysing real-world applications. Results show that large applications cause more crashes, as do applications with a higher min SDK target. Goodware also appears to generate more analysis failures than malware.

1584 **CLASS LOADERS IN THE MIDDLE: CON-** 1585 **FUSING ANDROID STATIC ANALYSERS**

1586 Things that try to look like things often do look more like things than things.

1587 — Esmerelda Weatherwax, Wyrd Sisters, Terry Pratchett

1588 The dynamic linking and loading of the different classes by the
1589 ART is a complex task that can eventually be exploited by
1590 an attacker. In particular, if the developer adds a class whose
1591 name collides with the name of a class of the Android operating
1592 system or another class in the application, they may confuse a
1593 reverse engineer in charge of studying such an application. In
1594 this chapter, we explore the consequences of those collisions. We
1595 highlight three attacks that we call shadow attacks because the
1596 class implementation that a reverser would find shadows a second
1597 implementation with a higher priority. In particular, we show that
1598 static analysis tools used by a reverser choose the shadow imple-
1599 mentation for most of the evaluated tools, and output a wrong
1600 result. In a dataset of 49 975 applications, we also investigate
1601 whether shadow attacks are used in the wild and show that, most
1602 of the time, there is no malicious behaviour behind them.

1603 **4.1 Introduction**

1604 In this chapter, we study how Android handles the loading of classes in the case of multiple
1605 versions of the same class. Such collisions can exist inside the APK file or between the APK

1606 file and Android SDK classes. We intend to understand if a reverser would be impacted during
1607 a static analysis when dealing with such obfuscated code. Because this problem is already
1608 complex enough with the current operations performed by Android, we exclude the case where
1609 the application uses class loaders dynamically (*e.g.*, dynamic code loading). We present a new
1610 technique that “shadows” a class, *i.e.*, embeds a class in the APK file and “presents” it to the
1611 reverser instead of the legitimate version. We show how these attacks can confuse the tools of
1612 the reverser when he performs a static analysis, and, in order to evaluate if such attacks are
1613 already used in the wild, we analysed 49 975 applications from 2023 that we extracted randomly
1614 from AndroZoo [3].

1615 The chapter is structured as follows. Section 4.2 investigates the internal mechanisms of class
1616 loading and presents how a reverser can be confused by these mechanisms. Then, in Section 4.3,
1617 we design obfuscation techniques and show their effect on static analysis tools. Next, Section 4.4
1618 evaluates whether these obfuscation techniques are used in the wild by scanning 49 975 APKs.
1619 Section 4.6 extends on the possible countermeasures against those shadow attacks, how they
1620 interact with other obfuscation techniques, as well as the limitations of this work and avenues
1621 left to explore. Finally, Section 4.6 concludes the chapter.

1622 4.2 Analysing the Class Loading Process

1623 For building obfuscation techniques based on the confusion of tools with class loaders, we
1624 manually studied the code of Android that handles class loading. In this section, we report
1625 the inner workings of ART, and we focus on the specificities of class loading that can bring
1626 confusion. Because the class loading implementation has evolved over time during the multiple
1627 iterations of the Android operating system, we mainly describe the behaviour of ART from
1628 Android version 14 (SDK 34).

1629 4.2.1 Class Loaders

1630 When ART needs to access a class, it queries an object implementing the `ClassLoader` class to
1631 retrieve its implementation. Each class has a reference to the `ClassLoader` that loaded it, and
1632 this class loader is the one that will be used to load supplementary classes used by the original
1633 class. For example, in Listing 1, when calling `A.f()`, the ART will load `B` with the class loader
1634 that was used to load `A`.

```
1 class A {  
2     public static void f() {  
3         B b = new B();  
4         b.do_something();  
5     }}
```

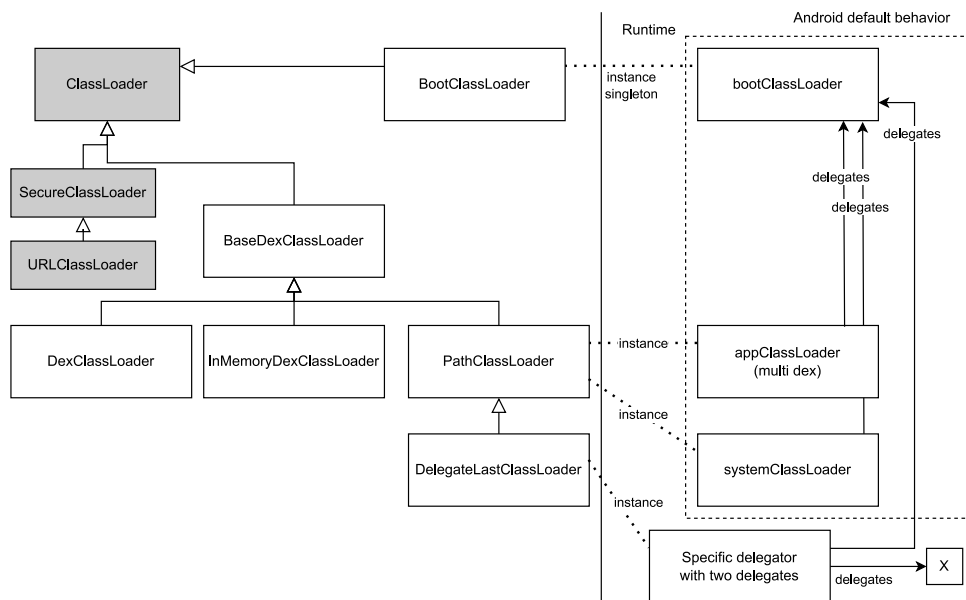
Listing 1: Class instantiation

1635 This behaviour has been inherited from Java, and most of the core classes regarding class loaders
 1636 have been kept in Android. Nevertheless, the Android implementation has slight differences,
 1637 and new class loaders have been added. For example, the Java class loader `URLClassLoader` is
 1638 still present in Android, but contrary to the official documentation, most of its methods have
 1639 been removed or replaced by a stub that just raises an exception. Moreover, rather than using
 1640 the Java class loaders `SecureClassLoader` or `URLClassLoader`, Android has several new class
 1641 loaders that inherit from `ClassLoader` and override the appropriate methods.

1642 The left part of Figure 12 shows the different class loaders specific to Android in white and
 1643 the stubs of the original Java class loaders in grey. The main difference between the original
 1644 Java class loaders and the ones used by Android is that they do not support the Java bytecode
 1645 format. Instead, the Android-specific class loaders load their classes from (many) different file
 1646 formats specific to Android. Usually, when used by a programmer, the classes are loaded from
 1647 memory or from a file using the DEX format (`.dex`). When used directly by ART, the classes
 1648 are usually stored in an application file (`.apk`) or in an optimised format (`OAR/ODEX`).

1649 4.2.2 Delegation

1650 The order in which classes are loaded at runtime requires special attention. All the specific
 1651 Android class loaders (`DexClassLoader`, `InMemoryClassLoader`, etc.) have the same behaviour
 1652 (except `DelegateLastClassLoader`), but they handle specificities for the input format. Each



grey – Java-based, white – Android-based
 Figure 12: The class loading hierarchy of Android

1653 class loader has a delegate class loader, represented in the right part of Figure 12 by black
1654 plain arrows for an instance of `PathClassLoader` and an instance of `DelegatLastClassLoader`
1655 (the other class loaders also have this delegate). This delegate is a concept specific to class
1656 loaders and has nothing to do with class inheritance. By default, class loaders will delegate
1657 to the singleton class `BootClassLoader`, except if a specific class loader is provided when
1658 instantiating the new class loader. When a class loader needs to load a class, except for
1659 `DelegatLastClassLoader`, it will first ask the delegate, i.e. `BootClassLoader`, and if the delegate
1660 does not find the class, the class loader will try to load the class on its own. This behaviour
1661 implements a priority and avoids redefining by error a core class of the system, for example,
1662 redefining `java.lang.String` that would be loaded by a child class loader instead of its delegates.
1663 `DelegatLastClassLoader` behaves slightly differently: it will first delegate to `BootClassLoader`,
1664 then it will check its files, and finally, it will delegate to its actual delegate (given when instan-
1665 tiating the `DelegatLastClassLoader`). This behaviour is useful for overriding specific classes of
1666 a class loader while keeping the other classes. A normal class loader would prioritise the classes
1667 of its delegate over its own.

1668 At runtime, Android instantiates for each application three instances of class loaders described
1669 previously: `bootClassLoader`, the unique instance of `BootClassLoader`, and two instances of
1670 `PathClassLoader`: `systemClassLoader` and `appClassLoader`. `bootClassLoader` is responsible for
1671 loading Android **platform classes**. It is the direct delegate of the two other class loaders
1672 instantiated by Android. `appClassLoader` points to the application `.apk` file, and is used to
1673 load the classes inside the application `systemClassLoader` is a `PathClassLoader` pointing to
1674 `'.'`, the working directory of the application, which is `'/'` by default. The documentation
1675 of `ClassLoader.getSystemClassLoader` reports that this class loader is the default context
1676 class loader for the main application thread. In reality, the platform classes are loaded
1677 by `bootClassLoader` and the classes from the application are loaded from `appClassLoader`.
1678 `systemClassLoader` is never used in production according to our careful reading of the AOSP
1679 sources.

1680 In addition to the class loaders instantiated by ART when starting an application, the developer
1681 of an application can use class loaders explicitly by calling ones from the Android SDK, or by
1682 recoding custom class loaders that inherit from the `ClassLoader` class. At this point, accurately
1683 modelling the complete class loading algorithm becomes impossible: the developer can program
1684 any algorithm of their choice. For this reason, this case is excluded from this chapter, and we
1685 focus on the default behaviour where the context class loader is the one pointing to the `.apk`
1686 file and where its delegate is `BootClassLoader`. With such a hypothesis, the delegation process
1687 can be modelled by the pseudo-code of method `load_class` given in Listing 2.

1688 In addition, it is important to distinguish the two types of platform classes handled by
1689 `BootClassLoader` and that both have priority over classes from the application at runtime:

```

1 def get_mutli_dex_classes_dex_name(index: int):
2     if index == 0:
3         return "classes.dex"
4     else:
5         return f"classes{index+1}.dex"
6
7 def load_class(class_name: str):
8     if is_platform_class(class_name):
9         return load_from_boot_class_loader(class_name)
10    else:
11        index = 0
12        dex_file = get_mutli_dex_classes_dex_name(index)
13        while file_exists_in_apk(dex_file) and \
14            not class_found_in_dex_file(class_name, dex_file):
15            index += 1
16            dex_file = get_mutli_dex_classes_dex_name(index)
17        if file_exists_in_apk(dex_file):
18            return load_from_file(dex_file, class_name)
19        else:
20            raise ClassNotFoundError()

```

Listing 2: Default Class Loading Algorithm for Android Applications

- 1690 • the ones available in the **Android SDK** (normally visible in the documentation).
1691 • the ones that are internal and that should not be used by the developer. We call them
1692 **hidden classes** [25, 34] (not documented).

1693 As a preliminary conclusion, we observe that a priority exists in the class loading mechanism
1694 and that an attacker could use it to prioritise an implementation over another one. This could
1695 mislead the reverser if they use the one that has the lowest priority. To determine if a class is
1696 impacted by the priority given to `BootClassLoader`, we need to obtain the list of classes that are
1697 part of Android *i.e.*, the platform classes. We discuss in the next section how to obtain these
1698 classes from the emulator.

1699 4.2.3 Determining Platform Classes

1700 Figure 13 shows how classes of Android are used in the development environment and at
1701 runtime. In the development environment, Android Studio uses `android.jar` and the specific
1702 classes written by the developer. After compilation, only the classes of the developer, and
1703 sometimes extra classes computed by Android Studio, are zipped in the APK file, using the
1704 multi-dex format. At runtime, the application uses `BootClassLoader` to load the platform classes
1705 from Android. Until our work, previous works [25, 34] considered both Android SDK and
1706 hidden classes to be in the file `/system/framework/framework.jar` found in the phone itself,

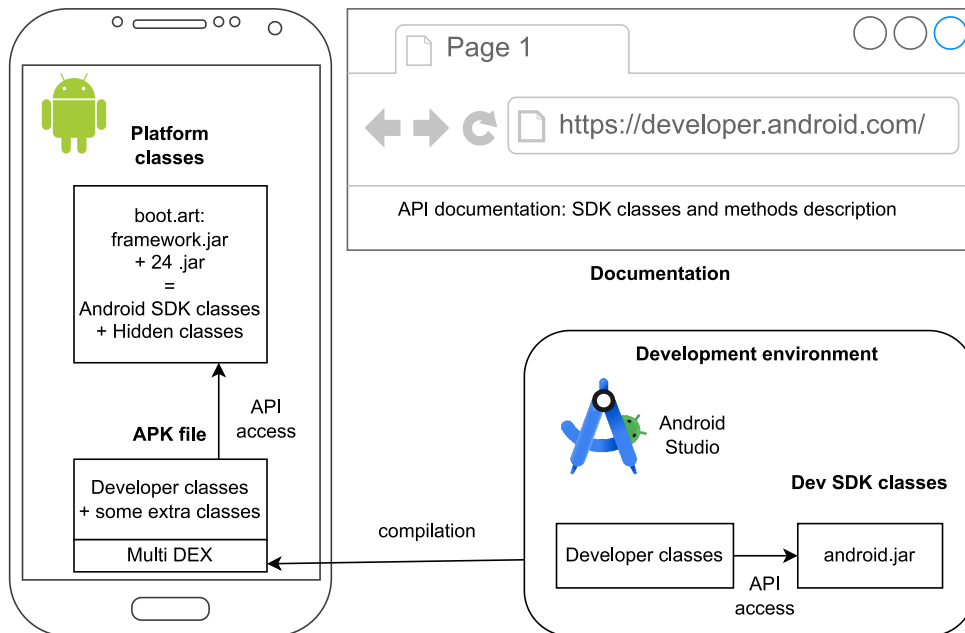


Figure 13: Location of SDK classes during development and at runtime

1707 but we found that the classes loaded by `bootClassLoader` are not all present in `framework.jar`.
 1708 For example, He *et al.* [25] counted 495 thousand APIs (fields and methods) in Android 12,
 1709 based on Google documentation on restrictions for non SDK interfaces¹. However, when looking
 1710 at the content of `framework.jar`, we only found 333 thousand APIs. Indeed, classes such as
 1711 `com.android.okhttp.OkHttpClient` are loaded by `bootClassLoader`, listed by Google, but not
 1712 in `framework.jar`.

1713 For optimisation purposes, classes are now loaded from `boot.art`. This file is used to speed up the
 1714 start-up time of applications: it stores a dump of the C++ objects representing the **platform**
 1715 **classes** (Android SDK and hidden classes) so that they do not need to be generated each time
 1716 an application starts. Unfortunately, this format is not documented and not retro-compatible
 1717 between Android versions and is thus difficult to parse. An easier solution to investigate the
 1718 platform classes is to look at the `BOOTCLASSPATH` environment variable in an emulator. This
 1719 variable is used to load the classes without the `boot.art` optimisation. We found 25 `.jar` files,
 1720 including `framework.jar`, in the `BOOTCLASSPATH` of the standard emulator for Android 12 (SDK
 1721 32), 31 for Android 13 (SDK 33), and 35 for Android 14 (SDK 35), containing respectively a total
 1722 of 499 837, 539 236 and 605 098 API methods and fields. Table 7 summarises the discrepancies
 1723 we found between Google’s list and the platform classes we found in Android emulators. Note
 1724 also that some methods may also be found *only* in the documentation. Our manual investigations

1725 1. <https://developer.android.com/guide/app-compatibility/restrictions-non-sdk-interfaces>

SDK version	Number of API methods			
	Documented	In emulator	Only documented	Only in emulator
32	495 713	499 837	1060	5184
33	537 427	539 236	1258	3067
34	605 106	605 098	26	18

Table 7: Comparison of API methods between documentation and emulators

1726 suggest that the documentation is not well synchronised with the evolution of the platform
1727 classes and that Google has almost solved this issue in API 34.

1728 We conclude that it can be dangerous to trust the documentation and that gathering infor-
1729 mation from the emulator or phone is the only reliable source. Gathering the precise list of
1730 classes and the associated bytecode is not a trivial task.

1731 4.2.4 Multiple DEX Files

1732 For the application class files, Android uses its specific format called DEX: all the classes
1733 of an application are loaded from the file `classes.dex`. With the increasing complexity of
1734 Android applications, the need arose to load more methods than the DEX format could
1735 support in one `.dex` file. To solve this problem, Android started storing classes in multiple
1736 files named `classesX.dex` as illustrated by the Listing 3 that generates the filenames read by
1737 class loaders. Android starts loading the file `GetMultiDexClassesDexName(0)` (`classes.dex`), then
1738 `GetMultiDexClassesDexName(1)` (`classes2.dex`), and continues until finding a value `n` for which
1739 `GetMultiDexClassesDexName(n)` does not exist. Even if Android emits a warning message when
1740 it finds more than 100 `.dex` files, it will still load any number of `.dex` files that way. This change
1741 has the unintended consequence of permitting two classes with the same name but different
1742 implementations to be stored in the same `.apk` file using two `.dex` files (*e.g.*, the class `Foo` can
1743 be defined both in `classes.dex` and `classes2.dex`).

1744 Android explicitly performs checks that prevent several classes from using the same name inside
1745 a `.dex` file. However, this check does not apply to multiple `.dex` files in the same `.apk` file, and
1746 a `.dex` can contain a class with a name already used by another class in another `.dex` file of the
1747 application. Of course, such a situation should not happen when multiple `.dex` files have been
1748 generated properly by Android Studio. Nevertheless, for an attacker controlling the process,
1749 this issue raises the question of which class is selected when several classes sharing the same
1750 name are present in `.apk` files.

1751 We found that Android loads the class whose implementation is found first when looking in the
1752 order of multiple `dexfiles`, as generated by the method `GetMultiDexClassesDexName`. We will

```

1  std::string DexFileLoader::GetMultiDexClassesDexName(size_t index) {
2      return (index == 0) ?
3          "classes.dex" :
4          StringPrintf("classes%zu.dex", index + 1);
5  }

```

Listing 3: The method generating the .dex filenames from the AOSP

1753 show later in Section 4.3.2 that this choice is not the most intuitive and can lead to fooling
 1754 analysis tools when reversing an application. As a conclusion, we model both the multi-dex and
 1755 delegation behaviours in the pseudo-code of Listing 2.

1756 4.3 Obfuscation Techniques

1757 In this section, we present new obfuscation techniques that take advantage of the complexity of
 1758 the class loading process. Then, in order to evaluate their efficiency, we reviewed some common
 1759 Android reverse analysis tools to see how they behave when collisions occur between classes
 1760 of the APK or between a class of the APK and classes of Android (Android SDK or hidden
 1761 classes). We call this collision “**class shadowing**”, because the attacker’s version of the class
 1762 shadows the one that will be used at runtime. To evaluate if such shadow attacks are working,
 1763 we handcrafted three applications implementing shadowing techniques to test their impact on
 1764 static analysis tools. Then, we manually inspected the output of the tools in order to check
 1765 their consistency with what Android is really doing at runtime. For example, for Apktool, we
 1766 look at the output disassembled code, and for Flowdroid [8], we check that a flow between
 1767 `Taint.source()` and `Taint.sink()` is correctly computed.

1768 4.3.1 Obfuscation Techniques

1769 From the results presented in Section 4.2, three approaches can be designed to hide the
 1770 behaviour of an application.

1771 **Self shadow:** *shadowing a class with another from APK* This method consists of hiding the
 1772 implementation of a class with another one by exploiting the possible collision of class names,
 1773 as described in Section 4.2.4 with multiple .dex files. If reversers or tools ignore the priority
 1774 order of a multi-dex file, they can take into account the wrong version of a class.

1775 **SDK shadow:** *shadowing a SDK class* This method consists of presenting to the reverser a
 1776 fake implementation of a class of the SDK. This class is embedded in the APK file and has the
 1777 same name as one of the SDK. Because `BootClassLoader` will give priority to the Android SDK
 1778 at runtime, the reverser or tool should ignore any version of a class that is contained in the
 1779 APK. The only constraint when shadowing an SDK class is that the shadowing implementation
 1780 must respect the signature of real classes. Note that, by introducing a custom class loader, the
 1781 attacker could invert the priority, but this case is out of the scope of this chapter.

```
1 public class Main {
2     public static void main(Activity ac) {
3         String personal_data = Taint.source();
4         String obfuscated_personal_data = Obfuscation.hide_flow(personal_data);
5         Taint.sink(ac, obfuscated_personal_data);
6     }
7 }
8
9 // customised for each obfuscation technique
10 public class Obfuscation {
11     public static String hide_flow(String personal_data) { ... }
12 }
```

Listing 4: Main body of test apps

1782 **Hidden shadow:** *shadowing a hidden class* This method is similar to the previous one,
1783 except the class that is shadowed is a hidden class. Because ART will give priority to the
1784 internal version of the class, the version provided in the APK file will be ignored. Such shadow
1785 attacks are more difficult to detect by a reverse engineer, who may not know the existence of
1786 this specific hidden class in Android.

1787 4.3.2 Impact on Static Analysis Tools

1788 We selected tools that are commonly used to unpack and reverse Android applications. In
1789 Chapter 3 (Section 3.2.2), we found only two tools to be still actively maintained: Androguard¹
1790 and Flowdroid². We also noticed that Apktool³ was a common dependency for a lot of the tools
1791 we tested in Chapter 3 (see Table 6), and is still used today. Consequently, we will test the
1792 impact of shadow attacks on those three tools. Lastly, because it is a state-of-the-art decompiler
1793 for Android applications, we added Jadx⁴ to the list of tools we tested.

1794 To evaluate the tools, we designed a single application that we can customise for different tests.
1795 Listing 4 shows the main body implementing:

- 1796 • a possible flow to evaluate FlowDroid: a flow from a method `Taint.source()` to a method
1797 `Taint.sink(Activity, String)` through a method `Obfuscation.hide_flow(String)`.
- 1798 • a possible use of a SDK or hidden class inside the class `Obfuscation` to evaluate platform
1799 classes shadowing for other tools.

1800 We used 4 versions of this application:

1801 1. <https://github.com/androguard/androguard>
1802 2. <https://github.com/secure-software-engineering/FlowDroid>
1803 3. <https://apktool.org/>
1804 4. <https://github.com/skylot/jadx>

- 1805 1. A control application that does not do anything special: `Obfuscation.hide_flow(String`
 1806 `personal_data)` returns `personal_data`. It will be used for checking the expected result of
 1807 tools.
- 1808 2. A version that implements self-shadowing: the class `Obfuscation` is duplicated: one is the
 1809 same as the one in the control app (`Obfuscation.hide_flow(String)` returns its arguments),
 1810 and the other version returns a constant string. These two versions are embedded in several
 1811 DEX of a multi-dex application.
- 1812 3. The third version implements SDK shadowing and needs an existing class of the SDK. We
 1813 used the SDK class `Pair` as the class to shadow. We put data in a new `Pair` instance and
 1814 reread the data from the `Pair`. The colliding `Pair` class we created discards the data at the
 1815 initialisation and stores `null` instead of the argument values. This decoy class breaks the
 1816 flow of information: Flowdroid will detect the information flow if it uses the actual SDK
 1817 implementation of `Pair` to compute the DFG, but not if it uses the decoy.
- 1818 4. The last version tests for Hidden API shadowing. Like for the third one, we similarly store
 1819 data in `com.android.okhttp.Request` and then retrieve it. Again, the shadowing implemen-
 1820 tation discards the data.

1821 We used the 4 selected tools on the 4 versions of the application and compared the results
 1822 on the control application to the results on the other application implementing the different
 1823 obfuscation techniques. We found that these static analysis tools do not consider the class
 1824 loading mechanism, either because the tools only look at the content of the application file
 1825 (*e.g.*, a disassembler) or because they consider class loading to be a dynamic feature and thus
 1826 out of their scope. In Table 8, we report on the types of shadowing that can trick each tool.
 1827 A plain circle is a shadow attack that leads to a wrong result. A white circle indicates a tool

Tool	Version	Shadow Attack		
		Self	SDK	Hidden
Jadx	1.5.0	○	●	●
Apktool	2.9.3	○	●	●
Androguard	4.1.2	○	●	●
Flowdroid	2.13.0	●	×	●

●: working

○: works but producing warning or can be seen by the reverser

×: not working

Table 8: Working attacks against static analysis tools

1828 emitting warnings or that displays the two versions of the class. A cross is a tool not impacted
1829 by a shadow attack.

1830 **4.3.2.1 Jadx**

1831 Jadx processes all the classes present in the application, but only saves/displays one class by
1832 name, even if two versions are present in multiple `.dex` files. Nevertheless, when multiple classes
1833 with the same name are found, Jadx reports it in a comment added to the generated Java source
1834 code. This warning stipulates that a possible collision exists and lists the files that contain the
1835 different versions of the class. Unfortunately, after reviewing the code of Jadx, we believe that
1836 the selection of the displayed class is an undefined behaviour. At least for version 1.5.0 that we
1837 tested, we found that Jadx selects the wrong implementation when a class with the same name
1838 is present. For example, in `classes2.dex` and `classes3.dex`. We report it with a “o” because
1839 warnings are issued.

1840 Shadowing Android SDK and hidden classes is possible in Jadx: there is only one implemen-
1841 tation of the class in the application, and Jadx does not have a list of the internal classes of
1842 Android: no warning is issued to the reverser that the displayed class is not the one used by
1843 Android.

1844 **4.3.2.2 Apktool**

1845 Apktool will store the disassembled classes in a folder that matches the `.dex` file that stores
1846 the bytecode. This means that when shadowing a class with two versions in two `.dex` files, the
1847 shadow implementations will be disassembled into two directories. No indication is displayed
1848 that a collision is possible. It is up to the reverser to have a chance to open the good one.

1849 Similarly to Jadx, using an Android SDK or hidden class will not be detected by the tool that
1850 will unpack the fake shadow version.

1851 **4.3.2.3 Androguard**

1852 Androguard has different usages, with different levels of analysis. The documentation highlights
1853 the analysis commands that compute three types of objects: an APK object, a list of DEX
1854 objects, and an Analysis object. The APK and the list of `.dex` files are a one-to-one represen-
1855 tation of the content of an application, and have the same issues that we discussed with Apktool:
1856 they provide the different versions of a shadow class contained in multiple `.dex` files.

1857 The Analysis object is used to compute a method call graph, and we found that this algorithm
1858 may choose the wrong version of a shadowed class when using the cross-references that are
1859 computed. This leads to an invalid call graph, as shown in Figure 14 b): the two methods
1860 `doSomething()` are represented in the graph, but the one linked to `main()` on the graph is the one
1861 calling the method `good()` when in fact the method `bad()` is called when running the application.

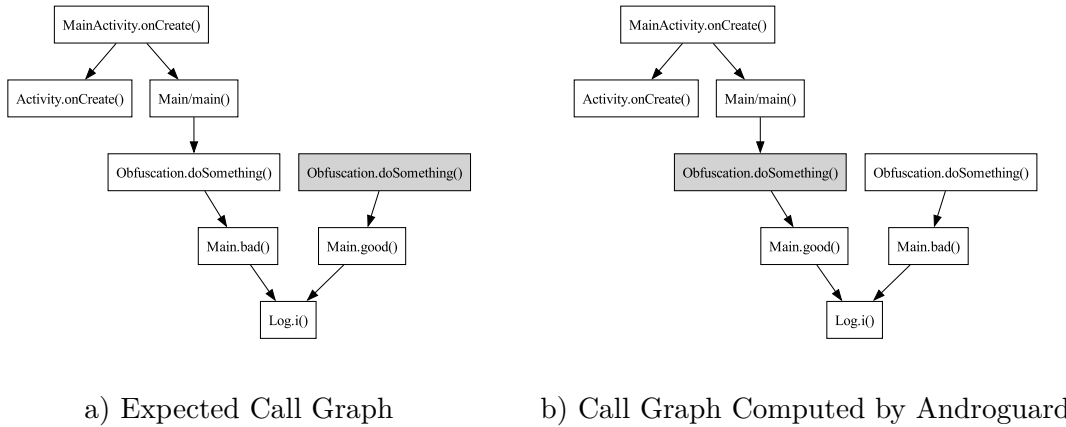


Figure 14: Call Graphs of an application calling `Main.bad()` from a shadowed `Obfuscation` class

1862 Androguard has a method `.is_external()` to detect if the implementation of a class is not
 1863 provided inside the application and a method `.is_android_api()` to detect if the class is part
 1864 of the Android API. Regrettably, the documentation of `.is_android_api()` explains that the
 1865 method is still experimental and just checks a few package names. This means that although
 1866 those methods are useful, the only indication of the use of an Android SDK or hidden classes
 1867 is the fact that the class is not in the APK file. Because of that, like for Apktool and Jadx,
 1868 Androguard has no way to warn the reverser that the shadow of an Android SDK or hidden
 1869 classes is not the class used when running the application.

1870 4.3.2.4 Flowdroid

1871 We found that when selecting the classes implementation in a multi-dex APK, Soot uses an
 1872 algorithm close to what ART is performing: Soot sorts the `.dex` bytecode file with a specified
 1873 `prioritizer` (a comparison function that defines an order for `.dex` files) and selects the first
 1874 implementation found when iterating over the sorted files. Unfortunately, the `prioritizer` used
 1875 by Soot is not exactly the same as the one used by the ART. The Soot `prioritizer` will give
 1876 priority to `classes.dex` and then give priority to files whose name starts with `classes` over other
 1877 files, and finally will use alphabetical order. This order is good enough for application with a
 1878 small number of `.dex` files generated by Android Studio, but because it uses the alphabetical
 1879 order and does not check the exact format used by Android, a malicious developer could hide the
 1880 implementation of a class in `classes2.dex` by putting a false implementation in `classes0.dex`,
 1881 `classes1.dex` or `classes12.dex`. Because Flowdroid is based on Soot, it inherits this issue from
 1882 it.

1883 In addition to self-shadowing, Flowdroid is sensitive to the use of platform classes, as it needs
1884 the bytecode of those classes to be able to track data flows. Flowdroid does have a record
1885 of SDK classes, and gives priority to the actual SDK classes over the classes implemented in
1886 the application, thus defeating SDK shadow attacks. Unfortunately, Flowdroid does not have
1887 a record of all platform classes, meaning that using hidden classes breaks the flow tracking.
1888 Solving this issue would require finding the bytecode of all the platform classes of the Android
1889 version targeted, and, as we said previously, it requires extracting this information from the
1890 emulator or phone.

1891 ∴

1892 We have seen that tools can be impacted by shadow attacks. In the next section, we will
1893 investigate whether these attacks are used in the wild.

1894 **4.4 Shadow Attacks in the Wild**

1895 In this section, we evaluate in the wild if applications that can be found in the Play Store or
1896 other markets use one of the shadow techniques. Our goal is to explore the usage of shadow
1897 techniques in real applications. Because we modelled the behaviour of a recent version of
1898 Android (SDK 34), we decided not to use our dataset from Chapter 3. The applications in
1899 the RASTA dataset span over more than 10 years, and we cannot guarantee that shadow
1900 attacks behaved the same during those 10 years. At the very least, self-shadowing would not be
1901 possible before the introduction of multi-dex in 2014 – about a fourth of the applications in the
1902 RASTA dataset. Instead, we sampled another dataset of recent applications. This way, we can
1903 also include malicious applications (in case such techniques would be used to hide malicious
1904 code), so we selected 50 000 applications randomly from AndroZoo [3] that appeared in 2023.
1905 Malicious applications are spotted in our dataset by using a threshold of 3 over the number of
1906 VirusTotal engines reporting an application as malware. This number is provided by AndroZoo
1907 for scans performed between January 2023 and January 2024, depending on the application. A
1908 few applications over the total could not be retrieved or parsed, leading to a final dataset of 49
1909 975 applications. We automatically disassembled the applications to obtain the list of included
1910 classes. Then, we check if any shadow attack occurs in the APK itself or with platform classes
1911 of SDK 34.

1912 **4.4.1 Results**

1913 [TODO: cl-shadow]

1914 We report in the upper part of Table 9 the statistics about the whole dataset and the three
1915 shadow attacks: “self” when a class shadows another one in the APK, “SDK” when a class of
1916 the SDK shadows one of the APK, and “Hidden” when a hidden class of Android shadows one
1917 of the APK. We observe that, on average, a few classes are shadowed by another class. Note

	Number of apps			Shadow classes	Average Median	Target SDK	Min SDK	Identical Code
	%	% malware						
For all applications of the dataset								
Self	49 975	100.0%	0.53%	2.1	0	32.1	21.7	74.8%
Sdk	49 975	100.0%	0.53%	6.5	0	32.1	21.7	8.04%
Hidden	49 975	100.0%	0.53%	0.5	0	32.1	21.7	17.42%
Total	49 975	100.0%	0.53%	9	0	32.1	21.7	23.76%
For applications with at least 1 shadow case								
Self	234	0.47%	5.98%	438.1	18	31.4	22.4	74.8%
Sdk	11 755	23.52%	0.38%	27.6	5	32.4	22	8.04%
Hidden	1556	3.11%	0.71%	16.1	1	32.1	22.2	17.42%
Total	12 301	24.61%	0.42%	36.7	6	32.4	22	23.76%

Table 9: Shadow classes compared to SDK 34 for a dataset of 49 975 applications

1918 that the median value is 0, meaning that few apps shadow a lot of classes, but the majority
1919 of apps do not shadow anything. The number of applications shadowing a hidden API is low,
1920 which is an expected result as these classes should not be known by the developer. We observe a
1921 consequent number of applications, 23.52%, that perform SDK shadowing. It can be explained
1922 by the fact that some classes that newly appear are embedded in the APK for end users that
1923 have old versions of Android: it is suggested by the average value of Min SDK which is 21.7
1924 for the whole dataset: on average, an application can be run inside a smartphone with API 21,
1925 which would require to embed all new classes from 22 to 34. This hypothesis about missing
1926 classes is further investigated later in this section.

1927 In the bottom part of Table 9, we give the same statistics, but we excluded applications that
1928 do not perform any shadowing. For those pairs of shadow classes, we disassembled them using
1929 Apktool to perform a comparison using instructions represented in the Smali language. For self-
1930 shadow, we compare the pair. For the shadowing of the SDK or Hidden class, we compare the
1931 code found in the APK with implementations found in the emulator and `android.jar` of SDK
1932 32, 33, and 34.

1933 *Self-shadowing* We observe a low number of applications doing self-shadow attacks. For each
1934 class that is shadowed, we compared its bytecode with the shadowed one (we compared the
1935 Smali instructions generated by Apktool for each method). We observe that 74.8% are identical,
1936 which suggests that the compilation process embeds the same class multiple times but makes
1937 variations in headers or metadata values. We investigate later in Section 4.4.2 the case of
1938 malicious applications.

1939 *SDK shadowing* For the shadowing of SDK classes, we observe a low ratio of identical classes.
1940 This result could lead to the wrong conclusion that developers embed malicious versions of the

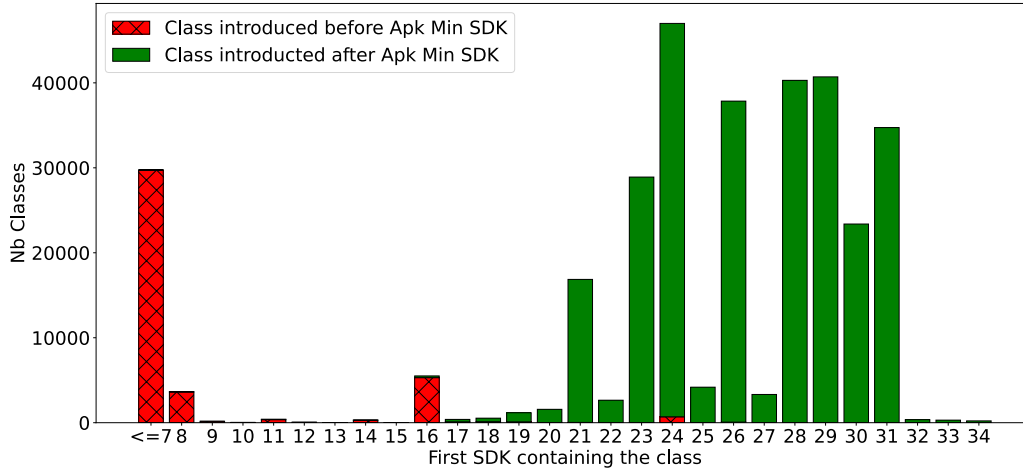


Figure 15: Redefined SDK classes, sorted by the first SDK they appeared in

1941 SDK classes, but our manual investigation shows that the difference is slight and probably due
 1942 to compiler optimisation. To go further in the investigation, in Figure 15, we represent these
 1943 redefined classes with the following rules:

- 1944 • The class is classified on the X abscissa in the figure according to the SDK it first appeared in.
- 1945 • The class is counted as “green” (solid) if it first appeared in the SDK **after** the APK min
- 1946 SDK (retro compatibility purpose).
- 1947 • The class is counted as “red” (hatched) if it first appeared in the SDK **before** the APK min
- 1948 SDK (which is useless for the application as the SDK version is always available).

1949 We observe that the majority of classes are legitimate retro-compatibility additions of classes,
 1950 especially after SDK 21 (which is the average min SDK, cf. Table 9). Abnormal cases are
 1951 observed for classes that appeared in API versions 7 and before, 8, and 16. Table 10 reports
 1952 the top ten classes that shadow the SDK for the three mentioned versions. For SDK before 7,
 1953 it mainly concerns HTTP classes: for example, the class `HttpParams` is an interface, containing
 1954 limited bytecode that mostly matches the class already present on the emulator (98.03% of
 1955 shadowed classes are identical). `HttpConnectionParams`, on the other hand, differs from the
 1956 platform class, and we observe only 4.99% of identical classes. Manual inspection of some
 1957 applications revealed that the two main reasons are:

- 1958 • Instead of checking if the method’s attributes are null inline, like Android does, applications
- 1959 use the method `org.apache.http.util.Args.notNull()`. According to comments in the source
- 1960 code of Android¹, the class was forked in 2007 from the Apache ‘httpcomponents’ project.
- 1961 Looking at the history of the project, the use of `Args.notNull()` was introduced in 2012².

- 1962 This shows that applications are embedding code from more recent versions of this library
 1963 without realising their version will not be the one used.
 1964 • Very small changes that we found can be attributed to the compilation process (e.g. swapping
 1965 registers: `v0` is used instead of `v1` and `v1` instead of `v0`), but even if we consider them different,
 1966 they are very similar.

Class	Occurrences	Identical ratio
redefined for SDK ≤ 7		
Lorg/apache/http/params/HttpParams;	1318	98.03%
Lorg/apache/http/params/HttpConnectionParams;	1202	4.99%
Lorg/apache/http/conn/ConnectTimeoutException;	1200	35.0%
Lorg/apache/http/params/CoreConnectionPNames;	1190	99.92%
Lorg/xmlpull/v1/XmlPullParser;	1111	52.57%
Lorg/apache/http/conn/scheme/SocketFactory;	1074	87.52%
Lorg/apache/http/conn/scheme/HostNameResolver;	1072	87.59%
Lorg/apache/http/conn/scheme/LayeredSocketFactory;	963	89.41%
Lorg/json/JSONException;	945	0.0%
Lorg/apache/http/conn/ssl/X509HostnameVerifier;	886	0.79%
redefined for SDK = 8		
Ljavax/xml/namespace/QName;	297	0.0%
Ljavax/xml/namespace/NamespaceContext;	226	98.23%
Landroid/net/http/SslError;	221	31.67%
Lorg/w3c/dom/UserDataHandler;	82	92.68%
Ljavax/xml/transform/TransformerConfigurationException;	73	69.86%
Ljavax/xml/transform/TransformerException;	73	0.0%
Lorg/w3c/dom/ls/LSEException;	61	63.93%
Lorg/w3c/dom/TypeInfo;	54	88.89%
Lorg/w3c/dom/DOMConfiguration;	54	46.3%
Ljavax/xml/transform/TransformerFactoryConfigurationError;	52	0.0%
redefined for SDK = 16		
Landroid/annotation/SuppressLint;	2634	98.48%
Landroid/annotation/TargetApi;	2634	98.48%
Landroid/media/MediaCodec\$CryptoException;	11	18.18%
Landroid/media/MediaCryptoException;	10	20.0%
Landroid/view/accessibility/AccessibilityNodeProvider;	9	0.0%
Landroid/view/ActionProvider\$VisibilityListener;	8	12.5%
Landroid/app/Notification\$BigTextStyle;	7	0.0%
Landroid/app/Notification\$Style;	7	0.0%
Landroid/util/LongSparseArray;	7	0.0%
Landroid/media/MediaPlayer\$TrackInfo;	7	0.0%

Table 10: Shadow classes compared to SDK 34 for a dataset of 49 975 applications

- 1967 1. <https://cs.android.com/android/platform/superproject/main/+main:frameworks/base/core/java/org/apache/http/params/HttpConnectionParams.java;drc=3bdd327f8532a79b83f575cc62e8eb09a1f93f3d?>
 1968 2. <https://github.com/apache/httpcomponents-core/commit/9104a92ea79e338d876b1b60f5cd2b243ba7069f?>
 1969

1970 The remaining 4.99% of classes that are identical to the Android version are classes where the
1971 body of the methods is replaced by stubs that throw `RuntimeException("Stub!")`. This code
1972 corresponds to what we found in `android.jar`, but not the code we found in the emulator, which
1973 is surprising. Nevertheless, we decided to count them as identical, because `android.jar` is the
1974 official jar file for developers, and stubs are replaced in the emulator: it is intended by Google
1975 developers.

1976 Other results of Table 10 can be similarly discussed: either they are identical with a high ratio, or
1977 they are different because of small variations. When substantial differences appear, it is mainly
1978 because different versions of the same library have been used or an SDK class is embedded for
1979 retro-compatibility.

1980 *Hidden shadowing* For applications redefining hidden classes, on average, 16.1 classes are
1981 redefined (cf bottom part of Table 9). The top 3 packages whose code actually differs from the
1982 ones found in Android are `java.util.stream`, `org.ccil.cowan.tagsoup` and `org.json`:

- 1983 • `stream`: when looking in more detail, we found that `java.util.stream` was only redefined
1984 by 6 applications, but the large number of classes redefined artificially puts the package at
1985 the top of the list. It is explained by the fact that developers have included this library,
1986 containing a lot of classes that collide with Android.
- 1987 • `tagsoup`: `TagSoup` is a library for parsing HTML. Developers do not know that it is part of
1988 Android as hidden classes.
- 1989 • `json`: there is only one hidden class in `org.json`, redefined by 821 applications: `JSONObject$1`.
1990 `org.json` is a package in Android SDK, not a hidden one. However, `JSONObject$1` is an
1991 anonymous class not provided by `android.jar` because its class `JSONObject` is an empty
1992 stub, and thus, does not use `JSONObject$1`. Thus, this class falls in the category of hidden
1993 platform classes.

1994 All these hidden shadow classes are libraries included by the developers who probably did not
1995 know that they were already embedded in Android.

1996 4.4.2 Shadowing in Malware Applications

1997 The last column of Table 9 shows the proportion of applications considered as malware because
1998 we arbitrarily fixed a threshold of 3 positive detections from VirusTotal reports. For the whole
1999 dataset, we have 0.53% of applications considered as malware. We can see that an application
2000 that uses self-shadowing is 10 times more likely to be malware, when the proportion of malware
2001 among application shadowing platform classes is the same as in the rest of the dataset. Thus, we
2002 manually reversed self-shadowing malware and found that the self-shadowing does not look to
2003 be voluntary. The colliding classes are often the same implementation, occasionally with minor
2004 differences, like different versions of a library. Additionally, we noticed multiple times internal

```

1 public class Reflection {
2     private static final int ERROR_SET_APPLICATION_FAILED = -20;
3     private static final String TAG = "Reflection";
4     // ...
5
6     static {
7         try {
8             Method declaredMethod = Class.class.getDeclaredMethod("forName",
String.class);
9             Method declaredMethod2 =
Class.class.getDeclaredMethod("getDeclaredMethod", String.class,
Class[].class);
10            Class cls = (Class) declaredMethod.invoke(null,
"dalvik.system.VMRuntime");
11            Method method = (Method) declaredMethod2.invoke(cls,
"getRuntime", null);
12            setHiddenApiExemptions = (Method) declaredMethod2.invoke(cls,
"setHiddenApiExemptions", new Class[]{String[].class});
13            sVmRuntime = method.invoke(null, new Object[0]);
14        } catch (Throwable th) { Log.e(TAG, "reflect bootstrap failed:",
th); }
15        System.loadLibrary("free-reflection");
16        // ...
17    }
18    // ...
19 }

```

Listing 5: Implementation of Reflection found in classes11.dex (shadows Listing 6)

2005 classes from `com.google.android.gms.ads` colliding with each other, but we believe that it is
2006 due to bad processing during the compilation of the application.

2007 The most notable case we found was an application that still exists on the Google
2008 Play Store with the same package name¹. This application contains a self-shadow class
2009 `me.weishu.reflection.Reflection` that can be found in Github, in the repository `tiann/`
2010 `FreeReflection`². This class is used to disable Android restrictions on hidden API. At first
2011 glance, we believed the shadowing to be done voluntarily for obfuscation purposes. The shadow
2012 class that would be seen by a reverser is given in Listing 5: it contains some Java bytecode
2013 performing reflection and loading a native library named “free-reflection” (the associated `.so`
2014 is missing). The shadowed class that is really executed is summarised in Listing 6. It contains

2015 1. SHA256: C46A65EA1A797119CCC03C579B61C94FE8161308A3B6A8F55718D6ADAD112546

2016 2. <https://github.com/tiann/FreeReflection>

```

1 public class Reflection {
2     private static final String DEX = "ZGV4CjAzNQCl4EprGS2pXI/
v30wlBrlfRnX5rmkKVdN0CwAAcA ... AoAAA==";
3     private static final String TAG = "Reflection";
4
5     private static native int unsealNative(int i);
6
7     public static int unseal(Context context) {
8         return (Build.VERSION.SDK_INT < 28 || BootstrapClass.exemptAll()
|| unsealByDexFile(context)) ? 0 : -1;
9     }
10
11     private static boolean unsealByDexFile(Context context) {
12         // Decode DEX from base64 and load it as bytecode.
13         // ...
14     }
15     // ...
16 }

```

Listing 6: Implementation of Reflection executed by ART (shadowed by Listing 5)

2017 a more obfuscated code: a DEX field storing base64 encoded DEX bytecode that is later used to
2018 load some new code. When looking at this new code stored in the field, we found that it does
2019 almost the same thing as the code in the shadow class. Thus, we believe that the developer has
2020 upgraded their obfuscation techniques, replacing a native library with inline base64 encoded
2021 bytecode. The shadow attack could be unintentional, but it strengthens the masking of the new
2022 implementation.

2023 ∴

2024 As a conclusion, we observed that:

- 2025 • SDK shadowing is performed by 23.52% of applications, but is unintentional: these classes
2026 are embedded for retro-compatibility purposes or because the developer added a library
2027 already present in Android.
- 2028 • Hidden shadowing rarely occurs and is mainly due to the usage of libraries that Android
2029 already contains.
- 2030 • Malware performs more self-shadowing than goodware applications, and we found a sample
2031 where self-shadowing would clearly mislead the reverser.

2032 4.5 Discussion

2033 TODO: small intro

2034 4.5.1 Countermeasures

2035 Countermeasures against shadow attacks depend on each tool and its objectives. The first
2036 important recommendation is to implement the class selection algorithm according to the
2037 algorithm described in Listing 2. It should solve any case of self-shadowing, except for tools
2038 like Apktool, that do not have to select a class for computing the result, but show the whole
2039 application’s content. For those tools, a clear warning should be added, pointing out that
2040 multiple implementations have been found and displaying the one that will be used at runtime.

2041 Countermeasures against SDK shadow and Hidden shadow attacks are more complex to handle:
2042 they require the list of platform classes on the target smartphone and, in some cases, their
2043 implementation. The list of SDK classes can be extracted easily from android.jar, but hidden
2044 classes need to be obtained by other means. They could be listed directly from the AOSP
2045 tree of the Android source code, obtained from Android documentation, or extracted from the
2046 phone itself. The first approach requires statically analysing the source code, which can be
2047 difficult to achieve as several programming languages are used, and the code base is large and
2048 fragmented. Ideally, the documentation would be the best solution, but as discussed earlier in the
2049 chapter, it can lack some classes. For this solution to be viable, Google would need to keep the
2050 documentation closer to the released version of Android than it currently is. Also, smartphone
2051 manufacturers might add additional classes that would not appear in Google documentation.
2052 In fact, neither the documentation nor the source code approach can be generalised for all
2053 possible versions of Android, as the exact list will depend on the exact targeted device, possibly
2054 modified by the manufacturer. Thus, to counter Shadow attacks, the static analysis tools that
2055 we evaluated need to embed multiple lists of platform classes, one for each Android version.
2056 Then, the best heuristic would be to use the list of platform classes that is closest to the target
2057 SDK of the analysed application.

2058 Some tools, like Flowdroid, would require additional countermeasures: to compute the exact
2059 flow of data, Flowdroid also needs to analyse the code of platform classes. For the SDK classes,
2060 Flowdroid has already analysed them, but the hidden classes have not. In addition to the data
2061 flow in hidden classes, Flowdroid needs a list of data sources and sinks from those classes. Other
2062 analysis tools may require additional data from platform classes, which may be too difficult to
2063 obtain.

2064 We believe that analysis tools can handle shadow attacks to some degree. The implementation
2065 of the solution will differ depending on the nature of the tool and may not always require the
2066 same implementation effort.

2067 4.5.2 Relation with Obfuscation Techniques

2068 As described in the state of the art, reverse engineers face other techniques of obfuscation, such
2069 as packers or native code. These techniques rely on custom class loaders that load new parts of

2070 the application from ciphered assets or from the network. The reverse engineers have to study
2071 the application dynamically to recover new classes and eventually go back to a static phase to
2072 understand the behaviour of the application. In this section, we compare shadow attacks with
2073 these techniques and discuss how they interact with them.

2074 Advanced obfuscation techniques relying on packers have a higher impact on the difficulty
2075 of performing a static analysis compared to shadow attacks. Most of the time, the reverse
2076 engineer cannot deobfuscate the application without performing a dynamic analysis. For these
2077 reasons, approaches have been designed to assist the capture of the bytecode that is loaded
2078 dynamically, after the precise time where the deobfuscation methods have been executed [68,
2079 70, 74]. On the contrary, a shadow attack can be easily defeated by implementing our algorithm
2080 in the static analysis tool, as discussed earlier in Section 4.5.1. Nevertheless, shadow attacks are
2081 stealthier than packers or native code. Packers can be easily spotted by artefacts left behind
2082 in the application or by detecting classes implementing a custom class loading mechanism. On
2083 the contrary, an extra class implementing a shadow attack, that would not be executed, could
2084 contain voluntarily little code, compared to the executed class of Android. Such an attack would
2085 be more discreet than a packer that adds in the application a lot of possibly native code.

2086 Combining regular obfuscation techniques with shadow attacks can be achieved in two ways.

2087 First, the attacker could hide the code of a packer or a native call by using a shadow attack. For
2088 example, by colliding a class of the SDK, a control flow analysis could be wrongly computed,
2089 leading to considering that part of the code to be dead, which would mislead the reverse engineer
2090 about the use of this part that contains a packer. At runtime, this code would be triggered,
2091 unpacking new code.

2092 Second, the attacker could use a packer to unpack code at runtime in a first phase. The
2093 reverse engineer would have to perform a dynamic analysis, for example, using a tool such as
2094 Dexhunter [74], to recover new DEX files that are loaded by a custom class loader. Then, the
2095 reverse engineer would go back to a new static analysis and could have the problem of solving
2096 shadow attacks, for example, if a class is defined multiple times in the loaded DEX files.

2097 Because the interactions between shadow attacks and other obfuscation techniques often rely on
2098 a loading mechanism implemented by the developer, investigating these cases requires analysing
2099 the Java bytecode that is handling the loading. This problem is left as future work.

2100 **4.5.3 Limitations**

2101 During the analysis of the ART internals, we made the hypothesis that its different operating
2102 modes are equivalent: we analysed the loading process for classes stored as non-optimised `.dex`
2103 format, and not for the pre-compiled `.oat`. It is a reasonable hypothesis to suppose that the two
2104 implementations have been produced from the same algorithm using two compilation workflows.

2105 Similarly, we assumed that the platform classes stored in `boot.art` are the same as the ones in
2106 `BOOTCLASSPATH`. We confirm empirically our hypothesis on an Android Emulator, but we may
2107 have missed some edge cases.

2108 The comparison of Smali code can lead to underestimated values, for example, if the compilation
2109 process performs minor modifications such as instruction reordering. The ratios reported in
2110 this study for the comparison of code are thus a lower bound and would be higher with a
2111 more precise comparison. In addition, platform classes are stored differently in older versions of
2112 Android and cannot be easily retrieved. For this reason, we did not compare the classes found
2113 in applications to their versions older than SDK 32 to avoid producing unreliable statistics for
2114 those versions.

2115 4.5.4 Future Works

2116 As we just said, our Smali-based comparison of class implementation is quite naive and could
2117 use more work. It could be insightful to be able to detect exactly when two classes are from
2118 the same source file, or which version of a library a class belong to. More importantly, a better
2119 comparison technique would allow us to detect cases where the shadowed library has actual
2120 malicious bytecode added that we could have missed manually.

2121 Additionally, the question of dynamic class loaders, used manually by the application developer,
2122 is interesting. This is reaching the limits of static analysis; those cases involve dynamically
2123 loading bytecode, and in many cases, the classes loaded by those class loaders are not even
2124 available for analysis. However, even with dynamic analysis, the behaviour of class loaders can
2125 still be an issue, especially when the analysis is performed by alternating static and dynamic
2126 analysis, as is often the case when manually reversing an application. To handle those cases, it
2127 could be interesting to develop a method to model any arbitrary class loader, either by analysing
2128 its bytecode or by interacting with an instance of the class loader dynamically.

2129 In September 2024 (just after we finished this work), Android 15 introduced support for the
2130 new version 41 of the DEX format. We can expect this version of DEX to become the norm in a
2131 few years. The most notable change in version 41 is the new container format: instead of storing
2132 the bytecode in separate DEX files, the different files can now be concatenated into one unique
2133 file. There is also some permeability between the concatenated files: some structures stored in
2134 one file can be used by the next concatenated files. This significant change in the bytecode
2135 storage is similar to the introduction of the multi-dex format. Considering that self-shadowing
2136 is only possible because of the multi-dex format, we can expect this change to have the potential
2137 to introduce new, similar issues. Thus, we believe that the implementation details of this new
2138 version should be studied and modelled properly to avoid introducing new issues when updating
2139 analysis tools to support it. Just by reading the specification¹, we believe that self-shadowing

2140 1. <https://source.android.com/docs/core/runtime/dex-format#container>

2141 between concatenated DEX files is possible, unless additional checks are enforced by the ART
 2142 when loading the file.
 2143

4.6 Conclusion

TODO: Ca serait bien de faire un PR ou deux a Jadx/Androguard/Soot quand même

This chapter has presented three shadow attacks that allow malware developers to fool static analysis tools when reversing an Android application. By including multiple classes with the same name or by using the same name as a class of the Android SDK, the developer can mislead a reverse engineer or impact the result of a flow analysis, such as those of Androguard or Flowdroid.

We explored whether such shadow attacks are present in a dataset of 49 975 applications. We found that on average, 23.52% of applications are shadowing the SDK, mainly for retro-compatibility purposes and library embedding. More suspiciously, 3.11% of applications are shadowing a hidden class, which could lead to unexpected execution as these classes can appear/disappear with the evolution of Android internals. Investigations for applications that defined classes multiple times suggest that the compilation process or the inclusion of different versions of the same library is the main explanation. Finally, when investigating malware samples, we found a specific sample containing a shadow attack that would hide a part of the critical code from a reverse engineer studying the application.

2161
 2162

Pb2: *What is the default Android class loading algorithm, and does it impact static analysis?*

2163
 2164
 2165
 2166

Listing 2 model the class loading algorithm: platform classes have priority over classes stored in `classes.dex` which have priority over `classes<n>.dex` (where $n \in \llbracket 2, +\infty \llbracket$ and $\forall i \in \llbracket 2, n \llbracket, \exists$ `classes<i>.dex`) which has priority over `classes<n+1>.dex`.

2167
 2168
 2169
 2170

Failing to implement this model (*i.e.*, by ignoring some platform classes or by sorting the `classes<n>.dex` alphabetically instead of numerically) can cause static analysis tools to compute an incorrect representation of the analysed application.

2171

2172

THE APPLICATION OF THESEUS: AFTER ADDING RUNTIME DATA, IT IS STILL YOUR APPLICATION

2173

2174

2175

Despite everything, it's still you.

2176

— Undertale, Toby Fox

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

Some applications use dynamic code loading and reflection calls that prevent static analysis tools from analysing the complete application. This can be detected with dynamic analysis; however, the collected data is not enough to analyse the application further: most tools do not have a way to process this additional data. In this chapter, we propose to use dynamic analysis to collect information related to dynamic code loading and reflection, and to encode this information in the bytecode of the application to allow further analysis. We compared the results of analysis on applications before and after the transformation, using tools like Flowdroid or Androguard, and found that the additional information is indeed processed by the static analysis tools. We also compared the finishing rate of the tools, using the same experiment as in Chapter 3, and found that the finishing rate is generally only slightly negatively impacted by the transformation.

2192 5.1 Introduction

2193 In the previous chapter, we studied the static impact of class loaders. However, as we focused
2194 on the default behaviour of Android, we ignored the main use of class loaders for developers:
2195 dynamic code loading. In this chapter, we address this issue, as well as the issue of reflection
2196 that often accompanies dynamic code loading. Dynamic code loading is the practice of loading
2197 at runtime bytecode that was not already part of the original bytecode of the application. This
2198 bytecode can be stored as assets of the application, downloaded from a remote server, or even
2199 generated algorithmically by the application. This is a problem for analysis: when the bytecode
2200 is not already visible in the application, it cannot be analysed statically. Meanwhile, reflection
2201 is the action of using code to manipulate objects representing structures of the code itself, like
2202 classes or methods. The main issue for analysis occurs when it is used to call methods. A static
2203 analysis will show calls to `Method.invoke()`, but not the actual method invoked.

2204 In both cases, static analysis falls short, as the information to analyse may be generated just
2205 in time for its use. For such cases, dynamic analysis is a more appropriate approach. It can be
2206 used to collect the missing information while the application is running. However, having this
2207 information does not mean that the application can now be analysed in its entirety. Generic
2208 analysis tools rarely have an easy way to read additional information about an application
2209 before analysing, and when they do, it is not standard. The usual approach for hybrid analysis
2210 (analyses that mix static and dynamic analysis) is to select one specific static tool and modify
2211 its code to take into account the additional data collected by dynamic analysis. This limits the
2212 reverse engineer to a few tools that they took the time to study and modify for the task. In this
2213 chapter, we propose to modify the code of the application to add the information needed for
2214 analysis in a format that any analysis tool can use. This way, the analyst is no longer limited
2215 in their choice of tool and can focus on the actual analysis of the application.

2216 We structured this chapter as follows: We first present an overview of our method in Section 5.2.
2217 We then present the transformations we apply to the application in Section 5.3 and the dynamic
2218 analysis we perform in Section 5.4. In Section 5.5 compare the results of different tools on the
2219 initial application versus the modified application. To complete this chapter, in Section 5.6 we
2220 discuss the limits of our solution, as well as directions for future work. Finally, we conclude in
2221 Section 5.7.

2222 5.2 Overview

2223 Our objective is to make available some dynamic information to any analysis tool able to analyse
2224 an Android APK. To do so, we elected to follow the same approach as a few contributions we
2225 presented in Chapter 2, such as DroidRA [32], and use instrumentation. As a reminder, DroidRA
2226 is a tool that uses COAL to compute reflection data statically, then instruments the application
2227 to directly call the methods. Contrary to DroidRA, we chose to use dynamic analysis. This




Figure 16: Process to add runtime information to an APK

2228 allows us to collect information that is simply not available statically (*e.g.*, a string sent from a
 2229 remote command and control server). The tradeoff here is the lack of exhaustiveness: dynamic
 2230 analysis is known to have code coverage issues.

2231 Figure 16 summarises our process. We first take an application that we analyse dynamically.
 2232 To improve code coverage, either a reverse engineer or an automated runner will interact with
 2233 the application. During this analysis, we use Frida to capture dynamic information like the
 2234 names of the methods called using reflection and bytecode loaded at runtime. This analysis is
 2235 described in Section 5.4.

2236 The data collected by this analysis is then combined with the application, transforming the
 2237 application into another one that can then be analysed further. We present the details of this
 2238 transformation in Section 5.3. Since the transformation drives the data we need to collect, we
 2239 have decided to place this section first in this chapter.

2240 5.3 Code Transformation

2241 In this section, we will see how we can transform the application code to make dynamic code
 2242 loading and reflective calls more analysable by static analysis tools.

2243 5.3.1 Transforming Reflection

2244 In Android, reflection allows applications to instantiate a class or call a method without having
 2245 this class or method appear in the bytecode. Instead, the bytecode uses the generic classes `Class`,
 2246 `Method` and `Constructor`, which represent any existing class, method or constructor. Reflection
 2247 often starts by retrieving the `Class` object representing the class to use. This class is usually
 2248 retrieved using a `ClassLoader` object (though there are other ways to get it). Once the class is
 2249 retrieved, it can be instantiated using the deprecated method `Class.newInstance()`, as shown
 2250 in Listing 7, or a specific method can be retrieved. The current approach to instantiate a class
 2251 is to retrieve the specific `Constructor` object, then call `Constructor.newInstance(..)` like in
 2252 Listing 8. Similarly, to call a method, the `Method` object must be retrieved, then called using
 2253 `Method.invoke(..)`, as shown in Listing 9.

```

1 ClassLoader cl = MainActivity.class.getClassLoader();
2 Class clz = cl.loadClass("com.example.Reflectee");
3 Object obj = clz.newInstance();

```

Listing 7: Instantiating a class using `Class.newInstance()`

```

1 Constructor cst = clz.getDeclaredConstructor(String.class);
2 Object obj = cst.newInstance("Hello Void");

```

Listing 8: Instantiating a class using `Constructor.newInstance(..)`

```

1 Method mth = clz.getMethod("myMethod", String.class);
2 Object[] args = {(Object)"an argument"};
3 String retData = (String) mth.invoke(obj, args);

```

Listing 9: Calling a method using reflection

2254 Although the process seems to differ between class instantiation and method call from the Java
 2255 standpoint, the runtime operations are very similar. When instantiating an object with `Object`
 2256 `obj = cst.newInstance("Hello Void")`, the constructor method `<init>(Ljava/lang/String;)V`,
 2257 represented by the `Constructor` `cst`, is called on the object `obj`. Thus, even for instantiation, a
 2258 method is called at some point.

2259 One of the main reasons to use reflection is to access classes that are neither platform classes
 2260 nor in the application bytecode, as is often the case when dealing with classes from dynamically
 2261 loaded bytecode. Indeed, if the ART were to encounter an instruction referencing a class that
 2262 cannot be loaded by the current class loader, it would crash the application.

2263 To allow static analysis tools to analyse an application that uses reflection, we want to replace
 2264 the reflection call with a bytecode chunk that actually calls the method and can be analysed
 2265 by any static analysis tool. In Section 5.3.2, we deal with the issue of dynamic code loading so
 2266 that the classes used are, in fact, present in the application.

2267 A notable issue is that a specific reflection call can call different methods. Listing 10 illustrates
 2268 a worst-case scenario where any method can be called at the same reflection call. In those
 2269 situations, we cannot guarantee that we know all the methods that can be called (*e.g.*, the
 2270 name of the method called could be retrieved from a remote server). In addition, the method
 2271 we propose in Section 5.4 is a best effort approach to collect reflection data: like any dynamic
 2272 analysis, it is limited by its code coverage.

2273 To handle those situations, instead of entirely removing the reflection call, we can modify the
 2274 application code to test if the `Method` (or `Constructor`) object matches any of the methods

```
1 Object myInvoke(Object obj, Method mth, Object[] args) throws .. {  
2     return mth.invoke(obj, args);  
3 }
```

Listing 10: A reflection call that can call any method

2275 observed dynamically, and if so, directly call the method. If the object does not match any
2276 expected method, the code can fall back to the original reflection call. DroidRA [32] has a
2277 similar solution, except that reflective calls are always evaluated, and the static equivalent
2278 follows just after, guarded behind an opaque predicate that is always false at runtime. Listing 11
2279 demonstrates this transformation for the code originally in Listing 9. Let’s suppose that we
2280 observed dynamically a call to a method `Reflectee.myMethod(String)` at line 3 when monitoring
2281 the execution of the code of Listing 9. In Listing 11, at line 25, the `Method` object `mth` is checked
2282 using a method we generated and injected in the application (defined at line 2 in the listing).
2283 This method checks if the method name (line 5), its parameters (lines 6-9), its return type
2284 (lines 10-11) and its declaring class (lines 13-14) match the expected method. If it is the case,
2285 the method is used directly (line 26) after casting the arguments and associated object into the
2286 types/classes we just checked. If the check line 25 does not pass, the original reflective call is
2287 made (line 28). If we were to expect other possible methods to be called in addition to `myMethod`,
2288 we would add `else if` blocks between lines 26 and 27, with other check methods reflecting each
2289 potential method call.

2290 The check of the `Method` value is done in a separate method injected inside the application
2291 to avoid cluttering the application too much. Because Java (and thus Android) uses polymor-
2292 phic methods, we cannot just check the method name and its class, but also the whole
2293 method signature. We chose to limit the transformation to the specific instruction that calls
2294 `Method.invoke(..)`. This drastically reduces the risks of breaking the application, but leads
2295 to a lot of type casting. Indeed, the reflection call uses the generic `Object` class, but actual
2296 methods usually use specific classes (*e.g.*, `String`, `Context`, `Reflectee`) or scalar types (*e.g.*, `int`,
2297 `long`, `boolean`). This means that the method parameters and object on which the method is
2298 called must be downcasted to their actual type before calling the method, then the returned
2299 value must be upcasted back to an `Object`. Scalar types especially require special attention.
2300 Java (and Android) distinguish between scalar types and classes, and they cannot be mixed:
2301 a scalar cannot be cast into an `Object`. However, each scalar type has an associated class that
2302 can be used when doing reflection. For example, the scalar type `int` is associated with the class
2303 `Integer`, the method `Integer.valueOf()` can convert an `int` scalar to an `Integer` object, and
2304 the method `Integer.intValue()` converts back an `Integer` object to an `int` scalar. Each time
2305 the method called by reflection uses scalars, the scalar-object conversion must be made before

```

1 class T {
2     static boolean check_is_reflectee_mymethod_e398(Method mth) {
3         Class<?>[] paramTys = mth.getParameterTypes();
4         return (
5             meth.getName().equals("myMethod") &&
6             paramTys.length == 1 &&
7             paramTys[0].descriptorString().equals(
8                 String.class.descriptorString()
9             ) &&
10            meth.getReturnType().descriptorString().equals(
11                String.class.descriptorString()
12            ) &&
13            meth.getDeclaringClass().descriptorString().equals(
14                Reflectee.class.descriptorString()
15            )
16        )
17    }
18 }
19 ...
20 ...
21
22 Method mth = clz.getMethod("myMethod", String.class);
23 Object[] args = {(Object)"an argument"}
24 Object objRet;
25 if (T.check_is_reflectee_mymethod_e398abf7d3ce6ede(mth)) {
26     objRet = (Object) ((Reflectee) obj).myMethod((String)args[0]);
27 } else {
28     objRet = mth.invoke(obj, args);
29 }
30 String retData = (String) objRet;

```

Listing 11: Listing 9 after the de-reflection transformation

2306 calling it. And finally, because the instruction following the reflection call expects an `Object`,
 2307 the return value of the method must be cast into an `Object`.

2308 This back and forth between types might confuse some analysis tools. This could be improved
 2309 in future works by analysing the code around the reflection call. For example, if the result of
 2310 the reflection call is immediately cast into the expected type (*e.g.*, in Listing 9, the result is cast
 2311 to a `String`), there should be no need to cast it to `Object` in between. Similarly, it is common to
 2312 have the method parameter arrays generated just before the reflection call and never be used
 2313 again (This is due to `Method.invoke(..)` being a varargs method: the array can be generated

2314 by the compiler at compile time). In those cases, the parameters could be used directly without
 2315 the detour inside an array.

2316 5.3.2 Transforming Code Loading (or Not)

2317
 2318 An application can dynamically import code from several formats like DEX, APK, JAR or OAT,
 2319 either stored in memory or in a file. Because it is an internal, platform-dependent format, we
 2320 elected to ignore the OAT format. Practically, JAR and APK files are zip files containing DEX
 2321 files. This means that we only need to find a way to integrate DEX files into the application.

2322 We saw in Chapter 4 the class loading model of Android. When doing dynamic code loading, an
 2323 application defines a new `ClassLoader` that handles the new bytecode, and starts accessing its
 2324 classes using reflection. We also saw in Chapter 4 that Android now uses the multi-dex format,
 2325 allowing it to handle any number of DEX files in one class loader. Therefore, the simpler way
 2326 to give access to the dynamically loaded code to static analysis tools is to add the dex files in
 2327 the application as additional multi-dex bytecode files. This should not impact the class loading
 2328 model as long as there is no class collision (we will explore this in Section 5.3.3) and as long as
 2329 the original application does not try to access inaccessible classes (we will develop this issue in
 2330 Section 5.6).
 2331

2332 In the end, we decided **not** to modify the original code that loads the bytecode. Most tools
 2333 already ignore dynamic code loading, and, with the dynamically loaded bytecode added using
 2334 the multi-dex format, they already have access to it. At runtime, although the bytecode is
 2335 already present in the application, the application will still dynamically load the code. This
 2336 ensures that the application keeps working as intended, even if the transformation we applied is
 2337 incomplete. Specifically, to call dynamically loaded code, an application needs to use reflection,
 2338 and we saw in Section 5.3.1 that we need to keep reflection calls, and in order to keep reflection
 2339 calls, we need the class loader created when loading bytecode.

2340 To summarise, we do not modify the existing bytecode. Instead, we add the intercepted
 2341 bytecode to the application as additional DEX files using the multi-dex format, as represented
 2342 in Figure 17.

2343 5.3.3 Class Collisions

2344 We saw in Chapter 4 that having several classes with the same name in the same application
 2345 can be problematic. In Section 5.3.2, we are adding new code. By doing so, we increase the
 2346 probability of having class collisions: The developer may have reused a helper class in both the
 2347 dynamically loaded bytecode and the application, or an obfuscation process may have renamed
 2348 classes without checking for intersection between the two sources of bytecode. When loaded
 2349 dynamically, the classes are in a different class loader, and the class resolution is resolved at
 2350 runtime, like we saw in Section 4.2. We decided to restrain our scope to the use of class loaders

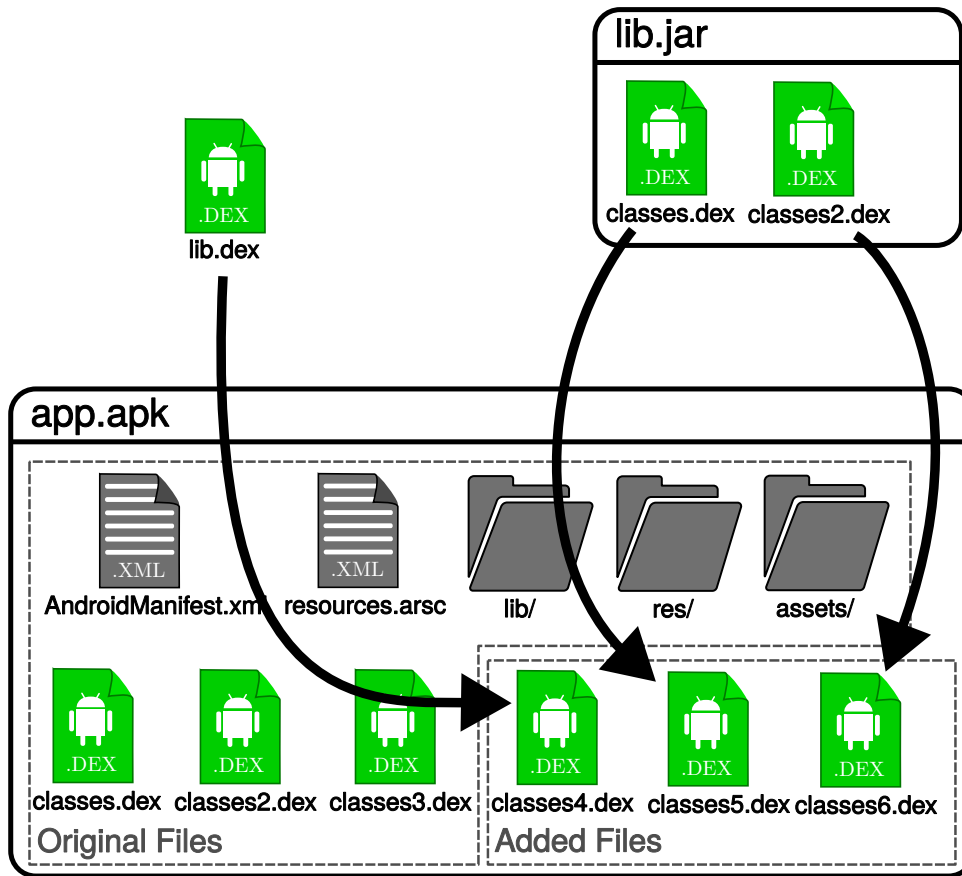


Figure 17: Inserting DEX files inside an APK

2351 from the Android SDK. In the absence of class collision, those class loaders behave seamlessly
 2352 and adding the classes to the application maintains the behaviour.

2353 When we detect a collision, we rename one of the colliding classes in order to be able to
 2354 differentiate between classes. To avoid breaking the application, we then need to rename all
 2355 references to this specific class and be careful not to modify references to the other class. To
 2356 do so, we regroup each class by the class loaders that define them. Then, for each colliding
 2357 class name and each class loader, we check the actual class used by the class loader. If the class
 2358 has been renamed, we rename all references to this class in the classes defined by this class
 2359 loader. To find the class used by a class loader, we reproduce the behaviour of the different
 2360 class loaders of the Android SDK. This is an important step: remember that the delegation
 2361 process can lead to situations where the class defined by a class loader is not the class that will
 2362 be loaded when querying the class loader. The pseudo-code in Listing 12 shows the three steps
 2363 of this algorithm:

Un exam-
 ple aiderait
 a compren-
 dre
 jm: j'en
 ai pas qui
 prennent
 pas 3 pages
 de listing

```
1 defined_classes = set()
2 redefined_classes = set()
3
4 # Rename the definition of redefined classes
5 for cl in class_loaders:
6     for clz in defined_classes.intersection(cl.defined_classes):
7         cl.rename_definition(clz)
8         redefined_classes.add(clz)
9     defined_classes.update(cl.defined_classes)
10
11 # Rename reference of redefined classes
12 for cl in class_loaders:
13     for clz in redefined_classes:
14         defining_cl = cl.resolve_class(clz).class_loader
15         cl.rename_reference(clz, defining_cl.new_name(clz))
16
17 # Merge the class loader into a flat APK
18 new_apk = Apk()
19 for cl in class_loaders:
20     for dex in cl.get_dex():
21         new_apk.add_dex(dex)
```

Listing 12: Pseudo-code of the renaming algorithm

- 2364 • First, we detect collisions and rename class definitions to remove the collisions.
- 2365 • Then we rename the reference to the colliding classes to make sure the right classes are called.
- 2366 • Ultimately, we merge the modified DEX files of each class loader into one Android appli-
- 2367 cation.

2368 5.3.4 Implementation Details

2369 Our initial idea was to use Apktool, but in Chapter 3, we found that many errors raised by
2370 tools were due to trying to parse Smali incorrectly. Thus, we decided to avoid Apktool.

2371 Most of the contributions of the state of the art that perform instrumentation rely on Soot.
2372 Soot works on an intermediate representation, Jimple, that is easier to manipulate. However,
2373 Soot can be cumbersome to set up and use, and we initially wanted better control over the
2374 modified bytecode. In addition, although it might be due to the fact that they performed more
2375 complex analysis, tools based on Soot showed a trend of consuming a lot of memory and failing
2376 with unclear errors, supporting us in our idea of avoiding Soot. For these reasons, we decided
2377 to make our own instrumentation library from scratch.

2378 That library, Androscapel, requires being able to parse, modify and generate valid DEX files.
2379 It was not as difficult as one would expect, thanks to the clear documentation of the Dalvik

2380 format from Google¹. In addition, when we had doubts about the specification, we had the
2381 option to check the implementation used by Apktool², or the code used by Android to check
2382 the integrity of the DEX files³.

2383 We chose to use Rust to implement this library. It has both good performance and ergonomics.
2384 For instance, we could parallelise the parsing and generation of DEX files without much effort.
2385 Because we are not using a high-level intermediate language like Jimple (used by Soot), the
2386 management of the Dalvik registers in the methods has to be done manually (by the user of
2387 the library), the same way it has to be done when using Apktool. This poses a few challenges.

2388 A method declares a number of internal registers it will use (let's call this number n), and
2389 has access to an additional number of registers used to store the parameters (let's call this
2390 number p). Each register is referred to by a number from 0 to 65535. The internal registers are
2391 numbered from 0 to $n - 1$, and the parameter registers from n to $n + p - 1$. This means that
2392 when adding new registers to the method when instrumenting it (let's say we want to add k
2393 registers), the new registers will be numbered from n to $n + k - 1$, and the parameter registers
2394 will be renumbered from $\llbracket n, n + p \rrbracket$ to $\llbracket n + k, n + k + p \rrbracket$. In general, this is not an issue, but
2395 some instructions can only operate on some registers (*e.g.*, `array-length`, which stores the length
2396 of an array in a register, only works on registers numbered between 0 and 8 excluded). This
2397 means that adding registers to a method can be enough to break a method. We solved this by
2398 adding instructions that move the content of registers $\llbracket n + k, n + k + p \rrbracket$ to the registers $\llbracket n, n +$
2399 $p \rrbracket$, and keeping the original register numbers ($\llbracket n, n + p \rrbracket$) for the parameters in the rest of the
2400 body of the method.

2401 The next challenge arises when we need to use one of the new registers with an instruction that
2402 only accepts registers lower than $n + p$. In such cases, a lower register must be used, and its
2403 content will be temporarily saved in one of the new registers. This is not as easy as it seems:
2404 the Dalvik instructions differ depending on whether the register stores a reference or a scalar
2405 value, and Android does check that the register types match the instructions. The type of the
2406 register can be computed from the control flow graph of the method (we added the computation
2407 of such a graph, with the type of each register, as a feature in Androscarpel). An edge case
2408 that must not be overlooked is that each instruction inside a `try` block is branching to each of
2409 the `catch` blocks. This is a problem: it prevents us from restoring the registers to their original
2410 values before entering the `catch` blocks (or, if we restore the values at the beginning of the `catch`
2411 blocks and an exception is raised before the value is saved, the register will be overwritten by

2412 1. <https://source.android.com/docs/core/runtime/dex-format>

2413 2. <https://github.com/JesusFreke/smali>

2414 3. [https://cs.android.com/android/platform/superproject/main/+/main:art/libdexfile/dex/dex_file_verifier.
2415 cc;drc=11bd0da6cfa3fa40bc61deae0ad1e6ba230b0954](https://cs.android.com/android/platform/superproject/main/+/main:art/libdexfile/dex/dex_file_verifier.cc;drc=11bd0da6cfa3fa40bc61deae0ad1e6ba230b0954)

2416 an invalid value). This means that when modifying the content of a `try` block, the block must
2417 be split into several blocks to prevent impromptu branching.

2418 One thing we noticed when manually instrumenting applications with Apktool is that sometimes
2419 the repackaged applications cannot be installed or run due to some files being stored incorrectly
2420 in the new application (*e.g.*, native library files must not be compressed). We also found that
2421 some applications deliberately store files with names that will crash the zip library used by
2422 Apktool. For this reason, we also used our own library to modify the APK files. We took special
2423 care to process the least possible files in the APKs, and only strip the DEX files and signatures,
2424 before adding the new modified DEX files at the end.

2425 Unfortunately, we did not have time to compare the robustness of our solution to existing tools
2426 like Apktool and Soot, but we did a quick performance comparison, summarised in Section 5.5.5.
2427 In hindsight, we probably should have taken the time to find a way to use smali/backsmali (the
2428 backend of Apktool) as a library or use SootUp to do the instrumentation, but neither option
2429 has documentation to instrument applications this way. At the time of writing, the feature is
2430 still being developed, but in the future, Androguard might also become an option to modify
2431 DEX files. Nevertheless, we published our instrumentation library, Androscalpel, for anyone
2432 who wants to use it (see Appendix A).

2433 ∴

2434 Now that we saw the transformations we want to make, we know the runtime information we
2435 need to do it. In the next section, we will propose a solution to collect that information.

2436 **5.4 Collecting Runtime Information**

2437 To perform the transformations described in Section 5.3, we need information like the name and
2438 signature of the method called with reflection, or the actual bytecode loaded dynamically. We
2439 decided to collect that information through dynamic analysis. We saw in Chapter 2 different
2440 contributions that collect this kind of information. In the end, we decided to keep the analysis as
2441 simple as possible, so we avoided using a custom Android build like DexHunter and instead used
2442 Frida to instrument the application and intercept calls to the methods of interest. Section 5.4.1
2443 presents our approach to collect dynamically loaded bytecode, and Section 5.4.2 presents our
2444 approach to collect the reflection data. Because using dynamic analysis raises the concern of
2445 coverage, we also need some interaction with the graphical user interface of the application
2446 during the analysis. Ideally, a reverse engineer would do the interaction. Because we wanted to
2447 analyse many applications in a reasonable time, we replaced this engineer with an automated
2448 runner that simulates the interactions. We discuss this option in Section 5.4.3.

2449 5.4.1 Collecting the Dynamically Loaded Bytecode

2450 Initially, we considered instrumenting the constructor methods of the class loaders of the
2451 Android SDK. However, this is a significant number of methods to instrument, and looking
2452 at older applications, we realised that we missed the `DexFile` class. `DexFile` is now deprecated
2453 but still usable class that can be used to load bytecode dynamically. We initially missed
2454 this class because it is neither a `ClassLoader` class nor an SDK class (anymore). To avoid
2455 running into this kind of oversight again, we decided to look at the ART source code
2456 and list all the places where the internal functions used to parse bytecode are called. We
2457 found that all those calls are from under either `DexFile.openInMemoryDexFilesNative(..)`
2458 or `DexFile.openDexFileNative(..)`, two hidden API methods. As a reference, in 2015, Dex-
2459 Hunter [74] already noticed `DexFile.openDexFileNative(..)` (although in the end DexHunter
2460 intruments another function, `DefineClass(..)`). `DefineClass(..)` is still a good function to
2461 instrument, but it is a C++ native method that does not have a Java interface, making it
2462 harder to work with using Frida, and we want to avoid patching the source code of the ART like
2463 DexHunter did. For this reason, we decided to hook `DexFile.openInMemoryDexFilesNative(..)`
2464 and `DexFile.openDexFileNative(..)` instead. Those methods take a list of Android code files
2465 as argument, either in the form of in-memory byte arrays or file paths, and a reference to the
2466 classloader associated with the code. Instrumenting those methods allows us to collect all the
2467 code files loaded by the ART and associate them with their class loaders.

2468 5.4.2 Collecting Reflection Data

2469 As described in Section 5.3.1, there are 3 methods that we need to instrument to capture
2470 reflection calls: `Class.newInstance()`, `Constructor.newInstance(..)` and `Method.invoke(..)`.
2471 Because Java has polymorphism, we need not only the method name and defining class, but
2472 also the whole signature of the method. In addition to that, in case there are several classes
2473 with the same name as the defining class, we also need the classloader of the defining class to
2474 distinguish it from the other classes.

2475 *Where* the reflection method is called is more difficult to find. In order to correctly modify
2476 the application, we need to know which specific call to a reflection method we intercepted.
2477 Specifically, we need the caller method (once again, we need the method name, full signature,
2478 defining class and its classloader), and the exact instruction that called the reflection method (in
2479 case the caller method uses reflection several times in different sites). This information is more
2480 difficult to collect than one would expect. It is stored in the stack, but before the SDK 34, the
2481 stack was not directly accessible programmatically. Historically, when a reverse engineer needed
2482 to access the stack, they would trigger and catch an exception and get the stack from that
2483 exception. The issue with this approach is that data stored in exceptions is meant for debugging.
2484 In particular, the location of the call in the bytecode has a different meaning depending on
2485 the debug information encoded in the bytecode. It can either be the address of the bytecode

2486 instruction invoking the callee method in the instruction array of the caller method, or the line
2487 number of the original source code that calls the callee method. Fortunately, in the SDK 34,
2488 Android introduced the `StackWalker` API. This API allows to programatically travel the current
2489 stack and retrieve information from it, including the bytecode address of the instruction calling
2490 the callee methods. Considering that the line number is not a reliable information, we chose
2491 to use the new API, despite the restrictions that come with choosing such a recent Android
2492 version (it was released in October 2023, around 2 years ago, and less than 50% of the current
2493 Android market share supports this API today¹).

2494 5.4.3 Application Execution

2495 Dynamic analysis requires actually running the application. In order to test multiple appli-
2496 cations automatically, we needed to simulate human interactions with the applications. In
2497 Chapter 2, we presented a few solutions to explore an application dynamically. We first
2498 eliminated Sapienz [41], as it relies on an application instrumentation library called ELLA,
2499 which has not been updated for 9 years. We also chose to avoid the Monkey because we
2500 noticed that it often triggers events that close the application (events like pressing the ‘home’
2501 button, or opening the general settings drop-down menu at the top of the screen). Stoa [60]
2502 and GroddDroid [1] use UI Automator to interact with the application. UI Automator is a
2503 standard Android API intended for automatic testing. Both Soat and GroddDroid perform
2504 additional analysis on the application to improve the exploration. In the end, we elected to use
2505 the most basic execution mode of GroddDroid that does not need this additional analysis. It
2506 explores the application following a depth-first search algorithm. We chose this option to keep
2507 the exploration lightweight and limit the chance of crashing the analysis (we saw in Chapter 3
2508 the issues brought by complex analysis). It might be interesting in future work to explore more
2509 advanced exploration techniques.

2510 Because we are using Frida, we do not need to use a custom version of Android with a modified
2511 ART or kernel. However, we decided not to inject Frida into the original application. This means
2512 we need to have root access to directly run Frida in Android, which is not a normal thing to have
2513 on Android. Because dynamic analysis can be slow, we also decided to run the applications on
2514 emulators. This makes it easier to run several analyses in parallel. The alternative would have
2515 been to run the application on actual smartphones, and would have required multiple phones
2516 to run the analysis in parallel. For simplicity, we chose to use Google’s Android emulator for
2517 our experiment. We spawned multiple emulators, installed Frida on them, took a snapshot of
2518 the emulator before installing the application to analyse. Then we run the application for five
2519 minutes with GroddRunner, and at the end of the analysis, we reload the snapshot in case the

2520 1. <https://gs.statcounter.com/android-version-market-share/mobile-tablet/worldwide/#monthly-202401-202508>
2521

2522 application modified the system in some unforeseen way. If at some point an emulator stops
2523 responding for too long, we terminate it and restart it.

2524 As we will see in Section 5.5.1, our experimental setup is quite naive and still requires
2525 improvement. TODO: Comment on dit proprement que c'est tout péché? For example, we do not
2526 implement any anti-evasion techniques, which can be a significant issue when analysing malware.
2527 Nonetheless, the benefit of our implementation is that it only requires an ADB connection to
2528 a phone with a rooted Android system to work. Of course, to analyse a specific application, a
2529 reverse engineer could use an actual smartphone and explore the application manually. It would
2530 be a lot more stable than our automated batch analysis setup.

2531 \therefore

2532 Now that we saw both the dynamic analysis setup and the transformation we want to perform
2533 on the APKs, we put our proposed approach into practice. In the next section, we will run our
2534 dynamic analysis on APKs and study the data collected, as well as the impact the instrumen-
2535 tation has on applications and different analysis tools.

2536 5.5 Results

2537 To study the impact of our transformation on analysis tools, we reused applications from the
2538 dataset we sampled in Chapter 3. Because we are running the application on a recent version of
2539 Android (SDK 34), we only took the most recent applications: the one collected in 2023. This
2540 represents 5000 applications over the 62 525 total of the initial dataset. Among them, we could
2541 not retrieve 43 from Androzoo, leaving us with 4957 applications to test.

2542 We will first look at the results of the dynamic analysis and look at the bytecode we intercepted.
2543 Then, we will study the impact the instrumentation has on static analysis tools, notably on
2544 their success rate. Additionally, we will study with the analysis of a handcrafted application to
2545 check whether the instrumentation does, in fact, improve the results of analysis tools.

2546 5.5.1 Dynamic Analysis Results

2547 After running the dynamic analysis on our dataset the first time, we realised our dynamic
2548 setup was quite fragile. We found that 43.09% of the executions failed with various errors.
2549 The majority of those errors were related to failures to connect to the Frida agent or start
2550 the activity from Frida. Some of those errors seemed to come from Frida, while others seemed
2551 related to the emulator failing to start the application. We found that relaunching the analysis
2552 for the applications that failed was the simplest way to fix those issues, and after 6 passes,
2553 we went from 2136 to 209 applications that could not be analysed. The remaining errors look
2554 more related to the application itself or Android, with 96 errors being a failure to install the
2555 application, and 110 others being a null pointer exception from Frida.

2556 Unfortunately, although we managed to start the applications, we can see from the list of
 2557 activities visited by GroddDroid that a majority (84.77%) of the applications stopped before
 2558 even starting one activity. Some applications do not have any activities and are not intended to
 2559 interact with a user, but those are clearly a minority and do not explain such a high number.
 2560 We expected some issues related to the use of an emulator, like the lack of `x86_64` library in
 2561 the applications, or countermeasures aborting the application if an emulator is detected. We
 2562 manually looked at some applications, but did not find a notable pattern. In some cases, the
 2563 application was just broken – for instance, an application was trying to load a native library
 2564 that simply does not exist in the application. In other cases, Frida is to blame: we found some
 2565 cases where calling a method from Frida can confuse the ART. `protected` methods cannot be
 2566 called from a class other than the one that defined the method or one of its children. The issue
 2567 is that Frida might be considered by the ART as another class, leading to the ART aborting
 2568 the application. Table 11 shows the number of applications that we analysed, if we managed
 2569 to start at least one activity and if we intercepted code loading or reflection. It also shows the
 2570 average number of activities visited (when at least one activity was started). This average is
 2571 slightly higher than 1, which seems reasonable: a lot of applications do not need more than one
 2572 activity, but some do, and we did manage to explore at least some of those additional activities.
 2573 As shown in the table, even if the application fails to start an activity, sometimes it will still
 2574 load external code or use reflection.

2575 We later tested the applications on a real phone (model Nothing (2a), Android 15), without
 2576 Frida but still using GroddRunner. This time, we managed to visit at least one activity for 2130
 2577 applications, 3 times more than in our actual experiment. This shows that our setup is indeed
 2578 breaking applications, but also that there is still another issue we did not find: more than half
 2579 of the tested applications did not display any activities at all.

2580 The high number of applications that did not start an activity means that our results will be
 2581 highly biased. The code/method that might be loaded/called by reflection from inside activities

	nb apk	nb failed		activities visited		average nb
		1 st pass	6 th pass	0	≥ 1	activities when > 0
All	4957	2136	209	4025	723	1.3
With Reflection	3948			3298	650	1.3
With Code Loading	598			453	145	1.2

Table 11: Summary of the dynamic exploration of the applications from the RASTA dataset collected by Androzoo in 2023

Nb Occurrences	SHA 256	Content	Format
273	bee390afa2...	Lcom/facebook/ads/*	DEX
98	7aae06433c...	Lcom/facebook/ads/*	DEX
70	920e465a87...	Lcom/facebook/ads/*	DEX
31	51dd5ff34a...	Lcom/google/android/ads/*	APK
12	d44cfb6b41...	Lcom/google/android/ads/*	APK
9	be87bb0a50...	Lcom/facebook/ads/*	DEX
9	26fb1a7903...	Lcom/facebook/ads/*	DEX
7	8395a0121e...	Lcom/facebook/ads/*	DEX
7	1eea5584eb...	Lcom/google/android/ads/*	APK
6	94f66aa1ae...	Lcom/appsflyer/internal/*	DEX
...			

Table 12: Most common dynamically loaded files

2582 is filtered out by the limit of our dynamic execution. This bias must be kept in mind while
 2583 reading the next subsection that studies the bytecode that we intercepted.

2584 5.5.2 The Bytecode Loaded by Application

2585 We collected a total of 640 files for 598 application that we detected loading bytecode
 2586 dynamically. 92 of them were loaded by a `DexClassLoader`, 547 were loaded by a
 2587 `InMemoryDexClassLoader`, and 1 was loaded by a `PathClassLoader`.

2588 Once we compared the files, we found that we only collected 70 distinct files, and that 273 were
 2589 identical. Once we looked more in detail, we found that most of those files are advertisement
 2590 libraries. In total, we collected 87 files containing Google ads libraries and 492 files containing
 2591 Facebook ads libraries. In addition, we found 48 files containing code that we believe to be
 2592 AppsFlyer, a company that provides “measurement, analytics, engagement, and fraud protec-
 2593 tion technologies”. The remaining 13 files were custom code from high security applications
 2594 (*i.e.*, banking, social security). Table 12 summarises the information we collected about the
 2595 most common bytecode files.

2596 5.5.3 Impact on Analysis Tools

2597 We took the applications associated with the 13 unique DEX files we found to see the impact
 2598 of our transformation.

2599 The applications were indeed obfuscated, making a manual analysis tedious. We did not find
 2600 visible DEX files or APK files inside the applications, meaning the applications are either

2601 downloading or generating them from variables and assets at runtime. To estimate the scope of
 2602 the code we made available, we use Androguard to generate the call graph of the applications,
 2603 before and after the instrumentation. Table 13 shows the number of edges of those call graphs.
 2604 The columns before and after show the total number of edges of the graphs, and the diff
 2605 column indicates the number of new edges detected (*i.e.*, the number of edges after instru-
 2606 mentation minus the number of edges before). This number include edges from the bytecode
 2607 loaded dynamically, as well as the call added to reflect reflection calls, and calls to “glue”
 2608 methods (method like `Integer.intValue()` used to convert objects to scalar values, or calls to
 2609 `T.check_is Xxx_xxx(Method)` used to check if a `Method` object represents a known method). The
 2610 last column, “Added Reflection”, is the list of non-glue method calls found in the call graph of
 2611 the instrumented application but neither in the call graph of the original APK, nor in the call
 2612 graphs of the added bytecode files that we computed separately. This corresponds to the calls
 2613 we added to represent reflection calls.

2614 The first application, 0019d7fb6a..., is noticeable. The instrumented APK has ten times more
 2615 edges to its call graph than the original, and only one reflection call. This is consistent with the
 2616 behaviour of a packer: the application loads the main part of its code at runtime and switches
 2617 from the bootstrap code to the loaded code with a single reflection call.

2618 Unfortunately, our implementation of the transformation is imperfect and sometimes fails, as
 2619 illustrated by 5d2cd1d10a... in Table 13. However, over the 4748 applications whose dynamic
 2620 analysis finished in our experiment, 4681 were patched. The remaining 1.41% failed either due
 2621 to some quirk in the zip format of the APK file, because of a bug in our implementation when
 2622 exceeding the method reference limit in a single DEX file, or in the case of 5d2cd1d10a...,

APK SHA 256	Number of Call Graph edges			
	Before	After	Diff	Added Reflection
0019d7fb6a...	641	60 170	59 529	1
274b677449...	537 613	540 674	3061	26
34599c2499...	336 740	339 616	2876	29
35065c6834...	343 245	346 694	3449	26
e7b2fb02ff...	464 642	465 389	747	91
efececc03c...	243 647	243 925	278	23
f34ce1e7a8...	704 095	706 576	2481	28
5d2cd1d10a...	<i>Instrumentation Crashed</i>			

Table 13: Edges added to the call graphs computed by Androguard by instrumenting the applications

2623 because the application reused the original application classloader to load new code instead of
 2624 instantiated a new class loader (a behavior we did not expect as possible using only the SDK,
 2625 but enabled by hidden APIs). Taking into account the failure from both dynamic analysis and
 2626 the instrumentation process, we have a 5.57% failure rate. This is a reasonable failure rate, but
 2627 we should keep in mind that it adds up to the failure rate of the other tools we want to use on
 2628 the patched application.

2629 To check the impact on the finishing rate of our instrumentation, we then run the same
 2630 experiment we ran in Chapter 3. We ran the tools on the APK before and after instrumentation,
 2631 and compared the finishing rates in Figure 18 (without taking into account APKs we failed
 2632 to patch¹).

2633 The finishing rate comparison is shown in Figure 18. We can see that in most cases, the finishing
 2634 rate is either the same or slightly lower for the instrumented application. This is consistent with
 2635 the fact that we add more bytecode to the application, hence adding more opportunities for
 2636 failure during analysis. They are two notable exceptions: Saaf and IC3. The finishing rate of IC3,
 2637 which was previously reasonable, dropped to 0 after our instrumentation, while the finishing
 2638 rate of Saaf jumped to 100%, which is extremely suspicious. Analysing the logs of the analysis
 2639 showed that both cases have the same origin: the bytecode generated by our instrumentation
 2640 has a version number of 37 (the version introduced by Android 7.0). Unfortunately, neither the




Figure 18: Exit status of static analysis tools on original APKs (left) and patched APKs (right)

2641 1. Due to a handling error during the experiment, the figure shows the results for 4274 APKs instead of 4681.
 2642 We also ignored the tool from Wognsen *et al.* [67] due to the high number of timeouts

2643 version of Apktool used by Saaf nor Dare (the tool used by IC3 to convert Dalvik bytecode to
2644 Java bytecode) recognises this version of bytecode, and thus failed to parse the APK. In the
2645 case of Dare and IC3, our experiment correctly identifies this as a crash. On the other hand,
2646 Saaf do not detect the issue with Apktool and pursues the analysis with no bytecode to analyse
2647 and returns a valid return file, but for an empty application.

2648 [TODO: Flowdroid results are inconclusive: some apks have more leak after and as many apks
2649 have less? also, runing flowdroid on the same apk can return a different number of leak???

2650 **5.5.4 Example**

2651 In this subsection, we use our approach on a unique APK to look in more detail into the
2652 analysis of the transformed application. We handcrafted this application for the purpose of
2653 demonstrating how this can help a reverse engineer in their work. Accordingly, this application
2654 is quite small and contains both dynamic code loading and reflection. We defined a method
2655 `Utils.source()` and `Utils.sink()` to model a method that collects sensitive data and a method
2656 that exfiltrates data respectively. Those methods are the ones we will use with Flowdroid to
2657 track data flows.

2658 A first analysis of the content of the application shows that the application contains one `Activity`
2659 that instantiates the class `Main` and calls `Main.main()`. Listing 13 shows most of the code of
2660 `Main` as returned by `Jadx`. We can see that the class contains another DEX file encoded in base
2661 64 and loaded in the `InMemoryDexClassLoader c1` (line 7). A class is then loaded from this class
2662 loader (line 11), and two methods from this class loader are called (line 14). The names of
2663 this class and methods are not directly accessible as they have been ciphered and are decoded
2664 just before being used at runtime. Here, the encryption key is available statically (line 6), and
2665 in theory, a very good static analyser implementing Android `Cipher` API could compute the
2666 actual methods called. However, we could easily imagine an application that gets this key from
2667 a remote command and control server. In this case, it would be impossible to compute those
2668 methods with static analysis alone. When running Flowdroid on this application, it computed a
2669 call graph of 43 edges on this application, and no data leaks. This is not particularly surprising
2670 considering the obfuscation methods used.

2671 Then we run the dynamic analysis we described in Section 5.4 on the application and apply
2672 the transformation described in Section 5.3 to add the dynamic information to it. This time,
2673 Flowdroid computes a larger call graph of 76 edges, and does find a data leak. Indeed,
2674 when looking at the new application with `Jadx`, we notice a new class `Malicious`, and the
2675 code of `Main.main()` is now as shown in Listing 14: the method called in the loop is either
2676 `Malicious.get_data`, `Malicious.send_data()` or `Method.invoke()` (lines 9, 11 and 12). Although

```

1 package com.example.theseus;
2
3 public class Main {
4     private static final String DEX = "ZGV4CjA [...] EAAABEAwAA";
5     Activity ac;
6     private Key key = new SecretKeySpec("_-__Secret Key_-__".getBytes(),
7 "AES");
8     ClassLoader cl = new
9 InMemoryDexClassLoader(ByteBuffer.wrap(Base64.decode(DEX, 2)),
10 Main.class.getClassLoader());
11
12     public void main() throws Exception {
13         String[] strArr = {"n6WGYJzjDrUvR9cYljNlw==", "dapES0wL/
14 iFIPuMnH3fh7g=="};
15         Class<?> loadClass =
16 this.cl.loadClass(decrypt("W5f3xRf3wCSYcYG7ckYGR5xuuESDZ2NcDUzGxsq3sIs="));
17         Object obj = "imei";
18         for (int i = 0; i < 2; i++) {
19             obj = loadClass.getMethod(decrypt(strArr[i]),
20 String.class, Activity.class).invoke(null, obj, this.ac);
21         }
22     }
23     public String decrypt(String str) throws Exception {
24         Cipher cipher = Cipher.getInstance("AES/ECB/PKCS5Padding");
25         cipher.init(2, this.key);
26         return new String(cipher.doFinal(Base64.decode(str, 2)));
27     }
28     ...
29 }

```

Listing 13: Code of the main class of the application, as shown by Jadx, before patching

2677 self-explanatory, verifying the code of those methods indeed confirms that `get_data()` calls
2678 `Utils.source()` and `send_data()` calls `Utils.sink()`.

2679 For a higher-level view of the method, we can also look at its call graph. We used Androguard
2680 to generate the call graphs in Figure 19 and Figure 20¹. Figure 19 shows the original call graph,
2681 and gives a good idea of the obfuscation methods used: we can see calls to `Main.decrypt(String)`
2682 that itself calls cryptographic APIs, as well as calls to `ClassLoader.loadClass(String)`,
2683 `Class.getMethod(String, Class[])` and `Method.invoke(Object, Object[])`. This indicates
2684 reflection calls based on ciphered strings, but does not reveal what the method actually does.

2685 ¹. We manually edited the generated `.dot` files for readability.

```

1 public void main() throws Exception {
2     String[] strArr = {"n6WGYJzjDrUvR9cYljlnlw==", "dapES0wl/
iFiPuMnH3fh7g=="};
3     Class<?> loadClass =
this.cl.loadClass(decrypt("W5f3xRf3wCSYcYG7ckYGR5xuuESDZ2NcDUzGxsq3sIs="));
4     Object obj = "imei";
5     for (int i = 0; i < 2; i++) {
6         Method method = loadClass.getMethod(decrypt(strArr[i]),
String.class, Activity.class);
7         Object[] objArr = {obj, this.ac};
8         obj =
T.check_is_Malicious_get_data_fe2fa96eab371e46(method) ?
9             Malicious.get_data((String) objArr[0], (Activity)
objArr[1]) :
10            T.check_is_Malicious_send_data_ca50fd7916476073(method) ?
11            Malicious.send_data((String) objArr[0], (Activity)
objArr[1]) :
12            method.invoke(null, objArr);
13     }
14 }

```

Listing 14: Code of Main.main(), as shown by Jadx, after patching

2686 In comparison, Figure 20, the call graph after instrumentation, still shows the cryptographic
2687 and reflection calls, as well as four new method calls. In grey on the figure, we can see the
2688 glue methods (T.check_is_Xxx_xxx(Method)). Those methods are part of the instrumentation




Figure 19: Call Graph of Main.main() generated by Androguard before patching




Figure 20: Call Graph of Main.main() generated by Androguard after patching

2689 process presented in Section 5.3, but do not bring a lot to the analysis of the call graph. In red
 2690 on the figure however, we have the calls that were hided by reflection in the first call graph,
 2691 and thank to the bytecode of the methods called being injected in the application, we can also
 2692 see that they call `Utils.source(String)` and `Utils.sink(String)`, the methods we defined for
 2693 this application as source of confidential data and exfiltration method.

2694 5.5.5 Androscalpel Performances

2695 Because we implemented our own instrumentation library, we wanted to compare it to other
 2696 existing options. Unfortunately, we did not have time to compare the robustness and correctness
 2697 of the generated applications. However, we did compare the performances of our library,
 2698 Androscalpel, to Apktool and Soot, over the first 100 applications of RASTA (in alphabetical
 2699 order of the SHA256).

2700 Due to time constraints, we could not test a complex transformation, as adding registers requires
 2701 complex operations for both Androscalpel and Apktool (see Section 5.3.4 for more details).
 2702 We decided to test two operations: travelling the instructions of an application (a read-only
 2703 operation), and regenerating an application, without modification (a read/write operation). It
 2704 should be noted that all three of the tested tools have multiprocessing support, but we disabled
 2705 the option when testing the generation of an application with Soot, as it raised errors.

2706 Table 14 compares the resources consumed by each tool for each operation. We can see that
 2707 for read-only operation, we are 16 times faster than Soot and 8 times faster than Apktool,
 2708 while keeping a smaller memory footprint. When generating an application, the gap lessens,
 2709 but we are still almost 8 times faster than Soot. Some of this difference probably comes from
 2710 implementation choices: Soot and Apktool are implemented in Java, which has a noticeable
 2711 overhead compared to Rust. However, a noticeable part of this difference can also be explained
 2712 by the specialised nature of our library; we did not implement all the features Soot has, and we

Tool		Soot	Apktool	Androscalpel
Read	Time (s)	73.95	33.09	4.38
	Mem (GB)	2.33	1.38	0.58
	Detected Crashes	0	0	0
Read/Write	Time (s)	156.58	74.89	20.34
	Mem (GB)	4.02	1.92	1.19
	Detected Crashes	10	4	1

Table 14: Average time and memory consumption of Soot, Apktool and Androscalpel

2713 do not parse Android resources like Apktool does. Having better performances does not means
2714 that our solution can replace the other in all cases.

2715 Nevertheless, it should be noted that over the 100 applications tested, Soot failed to regenerate
2716 10 of them, Apktool 4, and Androscapel only 1, showing that our efforts to limit crashes were
2717 successful.

2718 \therefore

2719 To conclude, we showed that our approach indeed improves the results of analysis tools without
2720 impacting their finishing rates much. Unfortunately, we also noticed that our dynamic analysis
2721 is suboptimal, either due to our experimental setup or due to our solution to explore the
2722 applications. In the next section, we will present in more detail the limitations of our solution,
2723 as well as future work that can be done to improve the contributions presented in this chapter.

2724 5.6 Limitations and Future Works

2725 The method we presented in this chapter has a number of underdeveloped aspects. In this
2726 section, we will present those issues and potential avenues of improvement related to the
2727 bytecode transformation, the dynamic analysis and DroidRA, a tool similar to our solution.

2728 5.6.1 Bytecode Transformation

2729 *Custom Class Loaders* The first obvious limitation of our bytecode transformation is that we
2730 do not know what custom class loaders (class loaders implemented by the application developer,
2731 as opposed to the class loaders in the SDK) do, so we cannot accurately reproduce statically
2732 their behaviour. For instance, we can imagine a class loader that loads all classes whose name
2733 starts with an A from one DEX file, and all classes whose name starts with a B from another. If
2734 both DEX files have colliding classes, our implementation will not select the right classes. We
2735 elected to fallback to the behaviour of the `BaseDexClassLoader`, which is the highest Android-
2736 specific class loader in the inheritance hierarchy, and whose behaviour is shared by all class
2737 loaders except `DelegateLastClassLoader`. The current implementation of the ART enforces some
2738 restrictions on the class loader's behaviour to optimise the runtime performance by caching
2739 classes. This gives us some guarantees that custom class loaders will keep some coherence
2740 with the classic class loaders. For instance, a class loaded dynamically must have the same
2741 name as the name used in `ClassLoader.loadClass()`. This makes `BaseDexClassLoader` a good
2742 approximation for legitimate class loaders. However, an obfuscated application could use the
2743 techniques discussed in Section 4.5.2, in which case our model would be entirely wrong.

2744 It would be interesting to explore if some form of static analysis, like symbolic execution, could
2745 be used to extract the behaviour of an ad hoc class loader and be used to model the class used
2746 appropriately. A more reasonable approach would be to improve the static analysis to intercept
2747 each call of `loadClass()` of each class loader, including implicit calls performed by the ART.

2748 This would allow us to collect a mapping (class loader, class name) → class that can then be
2749 used when renaming colliding classes.

2750 *Multiple Class Loaders for one Method.invoke()* Although we managed to handle calls to
2751 different methods from one `Method.invoke()` site, we do not handle calling methods from dif-
2752 ferent class loaders with colliding class definitions. The first reason is that it is quite challenging
2753 to compare class loaders statically. At runtime, each object has a unique identifier that can
2754 be used to compare them over the course of the same execution, but this identifier is reset
2755 each time the application starts. This means we cannot use this identifier in an `if` condition to
2756 differentiate the class loaders. Ideally, we would combine the hash of the loaded DEX files, the
2757 class loader class and parent to make a unique, static identifier, but the DEX files loaded by a
2758 class loader cannot be accessed at runtime without accessing the process memory at arbitrary
2759 locations. For some class loaders, the string representation returned by `Object.toString()` lists
2760 the location of the loaded DEX file on the file system. This is not the case for the commonly
2761 used `InMemoryClassLoader`. In addition, the DEX files are often located in the application's
2762 private folder, whose name is derived from the hash of the APK itself. Because we modify the
2763 application, the path of the private folder also changes, and so will the string representation of
2764 the class loaders. Checking the class loader of a class can also have side effects on class loaders
2765 that delegate to the main application class loader: because we inject the classes in the APK, the
2766 classes of the class loader are now already in the main application class loader, which in most
2767 cases will have priority over the other class loaders, and lead to the class being loaded by the
2768 application class loader instead of the original class loader. If we check for the class loader, we
2769 would need to consider such cases and rename each class of each class loader before reinjecting
2770 them in the application. This would greatly increase the risk of breaking the application during
2771 its transformation. Instead, we elected to ignore the class loaders when selecting the method
2772 to invoke. This leads to potential invalid runtime behaviour, as the first method that matches
2773 the class name will be called, but the alternative methods from other class loaders still appear
2774 in the new application, albeit in a block that might be flagged as dead code by a sufficiently
2775 advanced static analyser.

2776 *ClassNotFoundException may not be raised* In the very specific situation where the original
2777 application tries to access a class from dynamically loaded bytecode without actually accessing
2778 this bytecode (*e.g.*, by using the wrong class loader), the patched application behaviour will
2779 differ. The original application should raise a `ClassNotFoundException`, but in the patched
2780 application, the class will be accessible, and the exception will not be raised. In practice, there
2781 is not a lot of reasons to do such a thing. One could be to check if the APK has been tampered
2782 with, but there are easier ways to do this, like checking the application signature. Another
2783 would be to check if the class is already available, and if not, load it dynamically, in which
2784 case it does not matter, as code loaded dynamically is already present. In any case, because we

2785 remove neither the calls to the function that load the classes (like `ClassLoader.loadClass(...)`)
2786 nor the `try / catch` blocks, static analysis tools that can handle the original behaviour should
2787 still be able to access the old behaviour.

2788 **5.6.2 Dynamic Analysis**

2789 *Anti Evasion* Our dynamic analysis does not perform any kind of anti-evasive technique. Any
2790 application implementing even basic evasion will detect our environment and will probably not
2791 load malicious bytecode. Running the dynamic analysis in an appropriate sandbox, such as
2792 DroidDungeon [55], should improve the results significantly.

2793 *Code Coverage* In Section 5.5.1, we saw that our dynamic analysis performed poorly. It may
2794 be due to our experimental setup, and it is possible that a better sandbox will fix the issue.
2795 However, there is a larger code coverage issue. We tried to manually analyse a few applications
2796 marked as malware on MalwareBazaar to test our method. Although we did confirm that the
2797 applications were using reflection and dynamic code loading with a static analysis, we did not
2798 manage to trigger this behaviour at runtime, and other obfuscation techniques make it very
2799 difficult to determine statically the required condition to trigger them. Thus, we believe that
2800 techniques to improve code coverage are indeed needed when analysing applications. This could
2801 mean better exploration techniques, such as the one implemented by Stoa and GroddDroid,
2802 or more intrusive approaches, such as forced execution.

2803 **5.6.3 Comparison with DroidRA and Other Tools**

2804 It would be very interesting to compare our tool to DroidRA. DroidRA is a tool that computes
2805 reflection information using static analysis and patches the application to add those calls.
2806 Beyond the classic comparison of static versus dynamic, DroidRA has a similar goal and strategy
2807 to ours. Two notable comparison criteria would be the failure rate and the number of edges
2808 added to an application call graph. The first criterion indicates how much the results can be
2809 used by other tools, while the second indicates how effective the approaches are.

2810 Because we elected to make our own software to modify the bytecode of the APKs, it would
2811 be insightful to compare the finishing rate and performances of simple transformations with
2812 our tool, to the same transformation made with Apktool, Soot or SootUp (we only compared
2813 the performances for re-generating and application without transformations). An example of a
2814 transformation to test would be to log each method call and its return value. More than finding
2815 which solution is the best to instrument an application, this would allow us to compare the
2816 weaknesses of each tool and find if some recurring issues for some tools can be solved using
2817 a technical solution implemented by another tool (*e.g.*, some applications deliberately include
2818 files with names that crash the standard Java zip library).

2819 5.7 Conclusion

2820 In this chapter, we presented a set of transformations to encode reflection calls and code loaded
2821 dynamically inside the application. We also presented a dynamic analysis approach to collect
2822 the information needed to perform those transformations.

2823 We then applied this method to a recent subset of applications of our dataset from Chapter 3.
2824 When comparing the success rate of the tools of Chapter 3 on the applications before and after
2825 the transformation, we found that, in general, the success rate of those tools slightly decreases
2826 (except for a few tools). We also showed that our transformation allows static analysis tools to
2827 access and process that runtime information in their analysis. However, a more in-depth look
2828 at the results of our dynamic analysis showed that our code coverage is lacking, and that the
2829 great majority of dynamically loaded code we intercepted is from generic advertisement and
2830 telemetry libraries.

2831 **Pb3:** *Can we use instrumentation to provide dynamic code*
2832 *loading and reflection data collected dynamically to static analysis*
2833 *tools and improve their results?*

2834 We showed that instrumentation can be used to add direct calls to
2835 methods initially called through reflections, which, combined with
2836 the injection in the application of dynamically loaded bytecode,
2837 allows generic static analysis tools to access previously unavail-
2838 able code. However, we also found that the dynamic analysis can
2839 be a significant bottleneck in this approach.

2866 classes from the SDK are added either for retro-compatibility or due to the developer being
2867 unaware that a library was already present in the Android SDK, and the few cases where classes
2868 are present multiple times in the application appear to be mistakes during the compilation of
2869 the application. Still, 23.52% of the applications we tested were shadowing classes from the
2870 targeted SDK version.

2871 Lastly, we proposed a solution to reuse any static analysis tool on an application that uses
2872 dynamic code loading or reflection. To do so, we collect the relevant information dynamically,
2873 then instrument the application to encode the dynamic information inside a valid application
2874 mimicking the dynamic behaviour of the original one. This new application can then be analysed
2875 normally by any tool that accepts an application as input. We tested our method on a subset
2876 of recent applications from the dataset of our first contribution. The results of our dynamic
2877 analysis suggest that we failed to correctly explore many applications, hinting at weaknesses in
2878 our experimental setup. Nonetheless, we did obtain some dynamic data, allowing us to pursue
2879 our experiment. We compared the finishing rate of tools on the original application and the
2880 instrumented application using the same experiment as in our first contribution, and found
2881 that, in general, the instrumentation only slightly reduces the finishing rate of analysis tools.
2882 We also confirmed that the instrumentation improves the result of analysis tools, allowing them
2883 to compute more comprehensive call graphs of the applications, or to detect new data flows.

2884 **6.2 Perspectives for Future Work**

2885 In this section, we present what, in light of this thesis, we believe to be worthwhile avenues of
2886 work to improve the Android reverse engineering ecosystem.

2887 The main issues that appeared in all our work appear to be engineering ones. The errors we
2888 analysed in Chapter 3 showed that even something that should be basic, reading the content of
2889 an application, can be challenging. Chapter 4 also showed that reproducing the exact behaviour
2890 of Android is more difficult than it seems (in our specific case, it was the class loading algorithm,
2891 but we can expect other features to have similar edge cases). As long as those issues are not
2892 solved, we cannot build robust analysis tools.

2893 One avenue that is more research-oriented and that should be investigated would be to reuse
2894 for analysis purposes the code actually used by Android. For instance, the parsing of DEX,
2895 APK, and resource files could be done using the same code as the ART. This is possible thanks
2896 to AOSP being open-source. However, this is not straightforward. Dynamic analysis relying
2897 on patched versions of the AOSP showed that it is difficult to maintain this kind of software
2898 over time. Doing this would require limiting the modifications to the actual source code of
2899 Android to minimise the changes needed at each Android update. Another obstacle to overcome
2900 is to decouple the compilation of the tool from the rest of AOSP: it is a massive dependency
2901 that needs a lot of resources to build. Having such a dependency would be a barrier to entry,

2902 preventing others from modifying or improving the tool. Should those issues be solved, directly
2903 using the code from AOSP would allow such a tool to stay up to date with Android and limit
2904 discrepancies between what Android does and what the tool sees.

2905 An orthogonal solution to this problem of not being able to analyse edge cases is to create a new
2906 benchmark to test the capacity of a tool to handle real-life applications. Benchmarks are usually
2907 targeted at some specific technique (*e.g.*, taint tracking), and accordingly, test for issues specific
2908 to the targeted technique (*e.g.*, accurately tracking data that passes through an array). We
2909 suggest using a similar method to what we did in Chapter 3 to keep the benchmark independent
2910 from the tested tools. Instead of checking the correctness of the tools, this benchmark should test
2911 if the tool is able to finish its analysis. Applications in this benchmark could either be real-life
2912 applications that was proven to be difficult to analyse (for instance, applications that crashed
2913 many of the tested tools in Chapter 3), or hand-crafted applications reproducing corner cases
2914 or anti-reverse techniques encountered while analysing obfuscated applications (for instance, an
2915 application with gibberish binary file names inside `META-INF/` that can crash Jadx zip reader).
2916 The main challenge with such a benchmark is that it would need frequent updates to follow
2917 Android evolutions, and be diverse enough to encompass a large spectrum of possible issues.

2918 Lastly, our experience with dynamic analysis led us to believe that there is a need for a new
2919 protocol and/or API for automatic testing. Currently, except for intrusive methods like instru-
2920 mentations, interacting with an application is done through a mix of ADB and UI Automator.
2921 Unfortunately, those tools give very poor feedback. Information about the execution is mostly
2922 found in the Android logs, lost among other system events, and it is difficult to filter the events
2923 related to the application without losing critical data. For instance, exception logs are usually
2924 linked to the application in the Android logs, but some exceptions originating from the ART
2925 (and not the code of the application itself) are logged as system errors. Similarly, an application
2926 can fail to install, or successfully install but later be quarantined or uninstalled by the Android
2927 operating system, without clear feedback. Another issue is that more and more, Android will
2928 require interactions with the system (for instance, with a security permission pop-up), and those
2929 interactions are not handled by UI Automator. Similarly, when an application opens another
2930 one for a specific task (opening a document, for example), what is happening is unclear from
2931 the standpoint of UI Automator. We think that an API or protocol that merges and delivers in
2932 a structured way all those informations, and allows access to the relevant component (such as
2933 a system popup) would be beneficial both for automated testing and dynamic analysis.

2934 Integrating such a protocol into Android would open interesting perspectives. For instance,
2935 we could imagine Google requiring applications requesting critical permissions to provide test
2936 inputs with a high code coverage (maybe even 100% of coverage). Those tests would incentivise
2937 application developers to provide better quality code for applications handling sensitive data,
2938 but also to provide solutions for the coverage issue that comes with dynamic analysis. Requiring

2939 a high code coverage would force the developer to supply solutions for situations normally
2940 requiring human interaction. For example, if an application requires the user to authenticate
2941 themselves, the developer would need to provide a testing account that can then be used for tests
2942 and analysis. Of course, we can expect malicious applications to implement evasion techniques
2943 when they detect an analysis following the tests they provided, but code coverage can be
2944 checked, and imposing constraints on the coverage of the tests should mitigate evasion.

APPENDICES

RELEASED SOFTWARE AND ARTIFACTS

2948 In Chapter 3, we mentioned that we had some difficulties finding some software listed by Li *et*
2949 *al.* following the disappearance of the original websites hosting it. To limit the risk of having
2950 the same issue, we hosted the different pieces of software we released for this thesis in several
2951 locations. This appendix lists the software we released as well as the different places they can
2952 be found.

2953 A.1 RASTA

2954 The code used in Chapter 3 is available at those locations:

- 2955 • <https://gitlab.inria.fr/pirat-public/android/rasta>
- 2956 • <https://git.mineau.eu/these-android-re/rasta>
- 2957 • <https://github.com/histausse/rasta>
- 2958 • <https://doi.org/10.5281/zenodo.10137904>

2959 The exact version of the code used in Chapter 3 is tagged as `icsr2024` in the git repositories
2960 and corresponds to the one stored in Zenodo.

2961 The results of our experiment and the list of applications in the dataset are also available in
2962 the Zenodo archive¹.

2963 The container images used to run the different tools are available on Zenodo at <https://doi.org/10.5281/zenodo.10980349> as Singularity images, and on Dockerhub under the names:

- 2965 • `histausse/rasta-adagio:icsr2024`
- 2966 • `histausse/rasta-amandroid:icsr2024`
- 2967 • `histausse/rasta-anadroid:icsr2024`
- 2968 • `histausse/rasta-androguard-dad:icsr2024`
- 2969 • `histausse/rasta-androguard:icsr2024`
- 2970 • `histausse/rasta-apparecium:icsr2024`
- 2971 • `histausse/rasta-blueseal:icsr2024`
- 2972 • `histausse/rasta-dialdroid:icsr2024`
- 2973 • `histausse/rasta-didfail:icsr2024`
- 2974 • `histausse/rasta-droidsafe:icsr2024`

2975 1. <https://doi.org/10.5281/zenodo.10137904>

- 2976 • `histausse/rasta-flowdroid:icsr2024`
- 2977 • `histausse/rasta-gator:icsr2024`
- 2978 • `histausse/rasta-ic3-fork:icsr2024`
- 2979 • `histausse/rasta-ic3:icsr2024`
- 2980 • `histausse/rasta-iccta:icsr2024`
- 2981 • `histausse/rasta-malldroid:icsr2024`
- 2982 • `histausse/rasta-redexer:icsr2024`
- 2983 • `histausse/rasta-saaf:icsr2024`
- 2984 • `histausse/rasta-wognsen:icsr2024`

2985 **A.2 Shadow Attack Survey Dataset**

2986 The list of applications we scanned in Chapter 4, as well as the lists of platform classes, fields
2987 and, methods we extracted from the emulators for Android SDKs 32, 33, and 34, are stored on
2988 Zenodo at <https://doi.org/10.5281/zenodo.15846481>.

2989 The experiment we used to survey in-the-wild applications is available here:

- 2990 • https://gitlab.inria.fr/pirat-public/android/android_class_shadowing_scanner
- 2991 • https://git.mineau.eu/these-android-re/android_class_shadowing_scanner
- 2992 • https://github.com/histausse/android_class_shadowing_scanner

2993 **A.3 Theseus**

2994 The scripts we used for dynamic analysis and the code implementing the transformations
2995 described in Chapter 5 are available at the following locations under GPL license:

- 2996 • <https://gitlab.inria.fr/pirat-public/android/android-of-theseus>
- 2997 • https://git.mineau.eu/these-android-re/android_of_theseus
- 2998 • https://github.com/histausse/android_of_theseus

2999 The application transformations rely on Androscalpel, the crate we developed to manipulate
3000 Dalvik bytecode. Androscalpel can be found at the following locations:

- 3001 • <https://gitlab.inria.fr/pirat-public/android/androscalpel>
- 3002 • <https://git.mineau.eu/these-android-re/androscalpel>
- 3003 • <https://github.com/histausse/androscalpel>

3004 The dataset, results of the dynamic analysis and results of benchmark before and after the
3005 instrumentations with Theseus are available on Zenodo at [https://doi.org/10.5281/zenodo.](https://doi.org/10.5281/zenodo.17350991)
3006 17350991.

RÉSUMÉ SUBSTANTIEL EN FRANÇAIS

3009 Il y a deux réponses à cette question, comme à toutes les questions : celle du savant et celle
3010 du poète.

3011 — Ellana Caldin, *Le Pacte des Marchombres*, Tome 1: Ellana, de Pierre Bottero

3012 B.1 Introduction

3013 Android est le système d'exploitation pour téléphones portables le plus utilisé depuis 2014, et
3014 depuis 2017, il surpasse même Windows toutes plateformes confondues. Cette popularité en fait
3015 une cible de choix pour les acteurs malveillants. Il est donc important d'être capable d'analyser
3016 une application pour savoir exactement ce qu'elle fait. Ce processus est appelé l'ingénierie
3017 inverse.

3018 Beaucoup de travail a été fait dans ce domaine pour les programmes d'ordinateur. Toutefois,
3019 les applications Android présentent leur propres difficultés. Par exemple, les applications sont
3020 distribuées dans leur format spécifique, le format APK, et le code des applications est lui-même
3021 compilé dans un format de code à octets spécifique à Android: Dalvik. La première difficulté
3022 pour l'ingénieur·e inverse est donc d'avoir des outils qui comprennent les formats utilisés par
3023 Android. Dans le processus d'analyse, une première étape serait alors de lire le contenu de
3024 l'application. Des outils comme Apktool peuvent être utilisés pour convertir les fichiers binaire
3025 de l'application dans une version lisible par un·e humain·e. D'autres comme Jadx essaient de
3026 générer le code source Java depuis le code à octets. Toutefois, les applications Android peuvent
3027 être très grosses et il n'est pas toujours possible de les analyser manuellement. D'autres outils ont
3028 été développés pour extraire une représentation plus haut niveau du contenu de l'application.
3029 Par exemple, Flowdroid a pour objectif de détecter les fuites d'informations: l'utilisateur·ice
3030 définit une liste de méthodes qui génèrent des informations sensibles, et une liste de méthodes
3031 qui exfiltrent des informations vers l'extérieur. Flowdroid va alors calculer s'il existe des chemins

3032 dans l'application permettant de relier des méthodes de la première catégorie avec des méthodes
3033 de la seconde.

3034 Malheureusement, ces outils sont difficiles à utiliser, et même s'ils fonctionnent sur des applica-
3035 tions simples construites dans le but de tester les outils, il n'est pas rare que ces outils échouent
3036 sur de vraies applications. Cela pose la problématique suivante: *À quel point les outils d'analyse*
3037 *Android préalablement publiés sont-ils utilisables aujourd'hui, et quels facteurs impactent leur*
3038 *réutilisabilité?*

3039 Il y a deux familles d'analyse: l'analyse statique et l'analyse dynamique. L'analyse statique
3040 analyse l'application sans la lancer, alors que l'analyse dynamique examine le comportement de
3041 l'application pendant son exécution. Chacune a ses forces et ses faiblesses, et certains problèmes
3042 d'analyse sont traditionnellement associés à l'une ou l'autre pour les résoudre. L'un de ces
3043 problèmes est le chargement dynamique de code. Les applications Android sont initialement
3044 prévues pour être codées en Java, Android a donc hérité de nombreuses fonctionnalités de
3045 Java. En l'occurrence, Android a un système de chargeur de classes similaire à celui de Java,
3046 qui peut être utilisé pour charger, en cours d'exécution, du code extérieur à l'application. Etant
3047 donné que ce code chargé dynamiquement n'est pas nécessairement disponible dans l'application
3048 initialement, ce problème est relégué à l'analyse dynamique. Toutefois, il semblerait qu'une
3049 généralisation hâtive soit souvent faite, et que le système de chargement de classe dans son
3050 ensemble soit relégué à l'analyse dynamique. L'absence d'étude détaillée de ce mécanisme nous
3051 amène à notre seconde problématique: *Quel est l'algorithme de chargement de classe utilisé par*
3052 *défaut par Android, et est-ce qu'il impacte l'analyse statique?*

3053 Un autre problème usuellement associé à l'analyse dynamique est la réflexion. Android permet
3054 à une application de manipuler les classes et méthodes sous forme d'objet. En utilisant cette
3055 fonctionnalité, il est donc possible d'appeler une méthode en utilisant son nom sous forme
3056 de chaîne de caractère au lieu d'utiliser une instruction Dalvik avec une référence vers la
3057 méthode appelée. Ce cas est déjà compliqué à analyser statiquement quand la chaîne de
3058 caractère est lisible dans l'application, mais il devient impossible quand elle ne l'est pas (*ex.*
3059 *la chaîne est envoyée par un serveur externe lors de l'exécution, ou elle est stockée chiffrée*
3060 *et n'est déchiffrée qu'au dernier moment). L'analyse dynamique permet de capturer à la fois*
3061 *le code chargé dynamiquement et les méthodes appelées par réflexion. Toutefois, obtenir ces*
3062 *instructions est insuffisant. Il n'existe pas de solution standard pour transmettre ces données*
3063 *aux outils d'analyse statique, qui pourtant peuvent en avoir besoin pour analyser l'application*
3064 *dans son entièreté. Certaines contributions d'ingénierie inverse ont déjà proposé d'instrumenter*
3065 *(modifier) l'application pour y ajouter les résultats de leur analyse avant de l'analyser avec*
3066 *d'autres outils. Cette approche prometteuse motive notre troisième problématique: *Peut-on**
3067 *utiliser l'instrumentation pour fournir le code chargé dynamiquement et les informations de*
3068 *réflexion collectées dynamiquement aux outils d'analyse statique pour améliorer leurs résultats?*

3069 **B.2 Evaluation de la réutilisabilité des outils d’analyse statique pour** 3070 **Android**

3071 Dans ce chapitre, nous étudions la réutilisabilité d’outils d’analyse statique publiés entre 2011
3072 et 2017. Le but de cette étude n’est pas de quantifier la précision des outils, car ces outils ont
3073 des objectifs finaux différents. Au contraire, dans ce chapitre nous allons considérer comme
3074 correct tout résultat renvoyé par les outils, et uniquement compter les occurrences où les outils
3075 échouent à calculer un résultat quel qu’il soit.

3076 Les questions auxquelles nous voulons répondre sont:

3077 **QR1** Quels outils d’analyse statique pour Android vieux de plus de 5 ans peuvent encore être
3078 utilisés aujourd’hui avec un effort raisonnable?

3079 **QR2** Comment la réutilisabilité des outils évolue-t-elle avec le temps, en particulier pour
3080 l’analyse d’applications publiées avec plus de 5 ans d’écart avec l’outil?

3081 **QR3** Est-ce que la réutilisabilité des outils change quand on analyse une application bénigne
3082 comparé à un malicieux?

3083 Nous basons notre étude sur revue de littérature systématique de Li *et al.* qui liste les contri-
3084 butions accompagnées d’outils sous licence libre. Nous avons retrouvé les outils en questions,
3085 listé dans le Table 15. Nous avons éliminé les outils utilisant de l’analyse dynamique en plus
3086 de l’analyse statique, et vérifié la présence des sources, de la documentation, et d’un optionnel
3087 exécutable si jamais les sources ne peuvent pas être compilées.

3088 Nous avons ensuite sélectionné la version des outils à utiliser. Certains outils ont évolué depuis
3089 leur publication, soit en étant maintenus par leurs auteurs, soit suite à un branchement par un
3090 autre développeur. Nous avons décidé d’utiliser la dernière version stable en date de 2023 (date
3091 de l’étude). Le seul cas de branchement intéressant que nous avons trouvé est celui d’IC3, que
3092 nous avons décidé d’inclure en plus de la version originale. Le Table 16 résume cette étape.

3093 Nous avons ensuite exécuté ces outils sur deux jeux d’applications: Drebin, un jeu de malicieux
3094 connus pour être vieux et biaisé, et RASTA, un jeu que nous avons échantillonné nous-mêmes
3095 pour représenter l’évolution des caractéristiques des applications entre 2010 et 2023, d’un total
3096 de 62 525 APKs.

3097 Après avoir lancé les outils, nous avons collecté les différents résultats et traces d’exécution.
3098 Figure 21 et Figure 22 montrent les résultats des analyses sur les applications de Drebin et
3099 RASTA. L’analyse est considérée comme réussie (vert) si un résultat est obtenu, sinon elle
3100 a échoué (rouge). Quand l’analyse met plus d’une heure à finir, elle est avortée (bleue). On
3101 peut voir que les outils ont d’assez bon résultats sur Drebin, avec 11 outils qui ont un taux
3102 de terminaison au dessus de 85%. Sur RASTA par contre, 12 outils (54.55%) ont un taux de
3103 terminaison en dessous de 50%. Les outils qui avaient des difficultés avec Drebin ont aussi de

Outil	Disponibilité			Dépo type	Décision	Commentaires
	Bin	Src	Doc			
A3E	–	✓	✓	github	✗	Outil hybride (statique et dynamique)
A5	–	✓	✗	github	✗	Outil hybride (statique et dynamique)
Adagio	–	✓	✓	github	✓	
Amandroid	✓	✓	✓	github	✓	
Anadroid	✗	✓	✓	github	✓	
Androguard	–	✓	✓	github	✓	
Android-app-analysis	✗	✓	✓	google	✗	Outil hybride (statique et dynamique)
Apparecium	✓	✓	✗	github	✓	
BlueSeal	✗	✓	○	github	✓	
Choi <i>et al.</i>	✗	✓	○	github	✗	Nécessite les fichiers sources
DIALDroid	✓	✓	✓	github	✓	
DidFail	✓	✓	○	bitbucket	✓	
DroidSafe	✗	✓	✓	github	✓	
Flowdroid	✓	✓	✓	github	✓	
Gator	✗	✓	✓	edu	✓	
IC3	✓	✓	○	github	✓	
IccTA	✓	✓	✓	github	✓	
Lotrack	✗	✓	✗	github	○	Auteurs reconnaissent documentation insuffisante.
MalloDroid	–	✓	✓	github	✓	
PerfChecker	✗	✗	○	request	✓	Binaire obtenu des auteurs
Poeplau <i>et al.</i>	✗	○	✗	github	✗	Dédié à la sécurisation d'Android
Redexer	✗	✓	✓	github	✓	
SAAF	✓	✓	✓	github	✓	
StaDynA	✗	✓	✓	request	✗	Outil hybride (statique et dynamique)
Thresher	✗	✓	✓	github	○	Ne compile pas même avec l'aide des auteurs
Wognsen <i>et al.</i>	–	✓	✗	bitbucket	✓	

binaires, sources: –: non pertinent, ✓: disponible, ○: partiellement disponible, ✗: non fourni
documentation: ✓✓: excellente, MWE, ✓: quelques incohérences, ○: mauvaise qualité, ✗: non disponible

décision: ✓: considéré; ○: considéré mais pas compilé; ✗: sort du cadre de l'étude

Table 15: Outils considérés: disponibilité et réutilisabilité

3104 mauvais résultats sur RASTA, mais d'autres outils avec des résultats acceptables sur Drebin
 3105 voient leur résultats chuter avec RASTA.

3106 Ces résultats nous permettent de répondre à notre première question **QR1**:

3107 Sur un jeu d'applications récentes nous considérons que 54.55% des outils sont utilisables. De
 3108 plus pour les outils que nous avons pu lancer, 54.9% des analyses ont bien terminé correctement.

Outil	Original		Branchement Vivant		Date Dernière Modification	Auteurs Contactés	Environnement Langage – SE
	Etoiles	Vivant	Nb	Utilisable			
Adagio	74	✓	0	✗	2022-11-17	✓	Python – U20.04
Amandroid	161	✗	2	✗	2021-11-10	✓	Scala – U22.04
Anadroid	10	✗	0	✗	2014-06-18	✗	Scala/Java/Python – U22.04
Androguard	4430	✓	3	✗	2023-02-01	✗	Python – Python 3.11 slim
Apparecium	0	✗	1	✗	2014-11-07	✗	Python – U22.04
BlueSeal	0	✗	0	✗	2018-07-04	✓	Java – U14.04
DIALDroid	16	✗	1	✗	2018-04-17	✗	Java – U18.04
DidFail	4	✗			2015-06-17	✓	Java/Python – U12.04
DroidSafe	92	✗	3	✗	2017-04-17	✓	Java/Python – U14.04
Flowdroid	868	✓	1	✗	2023-05-07	✓	Java – U22.04
Gator					2019-09-09	✓	Java/Python – U22.04
IC3	32	✗	3	✓	2022-12-06	✗	Java – U12.04 / 22.04
IccTA	83	✗	0	✗	2016-02-21	✓	Java – U22.04
Lotrack	5	✗	2	✗	2017-05-11	✓	Java – ?
MalloDroid	64	✗	10	✗	2013-12-30	✗	Python – U16.04
PerfChecker		✗			–	✓	Java – U14.04
Redexer	153	✗	0	✗	2021-05-20	✓	Ocaml/Ruby – U22.04
SAAF	35	✗	5	✗	2015-09-01	✓	Java – U14.04
Thresher	31	✗	1	✗	2014-10-25	✓	Java – U14.04
Wognsen <i>et al.</i>				✗	2022-06-27	✗	Python/Prolog – U22.04

✓: oui, ✗: non, UX.04: Ubuntu X.04

Table 16: Outils sélectionnés, branchements, versions sélectionnées et environnements d'exécution

- 3109 Nous avons ensuite étudié l'évolution du taux de terminaison des outils au cours des ans. La
3110 Figure 23 montre cette évolution pour les outils codés en Java. On peut noter une tendance
3111 générale où le taux de terminaison diminue avec le temps.
- 3112 Plusieurs facteurs peuvent être responsables. Par exemple, la librairie standard d'Android et le
3113 format des applications ont évolué avec les versions d'Android. Un autre changement notable
3114 est la taille du code à octets des applications. Les applications les plus récentes ont notablement
3115 plus de code.
- 3116 Pour déterminer le facteur qui influence le plus le taux de terminaison, nous avons étudié des
3117 sous-ensembles de RASTA avec certains de ces paramètres fixés. Par exemple, nous avons tracé




Figure 21: Taux de terminaison pour le jeu d'applications Drebin

3118 l'évolution du taux de terminaison en fonction de l'année de publication des applications sur
 3119 l'ensemble des applications dont le code à octets fait entre 4.08 et 5.2 Mo.

3120 Nous en avons conclu la réponse à notre question de recherche **QR2**: Au cours du temps, le taux
 3121 de terminaison des outils diminue, allant de 78% à 61% cinq ans plus tard, à 45% dix ans plus
 3122 tard. Ce taux varie en fonction de la taille du code à octet, et, dans de moindre mesure, la
 3123 version d'Android.




Figure 22: Taux de terminaison pour le jeu d'application RASTA




Figure 23: Taux de terminaison des outils basé sur Java au cours des ans

3124 Pour répondre à notre dernière recherche question, nous avons comparé le taux de terminaison
 3125 entre les applications bénignes et les malicieux. Les résultats semblent indiquer que les malicieux
 3126 sont plus facilement analysés. Pour vérifier cette affirmation, nous avons comparé le taux de
 3127 terminaison mais aussi la taille du code à octets des applications et effectivement, il semblerait
 3128 que ce résultat soit vrai, y compris à taille égale.

3129 Nous avons donc une réponse à notre **RQ3**: Les malicieux causent moins d'erreurs lors de leur
 3130 analyse par des outils d'analyse statique.

3131 Finalement, nous avons une réponse à notre première problématique:

3132 Plus de la moitié des outils sélectionnés ne sont plus utilisables. Dans certains cas, cela est
 3133 dû à notre incapacité à les installer correctement, mais majoritairement, cela est dû au faible
 3134 taux de terminaison des outils lors de l'analyse des applications. Nos résultats montrent que
 3135 les applications avec beaucoup de code sont plus difficiles à analyser, et, en moindre mesure, la
 3136 version d'Android ainsi que la malignité de l'application peut avoir un impact.

3137 **B.3 Chargeurs de classes au milieu: Dérouter les analyseur statiques** 3138 **pour Android**

3139 Dans ce chapitre, nous étudions comment Android gère le chargement de classe en présence
 3140 de multiples versions de la même classe. Nous modélisons l'algorithme de chargement de classe
 3141 d'Android, et l'utilisons comme base pour une nouvelle famille de brouillage de code que nous

3142 appelons *masquage de classes*. Nous analysons ensuite des applications publiés en 2023 pour
3143 déterminer si cette technique de brouillage est actuellement utilisée.

3144 Le chargement de classe est une fonctionnalité de Java dont Android a hérité. Les développeurs
3145 interagissent avec elle le plus souvent au travers de classes héritant de `ClassLoader` pour charger
3146 dynamiquement du code. Toutefois, elle a un rôle bien plus général. À chaque fois qu'Android
3147 rencontre une nouvelle classe en exécutant une méthode, il va charger cette classe au travers
3148 du mécanisme de chargement de classe.

3149 L'intérêt de ce mécanisme est qu'il permet d'utiliser des classes provenant de différentes
3150 sources. Pour cela, Android associe à chaque classe un objet `ClassLoader`, celui qui a été
3151 utilisé pour charger cette classe. Par la suite, Android utilise ce `ClassLoader` pour charger
3152 toutes les classes référencées par cette première classe. Pour permettre aux classes provenant
3153 de différents `ClassLoader` d'interagir entre elles, les `ClassLoader` implémentent un mécanisme
3154 de délégation. Chaque `ClassLoader` a un "parent", un autre objet de type `ClassLoader`, auquel
3155 le `ClassLoader` va déléguer le chargement de classe. Si la classe n'est pas trouvée par le parent,
3156 alors le `ClassLoader` va la charger lui-même. Bien que ce système de délégation est utilisé par
3157 toutes les classes héritant de `ClassLoader` dans la librairie standard d'Android (à l'exception
3158 de `DelegateLastClassLoader` qui délègue dans un ordre légèrement différent), ce comportement
3159 est spécifié par l'implémentation de chaque classe `ClassLoader`. Une application peut très bien
3160 définir une nouvelle classe héritant de `ClassLoader` qui n'implémente pas ce processus. Toutefois,
3161 ce cas relève de l'analyse dynamique: un `ClassLoader` défini dans l'application ne peut pas être
3162 utilisé par Android sans exécuter du code de l'application pour l'instancier. Dans ce chapitre,
3163 nous nous concentrons sur le comportement par défaut d'Android, aussi nous n'avons besoin
3164 d'analyser que les `ClassLoader` instanciés par Android lui-même pour lancer l'application.

3165 Le premier `ClassLoader` utilisé par Android est `BootClassLoader`. Cette classe est une classe
3166 singleton, ce qui signifie qu'il ne peut y avoir qu'une seule instance de la classe par application.
3167 Elle est utilisée pour charger les classes de plateforme. Ces classes sont les classes implémentées
3168 par Android et qui peuvent être utilisées par une application sans qu'elles ne soient présentes
3169 dans l'application. Elles peuvent être séparées en deux catégories, les classes du KDL (Kit
3170 de Développement Logiciel) Android, et les classes de l'IPA (Interface de Programmation
3171 d'Application) cachée. Les premières forment la librairie standard d'Android. Elles sont docu-
3172 mentées et couramment utilisées par les développeurs. Les secondes sont des classes utilisées par
3173 Android en interne, mais que les applications ne sont pas supposées utiliser. Elles ne sont pas
3174 documentées, et depuis quelques années Android commence à faire des efforts pour empêcher
3175 les développeurs de les utiliser. Elles sont toutefois encore utilisées, et, au moins jusqu'à présent,
3176 les mesures d'Android ne suffisent pas à les rendre inaccessibles.

3177 Ce `BootClassLoader` est utilisé comme le parent par défaut par tous les `ClassLoader` définis
3178 dans les classes plateforme d'Android. Quand le parent d'un `ClassLoader` n'est pas défini
3179 (quand sa valeur est nulle), les `ClassLoader` vont déléguer le chargement au `BootClassLoader`
3180 à la place. L'autre type de `ClassLoader` utilisé par Android par défaut est le `PathClassLoader`.
3181 Cette classe est utilisée pour charger des classes stockées dans des fichiers. Android en définit
3182 deux par défaut, un `PathClassLoader` "système", et un `PathClassLoader` pour l'application. La
3183 documentation indique que le chargeur "système" est le chargeur par défaut pour le processus
3184 principal. Toutefois, il ne semble pas être utilisé en pratique. Le chargeur de l'application en
3185 revanche est utilisé pour les classes contenues dans l'application, c'est donc le chargeur utilisé
3186 par défaut pour toutes les classes codées par le développeur.

3187 En plus des chargeurs de classes, il y a un dernier critère à considérer. Les fichiers DEX contenant
3188 le code à octets des applications ont une limite du nombre de méthodes qui peuvent être
3189 référencées. Pour y remédier, Android a introduit un nouveau format d'application contenant
3190 plusieurs fichiers DEX. Pour notre étude, le point notable de ces applications est que bien
3191 qu'Android teste qu'un fichier DEX ne contient qu'une seule implémentation de chaque classe,
3192 ce test n'est fait que fichier par fichier: deux fichiers DEX peuvent contenir une implémentation
3193 d'une même classe chacun. Les fichiers DEX de ces applications "multi-dex" sont nommés
3194 `classes.dex`, puis `classesX.dex` où X est un entier supérieur ou égale à 2. Pour savoir quelle
3195 implémentation est utilisée par Android, il faut donc savoir dans quel ordre les fichiers sont
3196 visités par les `PathClassLoader`.

3197 Finalement, après avoir étudié le code source d'Android, nous concluons que l'algorithme utilisé
3198 est le même que celui que nous avons décrit dans le pseudo code Code 15. Cet algorithme a
3199 deux points notables. En premier lieu, les classes plateformes ont toujours la priorité sur les
3200 autres classes. Cela peut être intuitif pour les classes courantes comme `String`, mais il faut
3201 se rappeler que les classes de l'IPA cachée ne sont pas documentées. Ensuite, les classes sont
3202 sélectionnées parmi les fichiers DEX dans un ordre non trivial, et s'arrête à la première
3203 implémentation trouvée. Le premier fichier testé est `classes.dex`, suivi de `classes2.dex`, puis
3204 `classes3.dex` et ainsi de suite, jusqu'à ce qu'un fichier `classesX.dex` n'existe pas. La limite
3205 au nombre de fichiers DEX est très élevée (2^{64} sur les téléphones actuels), tant que le fichier
3206 suivant existe et que la classe n'est pas trouvée, Android va continuer. Aussi, le code contenu
3207 fichier `classes100.dex` peut être utilisé par Android, ou non, par exemple si `classes99.dex`
3208 n'existe pas. Plus surprenant, le code contenu dans un fichier `classes1.dex` ou `classes02.dex`
3209 ne sera pas utilisé. Lors de l'analyse statique d'applications, ces deux points peuvent mener à
3210 des complications que nous allons maintenant explorer.

3211 A partir de cet algorithme, nous avons mis au point plusieurs méthodes de brouillage de code que
3212 nous appelons *masquage de classe*: la classe utilisée est masquée par une autre implémentation
3213 fournie par le développeur. Nous nous concentrons sur l'obfuscation statique, mais cette stratégie

```

1 def obtenir_multi_dex_classes_nom_dex(indice: int):
2     if indice == 0:
3         return "classes.dex"
4     else:
5         return f"classes{indice+1}.dex"
6
7 def charge_classe(nom_classe: str):
8     if est_class_plateforme(nom_classe):
9         return charge_depuis_chargeur_class_boot(nom_classe)
10    else:
11        indice = 0
12        fichier_dex = obtenir_multi_dex_classes_nom_dex(indice)
13        while fichier_existe_dans_apk(fichier_dex) and \
14            not classe_non_trouvee_dans_fichier_dex(nom_classe, fichier_dex):
15            indice += 1
16        if fichier_existe_dans_apk(fichier_dex):
17            return charge_depuis_fichier(fichier_dex, nom_classe)
18        else:
19            raise ErreurClasseNonTrouvee()

```

Code 15: Algorithme de chargement de classe par défaut pour les applications Android

3214 peut être étendue à une approche dynamiquement en utilisant différents chargeurs de classes.
 3215 Nous proposons trois techniques dans cette catégorie:

3216 **Auto masquage** Ici, le développeur utilise le format multi-dex pour mettre plusieurs implé-
 3217 mentations différentes dans la même application. L'objectif est d'exploiter les divergences
 3218 entre l'algorithme de chargement de classes d'Android et la façon dont les outils d'analyse
 3219 sélectionnent l'implémentation à utiliser. De cette façon, la classe utilisée par Android ne
 3220 sera pas celle analysée.

3221 **Masquage de KDL** Cette fois, le développeur inclut une implémentation pour une classe
 3222 du KDL dans l'application. Un outil qui ne priorise pas les classes plateformes, ou ne les
 3223 connaît pas, va alors utiliser une implémentation invalide de la classe pour son analyse.

3224 **Masquage d'IPA Cachée** L'idée est la même que pour la technique précédente, mais cette
 3225 fois pour une classe de l'IPA cachée. Nous distinguons masquage de KDL et masquage
 3226 d'IPA caché car les IPA cachés n'étant pas documentés, il est possible que des outils soient
 3227 capables de résoudre la première technique mais pas la deuxième.

3228 Nous avons vérifié l'effet de ses techniques sur 4 outils de rétro-ingénierie Android courants:
 3229 Jadx, Apktool, Androguard et Flowdroid. Le Table 17 résume nos conclusions. Jadx est un
 3230 décompilateur d'application. Lorsqu'il est utilisé pour décompiler une application usant d'auto-
 3231 masquage, il va sélectionner la mauvaise classe, mais indiquer en commentaire la liste des fichiers

3232 de code à octet contenant une implementation de la classe. Apktool et Androguard listent toutes
 3233 les classes de l'application, il revient donc à l'analyste de choisir la bonne implémentation, et
 3234 pour les analyses plus poussées d'Androguard, Androguard choisit la mauvaise classe. Aucun
 3235 de ces trois outils n'indiquent en aucune façon qu'une classe est déjà définie dans le KDL ou
 3236 les IPA cachés. Flowdroid en revanche est capable de détecter les flux de données passant par
 3237 des classes du KDL, y compris en présence d'une réimplémentation dans l'application. Ce n'est
 3238 par contre pas le cas pour les classes d'IPA cachées. Il est intéressant de noter que Soot, la
 3239 librairie sur laquelle est basé Flowdroid, a bien un algorithme qui priorise les fichiers DEX, et
 3240 que cet algorithme est très proche de celui d'Android. Toutefois, les fichiers commençant par
 3241 `classes` sont ensuite priorisés par ordre alphabétique. Cela signifie que les classes contenues
 3242 dans `classes0.dex`, `classes02.dex` ou `classes10.dex` sont priorisées sur celles de `classes2.dex`.
 3243 Ce problème est hérité par Flowdroid, ce qui le rend sensible à la technique d'auto-masquage.
 3244 Pour savoir si ces techniques sont utilisées dans la nature, nous avons scanné 49 975 applications
 3245 publiées entre janvier 2023 et 2024. Pour vérifier que les différentes implémentations sont bien
 3246 distinctes, nous comparons la représentation smali (le langage assembleur pour le format Dalvik)
 3247 du code à octets des méthodes. La Table 18 résume ces résultats. Il est notable qu'un nombre
 3248 important d'applications (23.52%) ont au moins un cas de masquage. En étudiant en détail, nous
 3249 avons noté que la majorité des classes concernées sont des classes introduites entre la version
 3250 minimale et la version cible d'Android pour l'application. Cela laisse entendre que ces classes
 3251 ont été rajoutées pour permettre à l'application de fonctionner avec les versions d'Android où
 3252 ces classes n'existent pas. Le taux élevé de code identique pour les cas d'auto-masquage semble
 3253 également pointer vers des erreurs lors de la compilation de l'application. De plus, l'analyse

Outil	Version	Masquage		
		Auto	KDL	Cachée
Jadx	1.5.0	○	●	●
Apktool	2.9.3	○	●	●
Androguard	4.1.2	○	●	●
Flowdroid	2.13.0	●	×	●

●: Le masquage fonctionne

○: Le masquage fonctionne, mais un avertissement est émis ou les différentes implémentations sont visibles

×: Le masquage ne marche pas

Table 17: Résultats des techniques de masquage contre des outils d'analyse statique

	Nombre d'app			Classes Masquées	Moyenne Médiane	KDL	Min KDL	Code Identique
	%	% malicieux						
Pour toutes les applications du jeu								
Auto	49 975	100.0%	0.53%	2.1	0	32.1	21.7	74.8%
KDL	49 975	100.0%	0.53%	6.5	0	32.1	21.7	8.04%
Cachées	49 975	100.0%	0.53%	0.5	0	32.1	21.7	17.42%
Totale	49 975	100.0%	0.53%	9	0	32.1	21.7	23.76%
Pour les applications avec au moins 1 cas de masquage								
Auto	234	0.47%	5.98%	438.1	18	31.4	22.4	74.8%
KDL	11 755	23.52%	0.38%	27.6	5	32.4	22	8.04%
Cachées	1556	3.11%	0.71%	16.1	1	32.1	22.2	17.42%
Totale	12 301	24.61%	0.42%	36.7	6	32.4	22	23.76%

Table 18: Classes masquées comparées aux classes platform d'Android IPA 34 pour un jeu de 49 975 applications

3254 manuelle des cas où le smali diffère montre que les différences viennent de détails lors de la
 3255 compilation (par exemple, l'inversion de deux registres, ce qui n'a aucun effet sur l'exécution
 3256 du code).

3257 Notre conclusion est que le masquage de classes n'est pas activement exploité pour le brouillage
 3258 de code. En revanche, cette situation se produit naturellement dans les applications. Il est donc
 3259 important pour les outils d'analyses de modéliser correctement le processus de chargement de
 3260 classes. Avec l'algorithme décrit par le Code 15, cela répond à notre seconde problématique.

3261 B.4 L'Application de Thésée: Même après avoir ajouté les informa- 3262 tions d'exécution, c'est toujours votre application

3263 Dans ce dernier chapitre, nous nous penchons sur la question de comment permettre aux outils
 3264 d'analyse statique d'accéder à des résultats dynamiques. Des contributions précédentes ont
 3265 encodé leur résultats dans l'application elle-même pour transmettre leurs résultats à d'autres
 3266 outils. Nous allons utiliser cette approche pour permettre à des outils d'analyse statique
 3267 d'analyser le comportement d'applications utilisant de la réflexion ou du chargement de code
 3268 dynamique.

3269 Premièrement, il nous faut définir la transformation que nous voulons effectuer. Concernant
 3270 la réflexion, il y a 3 méthodes permettant d'appeler des méthodes arbitraires dans Android.
 3271 `Class.newInstance()` et `Constructor.newInstance()` permettent d'instancier un nouvel objet
 3272 et d'appeler l'un des ses constructeurs, tandis que `Method.invoke()` permet d'appeler une
 3273 méthode. Les objets `Class`, `Constructor` ou `Method` utilisés pour appeler ces méthodes peuvent
 3274 être obtenus de multiples façons. Nous n'allons donc pas chercher à modifier le code obtenant ces
 3275 objets. À la place, nous allons nous concentrer sur l'appel des méthodes. À différents moments de

3276 l'exécution, une même instruction appelant `Method.invoke()` peut appeler différentes méthodes.
3277 De plus, la collection des informations de réflexion ne sera jamais parfaite: il y a des situations
3278 où on ne peut jamais être certain d'avoir la liste complète des méthodes appelées. Par exemple,
3279 on peut imaginer une application qui appelle par réflexion une méthode dont le nom est obtenu
3280 depuis un serveur distant. Dans ce cas, sans accès au code du serveur il est impossible d'avoir
3281 la liste exhaustive des méthodes qui peuvent être utilisées. Pour prendre en compte ces deux
3282 cas, nous allons remplacer les appels par des blocs conditionnels. Pour chaque méthode dont
3283 on sait qu'elle peut être appelée, nous testons si l'objet `Method` correspond à cette méthode
3284 en comparant son nom et sa signature. Si c'est le cas, la méthode est appelée dans le bloc avec
3285 l'instruction Dalvik appropriée. Si la méthode ne correspond à aucune méthode connue, alors
3286 l'appel est fait par réflexion. Ainsi, le comportement de l'application est conservé.

3287 Pour le chargement de code dynamique, nous avons conclu qu'il n'est pas nécessaire de modifier
3288 le code. A la place, nous pouvons directement ajouter à l'application le fichier DEX en utilisant
3289 le format multi-dex. Toutefois, si jamais certaines classes contenues dans le code que nous
3290 injectons sont déjà présentes dans l'application, nous renommons la classe et ses références
3291 de sorte à reproduire statiquement le comportement de l'algorithme de chargement de classe
3292 d'Android. Cette approche a des limites, en particulier lors d'appel réflexifs, car il n'existe pas
3293 de solutions pour comparer les chargeurs de classes avec une valeur statique. Si un même site
3294 d'appel réflexif appelle deux méthodes avec des signatures et noms identiques mais associées
3295 à des classes provenant de chargeurs de classes différents, nous ne sommes pas en mesure de
3296 reproduire exactement le même comportement statiquement. Toutefois, les appels aux deux
3297 différentes méthodes apparaissent bien dans la nouvelle application, ce qui devrait permettre
3298 aux outils d'analyse statique de considérer les deux cas possibles.

3299 Pour pouvoir effectuer ces transformations, il nous faut certaines informations. Les noms,
3300 signature et chargeur de classes des méthodes appelées par réflexion, ainsi que la position
3301 exacte du site de l'appel réflexif, et le code à octets chargé dynamiquement. Pour obtenir ces
3302 informations, nous utilisons Frida, un outil permettant d'injecter des scripts dans les méthodes
3303 appelées pendant l'exécution d'une application Android. Pour la réflexion, nous avons bien
3304 entendu instrumenté `Class.newInstance()`, `Cstructor.newInstance()` et `Method.invoke()`.
3305 Pour le chargement de code, le choix est un peu moins évident car il existe de multiples façons de
3306 charger du code à octets. Nous avons finalement choisi `DexFile.openInMemoryDexFileNative()`
3307 et `DexFile.openDexFileNative()`, des méthodes de l'IPA cachée. Ces méthodes sont les dernières
3308 méthodes appelées dans l'environnement Java avant de passer en native pour analyser et charger
3309 le code à octets. Pour aider à l'exploration des applications, nous avons réutilisé une partie de
3310 GroddDroid, un outil dédié à l'exploration dynamique d'applications.

3311 Nous avons lancé notre analyse statique sur les 5000 applications publiées en 2023 du jeu
3312 d'applications RASTA. Malheureusement, les résultats semblent indiquer que notre environ-

3313 nement d'exécution est insuffisant et que beaucoup d'applications n'ont pas été visitées
 3314 correctement. Malgré tout, nous avons collecté 640 fichiers de code à octets. Toutefois, une fois
 3315 comparé, nous remarquons que parmi ces fichiers, il n'y a que 70 fichiers distincts. L'inspection
 3316 du contenu montre que ces fichiers sont principalement des bibliothèques de code publicitaire
 3317 ou analytiques. Seuls 13 fichiers parmi les 640 collectés ne proviennent ni de Google, ni de
 3318 Facebook, ni de AppsFlyer. Ces fichiers restants contiennent du code spécifique aux applications
 3319 les utilisant, principalement des applications exigeant un niveau important de sécurité comme
 3320 des applications bancaires ou d'assurance santé.

3321 La Table 19 montre le nombre d'arcs du graphe d'appel de fonction de ces quelques applications
 3322 qui chargent du code spécifique à leur usage dynamiquement. La colonne "Réflexion Ajoutées"
 3323 correspond au nombre d'appels réflexifs ajoutés à l'application. Les autres arcs ajoutés sont soit
 3324 des fonctions "colle" que nous avons ajoutées à l'application pour choisir la bonne méthode à
 3325 appeler réflexivement, soit des méthodes appelées par du code chargé dynamiquement auquel
 3326 Androguard n'avait pas accès avant l'instrumentation. On peut voir que notre méthode permet
 3327 à Androguard de calculer un plus grand graphe.

3328 Nous avons ensuite modifié les applications comme décrit précédemment, puis relancé les outils
 3329 de notre première contribution sur les applications modifiées pour comparer leur taux de
 3330 terminaison au taux sur les applications initiales. En fonction des outils, le taux de terminaison
 3331 est soit inchangé, soit légèrement plus faible pour les applications modifiées.

3332 Pour vérifier que notre approche fonctionne, nous avons créé une petite application de test
 3333 utilisant du chargement dynamique et des appels réflexifs. Code 16 montre la classe principale
 3334 de l'application. On peut voir par exemple que l'application utilise des chaînes de caractères
 3335 chiffrées pour stocker le nom des méthodes à appeler.

APK SHA 256	Nombre d'arcs du graphe d'appel			
	Avant	Après	Différences	Réflexion Ajoutées
0019d7fb6a...	641	60 170	59 529	1
274b677449...	537 613	540 674	3061	26
34599c2499...	336 740	339 616	2876	29
35065c6834...	343 245	346 694	3449	26
e7b2fb02ff...	464 642	465 389	747	91
efececc03c...	243 647	243 925	278	23
f34ce1e7a8...	704 095	706 576	2481	28

Table 19: Arcs ajoutés aux graphes d'appel de fonctions des applications instrumentées, calculé par Androguard

```

1 package com.example.theseus;
2
3 public class Main {
4     private static final String DEX = "ZGV4CjA [...] EAAABEAwAA";
5     Activity ac;
6     private Key key = new SecretKeySpec("__Secret Key__".getBytes(),
7 "AES");
8     ClassLoader cl = new
9 InMemoryDexClassLoader(ByteBuffer.wrap(Base64.decode(DEX, 2)),
10 Main.class.getClassLoader());
11
12     public void main() throws Exception {
13         String[] strArr = {"n6WGYJzjDrUvR9cYljLnlw==", "dapES0wL/
14 iFIPuMnH3fh7g=="};
15         Class<?> loadClass =
16 this.cl.loadClass(decrypt("W5f3xRf3wCSYcYG7ckYGR5xuuESDZ2NcDUzGxsq3sIs="));
17         Object obj = "imei";
18         for (int i = 0; i < 2; i++) {
19             obj = loadClass.getMethod(decrypt(strArr[i]),
20 String.class, Activity.class).invoke(null, obj, this.ac);
21         }
22     }
23     public String decrypt(String str) throws Exception {
24         Cipher cipher = Cipher.getInstance("AES/ECB/PKCS5Padding");
25         cipher.init(2, this.key);
26         return new String(cipher.doFinal(Base64.decode(str, 2)));
27     }
28     ...
29 }

```

Code 16: Code de la classe principale de l'application calculé par Jadx, avant modification

3336 Après avoir lancé notre analyse statique et instrumenté l'application pour y ajouter les
3337 informations dynamique, Jadx montre maintenant le Code 17. On peut voir que les méthodes
3338 `Malicious.get_data()` et `Malicious.send_data()` sont appelées. De plus, la classe `Malicious`
3339 qui n'était pas présente dans l'application originale est maintenant visible dans l'application
3340 modifiée.

3341 Dans le code de `Malicious`, `get_data()` retourne des données d'une source d'information sensible,
3342 et `send_data()` exfiltre les données qui lui sont passées. Une fuite d'information devrait donc être
3343 détectée par Flowdroid. Lancé sur l'application originale, Flowdroid calcule un graphe d'appel

```
1 public void main() throws Exception {
2     String[] strArr = {"n6WGYJzjDrUvR9cYLjlnlw==", "dapES0wL/
iFIPuMnH3fh7g=="};
3     Class<?> loadClass =
this.cl.loadClass(decrypt("W5f3xRf3wCSYcYG7ckYGR5xuuESDZ2NcDUzGxsq3sIs="));
4     Object obj = "imei";
5     for (int i = 0; i < 2; i++) {
6         Method method = loadClass.getMethod(decrypt(strArr[i]),
String.class, Activity.class);
7         Object[] objArr = {obj, this.ac};
8         obj = T.check_is_Malicious_get_data_fe2fa96eab371e46(method) ?
9             Malicious.get_data((String) objArr[0], (Activity)
objArr[1]) :
10            T.check_is_Malicious_send_data_ca50fd7916476073(method) ?
11            Malicious.send_data((String) objArr[0], (Activity)
objArr[1]) :
12            method.invoke(null, objArr);
13     }
14 }
```

Code 17: Code de Main.main() calculé par Jadx, après modifications

3344 de méthodes contenant 43 arcs, et ne détecte aucune fuite. En revanche, lancé sur l'application
3345 modifiée, Flowdroid calcule cette fois un graphe de 76 arcs, et détecte bien la fuite de données.
3346 Bien que nous n'ayons pas pu tester notre approche correctement à grande échelle dû aux
3347 limites de notre environnement d'analyse statique, nous avons bien montré qu'il est possible de
3348 transmettre des informations dynamiques à des outils d'analyse statique pour améliorer leurs
3349 résultats.

-
- 3351 [1] A. Abraham, R. Andriatsimandefitra, A. Brunelat, J.-F. Lalande, and V. Viet Triem Tong.
3352 2015. GroddDroid: a gorilla for triggering malicious behaviors. In *2015 10th International*
3353 *Conference on Malicious and Unwanted Software (MALWARE)*, October 2015. 119–127.
3354 <https://doi.org/10.1109/MALWARE.2015.7413692>
- 3355 [2] Boladji Vinny Adjibi, Fatou Ndiaye Mbodji, Tegawendé F. Bissyandé, Kevin Allix, and
3356 Jacques Klein. 2022. The Devil is in the Details: Unwrapping the Cryptojacking Malware
3357 Ecosystem on Android. In *2022 IEEE 22nd International Working Conference on Source*
3358 *Code Analysis and Manipulation (SCAM)*, October 2022. 153–163. [https://doi.org/10.](https://doi.org/10.1109/SCAM55253.2022.00023)
3359 [1109/SCAM55253.2022.00023](https://doi.org/10.1109/SCAM55253.2022.00023)
- 3360 [3] Kevin Allix, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon. 2016. AndroZoo:
3361 Collecting Millions of Android Apps for the Research Community. In *13th Working*
3362 *Conference on Mining Software Repositories (MSR)*, May 2016. 468–471.
- 3363 [4] Radoniaina Andriatsimandefitra and Valérie Viet Triem Tong. 2015. Detection and Identi-
3364 fication of Android Malware Based on Information Flow Monitoring. In *2015 IEEE 2nd*
3365 *International Conference on Cyber Security and Cloud Computing*, November 2015. 200–
3366 203. <https://doi.org/10.1109/CSCloud.2015.27>
- 3367 [5] Radoniaina Andriatsimandefitra, Stéphane Geller, and Valérie Viet Triem Tong. 2012.
3368 Designing information flow policies for Android's operating system. In *IEEE International*
3369 *Conference on Communications*, June 2012. IEEE Computer Society, Ottawa, Canada,
3370 976–981. <https://doi.org/10.1109/ICC.2012.6364161>
- 3371 [6] Daniel Arp, Michael Spreitzenbarth, Hugo Gascon, Konrad Rieck, Germany Siemens, and
3372 Cert Munich. 2014. Drebin: Effective and Explainable Detection of Android Malware in
3373 Your Pocket. In *Network and Distributed System Security Symposium*, February 2014. The
3374 Internet Society, San Diego, California, USA.
- 3375 [7] Steven Arzt, Siegfried Rasthofer, and Eric Bodden. 2013. Instrumenting Android and Java
3376 Applications as Easy as abc. In *Fourth International Conference on Runtime Verification*,
3377 September 2013. Springer Berlin Heidelberg, Rennes, France, 364–381. [https://doi.org/10.](https://doi.org/10.1007/978-3-642-40787-1_26)
3378 [1007/978-3-642-40787-1_26](https://doi.org/10.1007/978-3-642-40787-1_26)
- 3379 [8] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques
3380 Klein, Yves Le Traon, Damien Octeau, and Patrick McDaniel. 2014. FlowDroid: Precise
3381 Context, Flow, Field, Object-sensitive and Lifecycle-aware Taint Analysis for Android

-
- 3382 Apps. In *ACM SIGPLAN Conference on Programming Language Design and Implemen-*
3383 *tation*, June 05, 2014. ACM Press, Edinburgh, UK, 259–269. [https://doi.org/10.1145/](https://doi.org/10.1145/2666356.2594299)
3384 [2666356.2594299](https://doi.org/10.1145/2666356.2594299)
- 3385 [9] Tanzirul Azim and Iulian Neamtiu. 2013. Targeted and Depth-First Exploration for Sys-
3386 *tematic Testing of Android Apps*. In *Proceedings of the 2013 ACM SIGPLAN International*
3387 *Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA*
3388 *2013, Part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, 2013. ACM,
3389 641–660. <https://doi.org/10.1145/2509136.2509549>
- 3390 [10] Mario Luca Bernardi, Marta Cimitile, Damiano Distanto, Fabio Martinelli, and Francesco
3391 *Mercaldo*. 2019. Dynamic malware detection and phylogeny analysis using process mining.
3392 *International Journal of Information Security* 18, 3 (June 2019), 257–284. [https://doi.org/](https://doi.org/10.1007/s10207-018-0415-3)
3393 [10.1007/s10207-018-0415-3](https://doi.org/10.1007/s10207-018-0415-3)
- 3394 [11] Sam Blackshear, Bor-Yuh Evan Chang, and Manu Sridharan. 2013. Thresher: Precise
3395 *Refutations for Heap Reachability*. *ACM SIGPLAN Notices* 48, 6 (June 2013), 275–286.
3396 <https://doi.org/10.1145/2499370.2462186>
- 3397 [12] Amiangshu Bosu, Fang Liu, Danfeng (Daphne) Yao, and Gang Wang. 2017. Collusive Data
3398 *Leak and More: Large-scale Threat Analysis of Inter-app Communications*. In *Proceedings*
3399 *of the 2017 ACM on Asia Conference on Computer and Communications Security*, April 02,
3400 2017. ACM, Abu Dhabi United Arab Emirates, 71–85. [https://doi.org/10.1145/3052973.](https://doi.org/10.1145/3052973.3053004)
3401 [3053004](https://doi.org/10.1145/3052973.3053004)
- 3402 [13] Kwanghoon Choi and Byeong-Mo Chang. 2014. A Type and Effect System for Activation
3403 *Flow of Components in Android Programs*. *Information Processing Letters* 114, 11 (2014),
3404 620–627. <https://doi.org/10.1016/j.ipl.2014.05.011>
- 3405 [14] Yuning Cui, Yi Sun, and Zhaowen Lin. 2023. DroidHook: a novel API-hook based Android
3406 *malware dynamic analysis sandbox*. *Automated Software Engineering* 30, 1 (February
3407 2023), 10. <https://doi.org/10.1007/s10515-023-00378-w>
- 3408 [15] Anthony Desnos and Geoffroy Gueguen. 2011. Android: From Reversing to Decompila-
3409 *tion*. *Black Hat Abu Dhabi (2011)*. Retrieved from [https://media.blackhat.com/bh-ad-11/](https://media.blackhat.com/bh-ad-11/Desnos/bh-ad-11-DesnosGueguen-Android-Reversing_to_Decompileation_WP.pdf)
3410 [Desnos/bh-ad-11-DesnosGueguen-Android-Reversing_to_Decompileation_WP.pdf](https://media.blackhat.com/bh-ad-11-DesnosGueguen-Android-Reversing_to_Decompileation_WP.pdf)
- 3411 [16] Yue Duan, Mu Zhang, Abhishek Vasisht Bhaskar, Heng Yin, Xiaorui Pan, Tongxin Li,
3412 *Xueqiang Wang, and Xiaofeng Wang*. 2018. Things You May Not Know About Android
3413 *(Un)Packers: A Systematic Study based on Whole-System Emulation*. In *24th Annual*
3414 *Network and Distributed System Security Symposium*, 2018.

- 3415 [17] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick
3416 McDaniel, and Anmol N. Sheth. 2010. TaintDroid: An Information-Flow Tracking System
3417 for Realtime Privacy Monitoring on Smartphones. In *9th USENIX Symposium on Operat-*
3418 *ing Systems Design and Implementation*, October 2010. USENIX Association, Vancouver,
3419 BC, Canada, 393–407.
- 3420 [18] Farnood Faghihi, Mohammad Zulkernine, and Steven Ding. 2022. CamoDroid: An Android
3421 application analysis environment resilient against sandbox evasion. *Journal of Systems*
3422 *Architecture* 125, (April 2022), 102452. <https://doi.org/10.1016/j.sysarc.2022.102452>
- 3423 [19] Sascha Fahl, Marian Harbach, Thomas Muders, Lars Baumgärtner, Bernd Freisleben, and
3424 Matthew Smith. 2012. Why Eve and Mallory Love Android: An Analysis of Android SSL
3425 (in)Security. In *Proceedings of the 2012 ACM Conference on Computer and Communica-*
3426 *tions Security*, October 16, 2012. ACM, Raleigh North Carolina USA, 50–61. [https://doi.](https://doi.org/10.1145/2382196.2382205)
3427 [org/10.1145/2382196.2382205](https://doi.org/10.1145/2382196.2382205)
- 3428 [20] Eva Galperin. 2020. The State of the Stalkerware. January 2020. USENIX Association,
3429 San Francisco, CA.
- 3430 [21] Hugo Gascon, Fabian Yamaguchi, Daniel Arp, and Konrad Rieck. 2013. Structural
3431 Detection of Android Malware Using Embedded Call Graphs. In *Proceedings of the 2013*
3432 *ACM Workshop on Artificial Intelligence and Security*, November 04, 2013. ACM, Berlin
3433 Germany, 45–54. <https://doi.org/10.1145/2517312.2517315>
- 3434 [22] Dimitris Geneiatakis, Igor Nai Fovino, Ioannis Kounelis, and Pasquale Stirparo. 2015. A
3435 Permission Verification Approach for Android Mobile Applications. *Computers & Security*
3436 49, (March 2015), 192–205. <https://doi.org/10.1016/j.cose.2014.10.005>
- 3437 [23] Li Gong. 1998. Secure Java class loading. *IEEE Internet Computing* 2, 6 (November 1998),
3438 56–61. <https://doi.org/10.1109/4236.735987>
- 3439 [24] Michael I. Gordon, Deokhwan Kim, Jeff H. Perkins, Limei Gilham, Nguyen Nguyen, and
3440 Martin C. Rinard. 2015. Information Flow Analysis of Android Applications in DroidSafe.
3441 In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San*
3442 *Diego, California, USA, February 8-11, 2015*, 2015. The Internet Society.
- 3443 [25] Yi He, Yacong Gu, Purui Su, Kun Sun, Yajin Zhou, Zhi Wang, and Qi Li. 2023. A
3444 Systematic Study of Android Non-SDK (Hidden) Service API Security. *IEEE Transactions*
3445 *on Dependable and Secure Computing* 20, 2 (March 2023), 1609–1623. [https://doi.org/10.](https://doi.org/10.1109/TDSC.2022.3160872)
3446 [1109/TDSC.2022.3160872](https://doi.org/10.1109/TDSC.2022.3160872)
- 3447 [26] Johannes Hoffmann, Martin Ussath, Thorsten Holz, and Michael Spreitzenbarth. 2013.
3448 Slicing Droids: Program Slicing for Smali Code. In *Proceedings of the 28th Annual ACM*

-
- 3449 *Symposium on Applied Computing (SAC '13)*, March 18, 2013. Association for Computing
3450 Machinery, New York, NY, USA, 1844–1851. <https://doi.org/10.1145/2480362.2480706>
- 3451 [27] Jinseong Jeon, Kristopher K. Micinski, Jeffrey A. Vaughan, Ari Fogel, Nikhilesh Reddy,
3452 Jeffrey S. Foster, and Todd Millstein. 2012. Dr. Android and Mr. Hide: Fine-Grained
3453 Permissions in Android Applications. In *Proceedings of the Second ACM Workshop on*
3454 *Security and Privacy in Smartphones and Mobile Devices*, October 19, 2012. ACM, Raleigh
3455 North Carolina USA, 3–14. <https://doi.org/10.1145/2381934.2381938>
- 3456 [28] William Klieber, Lori Flynn, Amar Bhosale, Limin Jia, and Lujo Bauer. 2014. Android
3457 Taint Flow Analysis for App Sets. In *Proceedings of the 3rd ACM SIGPLAN International*
3458 *Workshop on the State of the Art in Java Program Analysis*, June 12, 2014. ACM,
3459 Edinburgh United Kingdom, 1–6. <https://doi.org/10.1145/2614628.2614633>
- 3460 [29] Pavel Kriz and Filip Maly. 2015. Provisioning of application modules to Android devices.
3461 In *2015 25th International Conference Radioelektronika (RADIOELEKTRONIKA)*, April
3462 2015. 423–426. <https://doi.org/10.1109/RADIOELEK.2015.7129009>
- 3463 [30] Li Li, Alexandre Bartel, Tegawende F. Bissyande, Jacques Klein, Yves Le Traon, Steven
3464 Arzt, Siegfried Rasthofer, Eric Bodden, Damien Ocateau, and Patrick McDaniel. 2015.
3465 IccTA: Detecting Inter-Component Privacy Leaks in Android Apps. In *2015 IEEE/ACM*
3466 *37th IEEE International Conference on Software Engineering*, May 2015. IEEE, Florence,
3467 Italy, 280–291. <https://doi.org/10.1109/ICSE.2015.48>
- 3468 [31] Li Li, Alexandre Bartel, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon. 2015.
3469 ApkCombiner: Combining Multiple Android Apps to Support Inter-App Analysis. In *ICT*
3470 *Systems Security and Privacy Protection*, 2015. Springer International Publishing, Cham,
3471 513–527. https://doi.org/10.1007/978-3-319-18467-8_34
- 3472 [32] Li Li, Tegawendé F. Bissyandé, Damien Ocateau, and Jacques Klein. 2016. DroidRA:
3473 taming reflection to support whole-program analysis of Android apps. In *Proceedings of*
3474 *the 25th International Symposium on Software Testing and Analysis (ISSTA 2016)*, July
3475 2016. Association for Computing Machinery, New York, NY, USA, 318–329. [https://doi.](https://doi.org/10.1145/2931037.2931044)
3476 [org/10.1145/2931037.2931044](https://doi.org/10.1145/2931037.2931044)
- 3477 [33] Li Li, Tegawendé F. Bissyandé, Mike Papadakis, Siegfried Rasthofer, Alexandre Bartel,
3478 Damien Ocateau, Jacques Klein, and Yves Le Traon. 2017. Static Analysis of Android Apps:
3479 A Systematic Literature Review. *Information and Software Technology* 88, (2017), 67–95.
3480 <https://doi.org/10.1016/j.infsof.2017.04.001>
- 3481 [34] Li Li, Tegawendé F. Bissyandé, Yves Le Traon, and Jacques Klein. 2016. Accessing
3482 Inaccessible Android APIs: An Empirical Study. In *2016 IEEE International Conference*

- 3483 *on Software Maintenance and Evolution (ICSME)*, October 2016. 411–422. <https://doi.org/10.1109/ICSME.2016.35>
3484
- 3485 [35] Sheng Liang and Gilad Bracha. 1998. Dynamic class loading in the Java virtual machine.
3486 *SIGPLAN Not.* 33, 10 (October 1998), 36–44. <https://doi.org/10.1145/286942.286945>
- 3487 [36] Shuying Liang, Andrew W. Keep, Matthew Might, Steven Lyde, Thomas Gilray, Petey
3488 Aldous, and David Van Horn. 2013. Sound and Precise Malware Analysis for Android
3489 via Pushdown Reachability and Entry-Point Saturation. In *Proceedings of the Third*
3490 *ACM Workshop on Security and Privacy in Smartphones & Mobile Devices (SPSM '13)*,
3491 November 08, 2013. Association for Computing Machinery, New York, NY, USA, 21–32.
3492 <https://doi.org/10.1145/2516760.2516769>
- 3493 [37] Yibin Liao, Jiakuan Li, Bo Li, Guodong Zhu, Yue Yin, and Ruoyan Cai. 2016. Automated
3494 Detection and Classification for Packed Android Applications. In *International Conference*
3495 *on Mobile Services*, June 2016. IEEE, San Francisco, USA, 200–203.
- 3496 [38] Max Lillack, Christian Kästner, and Eric Bodden. 2014. Tracking Load-Time Configu-
3497 ration Options. In *Proceedings of the 29th ACM/IEEE International Conference on*
3498 *Automated Software Engineering (ASE '14)*, September 15, 2014. Association for Comput-
3499 ing Machinery, New York, NY, USA, 445–456. <https://doi.org/10.1145/2642937.2643001>
- 3500 [39] Yepang Liu, Chang Xu, and Shing-Chi Cheung. 2014. Characterizing and Detecting
3501 Performance Bugs for Smartphone Applications. In *Proceedings of the 36th International*
3502 *Conference on Software Engineering*, May 31, 2014. ACM, Hyderabad India, 1013–1024.
3503 <https://doi.org/10.1145/2568225.2568229>
- 3504 [40] Linghui Luo, Felix Pauck, Goran Piskachev, Manuel Benz, Ivan Pashchenko, Martin Mory,
3505 Eric Bodden, Ben Hermann, and Fabio Massacci. 2022. TaintBench: Automatic Real-World
3506 Malware Benchmarking of Android Taint Analyses. *Empirical Software Engineering* 27, 1
3507 (January 2022), 16. <https://doi.org/10.1007/s10664-021-10013-5>
- 3508 [41] Ke Mao, Mark Harman, and Yue Jia. 2016. Sapienz: multi-objective automated testing
3509 for Android applications. In *Proceedings of the 25th International Symposium on Software*
3510 *Testing and Analysis (ISSTA 2016)*, July 2016. Association for Computing Machinery,
3511 New York, NY, USA, 94–105. <https://doi.org/10.1145/2931037.2931054>
- 3512 [42] Noah Mauthe, Ulf Kargén, and Nahid Shahmehri. 2021. A Large-Scale Empirical Study of
3513 Android App Decompilation. In *2021 IEEE International Conference on Software Analysis,*
3514 *Evolution and Reengineering (SANER)*, March 2021. 400–410. <https://doi.org/10.1109/SANER50967.2021.00044>
3515

-
- 3516 [43] René Mayrhofer, Jeffrey Vander Stoep, Chad Brubaker, and Nick Kralevich. 2021. The
3517 Android Platform Security Model. *ACM Trans. Priv. Secur.* 24, 3 (April 2021), 19:1–19:35.
3518 <https://doi.org/10.1145/3448609>
- 3519 [44] Jean-Marie Mineau and Jean-Francois Lalande. 2024. Evaluating the Reusability of
3520 Android Static Analysis Tools. In *Reuse and Software Quality*, 2024. Springer Nature
3521 Switzerland, Cham, 153–170.
- 3522 [45] Jean-Marie Mineau and Jean-François Lalande. 2025. Class Loaders in the Middle: Con-
3523 fusing Android Static Analyzers. *Digital Threats* 6, 3 (September 2025). [https://doi.org/](https://doi.org/10.1145/3754457)
3524 [10.1145/3754457](https://doi.org/10.1145/3754457)
- 3525 [46] Zhenyu Ning and Fengwei Zhang. 2018. DexLego: Reassembleable bytecode extraction for
3526 aiding static analysis. *Proceedings - 48th Annual IEEE/IFIP International Conference on*
3527 *Dependable Systems and Networks, DSN 2018* (2018), 690–701. [https://doi.org/10.1109/](https://doi.org/10.1109/DSN.2018.00075)
3528 [DSN.2018.00075](https://doi.org/10.1109/DSN.2018.00075)
- 3529 [47] Damien Ocateau, Daniel Luchaup, Matthew Dering, Somesh Jha, and Patrick McDaniel.
3530 2015. Composite Constant Propagation: Application to Android Inter-Component Com-
3531 munication Analysis. In *2015 IEEE/ACM 37th IEEE International Conference on Software*
3532 *Engineering*, May 2015. IEEE, Florence, Italy, 77–88. [https://doi.org/10.1109/ICSE.2015.](https://doi.org/10.1109/ICSE.2015.30)
3533 [30](https://doi.org/10.1109/ICSE.2015.30)
- 3534 [48] Damien Ocateau, Patrick McDaniel, Somesh Jha, Alexandre Bartel, Eric Bodden, Jacques
3535 Klein, and Yves Le Traon. 2013. Effective Inter-Component communication mapping in
3536 android: An essential step towards holistic security analysis. In *22nd USENIX Security*
3537 *Symposium (USENIX Security 13)*, 2013. 543–558.
- 3538 [49] Felix Pauck, Eric Bodden, and Heike Wehrheim. 2018. Do Android Taint Analysis Tools
3539 Keep Their Promises?. In *Proceedings of the 2018 26th ACM Joint Meeting on European*
3540 *Software Engineering Conference and Symposium on the Foundations of Software Engi-*
3541 *neering*, October 26, 2018. ACM, Lake Buena Vista FL USA, 331–341. [https://doi.org/](https://doi.org/10.1145/3236024.3236029)
3542 [10.1145/3236024.3236029](https://doi.org/10.1145/3236024.3236029)
- 3543 [50] Feargus Pendlebury, Fabio Pierazzi, Roberto Jordaney, Johannes Kinder, and Lorenzo
3544 Cavallaro. 2018. TESSERACT: Eliminating Experimental Bias in Malware Classification
3545 across Space and Time. (2018).
- 3546 [51] Sebastian Poeplau, Yanick Fratantonio, Antonio Bianchi, Christopher Kruegel, and Gio-
3547 vanni Vigna. 2014. Execute This! Analyzing Unsafe and Malicious Dynamic Code Loading
3548 in Android Applications. In *21st Annual Network and Distributed System Security Sympo-*
3549 *sium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*, 2014. The Internet
3550 Society.

- 3551 [52] Zhengyang Qu, Shahid Alam, Yan Chen, Xiaoyong Zhou, Wangjun Hong, and Ryan
3552 Riley. 2017. DyDroid: Measuring Dynamic Code Loading and Its Security Implications
3553 in Android Applications. In *2017 47th Annual IEEE/IFIP International Conference on*
3554 *Dependable Systems and Networks (DSN)*, June 2017. 415–426. [https://doi.org/10.1109/](https://doi.org/10.1109/DSN.2017.14)
3555 [DSN.2017.14](https://doi.org/10.1109/DSN.2017.14)
- 3556 [53] Bradley Reaves, Jasmine Bowers, Sigmund Albert Gorski III, Olabode Anise, Rahul
3557 Bobhate, Raymond Cho, Hiranava Das, Sharique Hussain, Hamza Karachiwala, Nolen
3558 Scaife, Byron Wright, Kevin Butler, William Enck, and Patrick Traynor. 2016. *droid:
3559 Assessment and Evaluation of Android Application Analysis Tools. *ACM Computing*
3560 *Surveys* 49, 3 (October 2016), 55:1–55:30. <https://doi.org/10.1145/2996358>
- 3561 [54] Atanas Rountev and Dacong Yan. 2014. Static Reference Analysis for GUI Objects in
3562 Android Software. In *Proceedings of Annual IEEE/ACM International Symposium on*
3563 *Code Generation and Optimization*, February 15, 2014. ACM, Orlando FL USA, 143–153.
3564 <https://doi.org/10.1145/2544137.2544159>
- 3565 [55] Antonio Ruggia, Dario Nisi, Savino Dambra, Alessio Merlo, Davide Balzarotti, and Simone
3566 Aonzo. 2024. Unmasking the Veiled: A Comprehensive Analysis of Android Evasive Mal-
3567 ware. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications*
3568 *Security (ASIA CCS '24)*, July 2024. Association for Computing Machinery, New York,
3569 NY, USA, 383–398. <https://doi.org/10.1145/3634737.3637658>
- 3570 [56] Jordan Samhi, Jun Gao, Nadia Daoudi, Pierre Graux, Henri Hoyez, Xiaoyu Sun, Kevin
3571 Allix, Tegawendé F. Bissyandé, and Jacques Klein. 2022. JuCify: a step towards Android
3572 code unification for enhanced static analysis. In *Proceedings of the 44th International*
3573 *Conference on Software Engineering (ICSE '22)*, July 2022. Association for Computing
3574 Machinery, New York, NY, USA, 1232–1244. <https://doi.org/10.1145/3510003.3512766>
- 3575 [57] Zhiyong Shan, Iulian Neamtiu, and Raina Samuel. 2018. Self-Hiding Behavior in Android
3576 Apps. In *40th International Conference on Software Engineering*, 2018. ACM Press, New
3577 York, New York, USA, 728–739. <https://doi.org/10.1145/3180155.3180214>
- 3578 [58] Feng Shen, Namita Vishnubhotla, Chirag Todarka, Mohit Arora, Babu Dhandapani, Eric
3579 John Lehner, Steven Y. Ko, and Lukasz Ziarek. 2014. Information Flows as a Permission
3580 Mechanism. In *Proceedings of the 29th ACM/IEEE International Conference on Automated*
3581 *Software Engineering*, September 15, 2014. ACM, Vasteras Sweden, 515–526. [https://doi.](https://doi.org/10.1145/2642937.2643018)
3582 [org/10.1145/2642937.2643018](https://doi.org/10.1145/2642937.2643018)
- 3583 [59] Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Andrew
3584 Dutcher, John Grosen, Siji Feng, Christophe Hauser, Christopher Kruegel, and Giovanni
3585 Vigna. 2016. SOK: (State of) The Art of War: Offensive Techniques in Binary Analysis.

-
- 3586 In *2016 IEEE Symposium on Security and Privacy (SP)*, 2016. 138–157. [https://doi.org/](https://doi.org/10.1109/SP.2016.17)
3587 10.1109/SP.2016.17
- 3588 [60] Ting Su, Guozhu Meng, Yuting Chen, Ke Wu, Weiming Yang, Yao Yao, Geguang Pu, Yang
3589 Liu, and Zhendong Su. 2017. Guided, stochastic model-based GUI testing of Android apps.
3590 In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*
3591 *(ESEC/FSE 2017)*, August 2017. Association for Computing Machinery, New York, NY,
3592 USA, 245–256. <https://doi.org/10.1145/3106237.3106298>
- 3593 [61] Thomas Sutter, Timo Kehrer, Marc Rennhard, Bernhard Tellenbach, and Jacques Klein.
3594 2024. Dynamic Security Analysis on Android: A Systematic Literature Review. *IEEE*
3595 *Access* 12, (2024), 57261–57287. <https://doi.org/10.1109/ACCESS.2024.3390612>
- 3596 [62] Kimberly Tam, Salahuddin Khan, Aristide Fattori, and Lorenzo Cavallaro. 2015. Copper-
3597 Droid: Automatic Reconstruction of Android Malware Behaviors. In *22nd Annual Network*
3598 *and Distributed System Security Symposium*, February 2015. The Internet Society, San
3599 Diego, California, USA.
- 3600 [63] Dennis Titze and Julian Schutte. 2015. Apparecium: Revealing Data Flows in Android
3601 Applications. In *2015 IEEE 29th International Conference on Advanced Information*
3602 *Networking and Applications*, March 2015. IEEE, Gwangju, South Korea, 579–586. [https://](https://doi.org/10.1109/AINA.2015.239)
3603 doi.org/10.1109/AINA.2015.239
- 3604 [64] Akihiko Tozawa and Masami Hagiya. 2002. Formalization and Analysis of Class Loading
3605 in Java. *Higher-Order and Symbolic Computation* 15, 1 (March 2002), 7–55. [https://doi.](https://doi.org/10.1023/A:1019912130555)
3606 [org/10.1023/A:1019912130555](https://doi.org/10.1023/A:1019912130555)
- 3607 [65] Timothy Vidas, Jiaqi Tan, Jay Nahata, Chaur Lih Tan, Nicolas Christin, and Patrick
3608 Tague. 2014. A5: Automated Analysis of Adversarial Android Applications. In *Proceedings*
3609 *of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*,
3610 November 07, 2014. ACM, Scottsdale Arizona USA, 39–50. [https://doi.org/10.1145/](https://doi.org/10.1145/2666620.2666630)
3611 [2666620.2666630](https://doi.org/10.1145/2666620.2666630)
- 3612 [66] Fengguo Wei, Sankardas Roy, Xinming Ou, and Robby. 2014. Amandroid: A Precise and
3613 General Inter-component Data Flow Analysis Framework for Security Vetting of Android
3614 Apps. In *ACM SIGSAC Conference on Computer and Communications Security*, Novem-
3615 ber 2014. ACM, Scottsdale Arizona USA, 1329–1341. [https://doi.org/10.1145/2660267.](https://doi.org/10.1145/2660267.2660357)
3616 [2660357](https://doi.org/10.1145/2660267.2660357)
- 3617 [67] Erik Ramsgaard Wognsen, Henrik Søndberg Karlsen, Mads Chr. Olesen, and René Rydhof
3618 Hansen. 2014. Formalisation and Analysis of Dalvik Bytecode. *Science of Computer*
3619 *Programming* 92, (October 2014), 25–55. <https://doi.org/10.1016/j.scico.2013.11.037>

- 3620 [68] Michelle Y Wong and David Lie. 2018. Tackling runtime-based obfuscation in Android
3621 with TIRO. In *USENIX Security Symposium*, August 2018. USENIX, Baltimore, USA,
3622 1247–1262.
- 3623 [69] Mingyuan Xia, Lu Gong, Yuanhao Lyu, Zhengwei Qi, and Xue Liu. 2015. Effective Real-
3624 Time Android Application Auditing. In *2015 IEEE Symposium on Security and Privacy*,
3625 May 2015. IEEE, San Jose, CA, 899–914. <https://doi.org/10.1109/SP.2015.60>
- 3626 [70] Lei Xue, Xiapu Luo, Le Yu, Shuai Wang, and Dinghao Wu. 2017. Adaptive unpacking
3627 of Android apps. In *International Conference on Software Engineering*, May 2017. IEEE,
3628 Buenos Aires, Argentina, 358–369.
- 3629 [71] Lok Kwong Yan and Heng Yin. 2012. DroidScope: Seamlessly Reconstructing the OS and
3630 Dalvik Semantic Views for Dynamic Android Malware Analysis. In *21st USENIX Security
3631 Symposium (USENIX Security 12)*, August 2012. USENIX Association, Bellevue, WA,
3632 569–584. Retrieved from [https://www.usenix.org/conference/usenixsecurity12/technical-
3633 sessions/presentation/yan](https://www.usenix.org/conference/usenixsecurity12/technical-sessions/presentation/yan)
- 3634 [72] Shengqian Yang, Dacong Yan, Haowei Wu, Yan Wang, and Atanas Rountev. 2015. Static
3635 Control-Flow Analysis of User-Driven Callbacks in Android Applications. In *2015 IEEE/
3636 ACM 37th IEEE International Conference on Software Engineering*, May 2015. IEEE,
3637 Florence, Italy, 89–99. <https://doi.org/10.1109/ICSE.2015.31>
- 3638 [73] Wenbo Yang, Yuanyuan Zhang, Juanru Li, Junliang Shu, Bodong Li, Wenjun Hu, and
3639 Dawu Gu. 2015. AppSpear: Bytecode Decrypting and DEX Reassembling for Packed
3640 Android Malware. In *Research in Attacks, Intrusions, and Defenses (Lecture Notes in
3641 Computer Science)*, 2015. Springer International Publishing, Cham, 359–381. [https://doi.
3642 org/10.1007/978-3-319-26362-5_17](https://doi.org/10.1007/978-3-319-26362-5_17)
- 3643 [74] Yueqian Zhang, Xiapu Luo, and Haoyang Yin. 2015. Dexhunter: toward extracting hidden
3644 code from packed android applications. In *European Symposium on Research in Computer
3645 Security*, November 2015. Springer, Vienna, Austria, 293–311.
- 3646 [75] Yury Zhauniarovich, Maqsood Ahmad, Olga Gadyatskaya, Bruno Crispo, and Fabio Mas-
3647 sacci. 2015. StaDynA: Addressing the Problem of Dynamic Code Updates in the Security
3648 Analysis of Android Applications. In *Proceedings of the 5th ACM Conference on Data and
3649 Application Security and Privacy*, March 02, 2015. ACM, San Antonio Texas USA, 37–48.
3650 <https://doi.org/10.1145/2699026.2699105>
- 3651 [76] Wenwen Zhou, Yang Yongzhi, and Jiejuan Wang. 2022. Dynamic Class Generating and
3652 Loading Technology in Android Web Application. In *2022 International Symposium on
3653 Networks, Computers and Communications (ISNCC)*, July 2022. 1–6. [https://doi.org/10.
3654 1109/ISNCC55209.2022.9851782](https://doi.org/10.1109/ISNCC55209.2022.9851782)



3655 **Titre :** Les difficultés de la rétro-ingénierie Android: de l'analyse large échelle au dé-brouillage dynamique

3656 **Mots clés :** Android, Analyse de Maliciels, Ingénierie Inverse, Chargement de Classe, Brouillage de Code

3657 **Résumé :** La place croissante des téléphones
3658 mobiles dans notre vie quotidienne en font une
3659 cible de choix pour les acteurs malveillants. Cette
3660 menace rend l'analyse d'application cruciale pour
3661 déterminer s'il s'agit ou non d'un maliciel et de son
3662 impact sur l'utilisateur s'il s'agit bien d'une applica-
3663 tion malveillante.

3664 Cette thèse explore les difficultés liées à
3665 l'ingénierie inverse d'applications Android. Dans un
3666 premier temps, elle reprend un effort de la commu-
3667 nauté qui a identifié les contributions entre 2011 et
3668 2017 portant sur l'analyse statique d'applications
3669 mobiles et propose une méthode pour évaluer
3670

la réutilisabilité des outils associés. Une étude
poussée des échecs lors de l'exécution de ces
outils montre que 54.55% d'entre eux ne sont plus
utilisables. Elle modélise ensuite le processus de
chargement des classes utilisées par une applica-
tion et présente une méthode de brouillage basée
sur les différences entre ce modèle et celui utilisé
par des outils d'analyses tels que Androguard
ou Flowdroid. Enfin, elle propose une approche
consistant à encoder les résultats d'une analyse
dynamique dans une nouvelle application valide
pour permettre aux outils existants d'analyser des
applications faisant usage de chargement de code
dynamique.

3671 **Title :** The Woes of Android Reverse Engineering: from Large Scale Analysis to Dynamic Deobfuscation

3672 **Keywords:** Android, Malware Analysis, Reverse Engineering, Class Loading, Code Obfuscation

3673 **Abstract:** The growing importance of smartphones
3674 in our daily lives makes them a prime target for
3675 malicious actors. This makes our ability to analyse
3676 an application critical, as it allows us to determine
3677 if an application is malware and its impact on the
3678 user.

3679 This thesis explores the difficulties of reverse engi-
3680 neering an Android application. First, we pursue
3681 a community effort that identified contributions
3682 between 2011 and 2017 about static analysis of
3683 Android applications, and we propose a method to
3684

evaluate the reusability of the associated tools. An
extensive analysis of the execution failures of those
tools shows that 54.55% of them are no longer
usable. Then, we model the mechanism that loads
the classes used by an application and present an
obfuscation method based on the discrepancies
between this model and the one used by analysis
tools like Androguard and Flowdroid. Finally, we
propose an approach that consists of encoding the
result of a dynamic analysis within a new valid
application, allowing existing tools to analyse appli-
cations that rely on dynamic code loading.