

SWAN: A Static Analysis Framework for Swift

Daniil Tiganov
University of Alberta
tiganov@ualberta.ca

Jeff Cho
University of Alberta
jeff.cho@ualberta.ca

Karim Ali
University of Alberta
karim.ali@ualberta.ca

Julian Dolby
IBM Research
dolby@us.ibm.com

ABSTRACT

Swift is an open-source programming language and Apple’s recommended choice for app development. Given the global widespread use of Apple devices, the ability to analyze Swift programs has significant impact on millions of users. Although static analysis frameworks exist for various computing platforms, there is a lack of comparable tools for Swift. While LLVM and Clang support some analyses for Swift, they are either primarily dynamic analyses or not suitable for deeper analyses of Swift programs such as taint tracking. Moreover, other existing tools for Swift only help enforce code styles and best practices.

In this paper, we present SWAN, an open-source framework that allows robust program analyses of Swift programs using IBM’s T.J. Watson Libraries for Analysis (WALA). To provide a wide range of analyses for Swift, SWAN leverages the well-established libraries in WALA. SWAN is publicly available at <https://github.com/themaplelab/swan>. We have also made a screencast available at <https://youtu.be/AZwfhOGqwFs>.

CCS CONCEPTS

• **Theory of computation** → **Program analysis.**

KEYWORDS

Swift, static analysis, taint analysis

1 INTRODUCTION

Static analysis reasons about the potential runtime behaviour of a program without necessarily executing it. Using this technique may help protect user privacy [1] and optimize applications [2]. Despite the potential benefits of static analysis, there is a lack of available tools for Swift [4], Apple’s recommended choice for development on iOS [15] and macOS [17]. In 2019, the web traffic analysis tool StatCounter estimated that iOS devices comprised approximately 24.79% of mobile devices in the world [10] and macOS devices accounted for 16.46% of desktop devices [9]. Trends also show that the popularity of both operating systems in 2020 has increased by 4.41% and 3.82%, respectively. Therefore, the ability to analyze Swift apps has significant impact on millions of users around the world.

To bridge the gap between the increasing popularity of Swift and the lack of available analysis tools, we introduce SWAN, an open-source static analysis framework for Swift. We designed SWAN with the app developer as our main target audience. Therefore, SWAN offers both a Command-Line Interface (CLI) and a Graphical User Interface (GUI). Both interfaces offer the same functionalities but address different use cases. For example, the CLI enables the

developer to integrate SWAN in their continuous integration workflow, providing analysis results at major development milestones. On the other hand, the GUI enables SWAN to provide the developer with up-to-date analysis results whenever they modify their code. This immediate feedback helps developers focus on the task at hand, which further helps them fix more bugs in less time [7].

While designing SWAN, we have also taken into consideration future contributions to its underlying analysis engine. To enable contributions from the wider static analysis community, we opted for building SWAN on top of the well-established IBM T.J. Watson Libraries for Analysis (WALA) [6]. We re-used various analysis components that WALA has built over the years, and has proven to work well for analyzing different programming languages such as Java, JavaScript, and Python. This design decision enables app developers to use the various analyses that SWAN offers out of the box (e.g., taint analysis, pointer analysis, call graph construction, and inter-procedural dataflow analysis) without having to implement their own analysis. Moreover, SWAN has a modular architecture that enables researchers to build their own analyses on top of it by leveraging its existing analysis infrastructure.

Through its suite of analyses, SWAN enables new directions of research for iOS and macOS that have long existed for other platforms such as Android [1], Java [13], and JavaScript [23].

2 HOW CAN APP DEVELOPERS USE SWAN?

App developers may use SWAN through either one of its frontends: a command-line interface and a VSCode [24] extension. We will demonstrate the features of both frontends through the built-in taint analysis of SWAN, which tracks data leaks in a given program.

2.1 Command-Line Interface

SWAN provides a CLI script called `run-swan-single`. This script analyzes a single Swift file¹. The user may define sources of private information (i.e., sources), potential locations where data may leak to (i.e., sinks), and methods that properly secure private information (i.e., sanitizers) in the source file as code comments. Figure 1 represents a sample Swift program that exhibits a tainted dataflow. The script runs on the input file that has a source and a sink defined as code comments. When SWAN finishes its analysis, it prints the results to the terminal. Figure 2 shows how SWAN formats the results in a tree structure, where each path consists of a source, sink, and path edges (i.e., intermediates).

2.2 VSCode Extension

SWAN provides a GUI via a custom VSCode extension for analyzing Swift programs and viewing its results. To use the extension, the user must first configure SWAN under `Settings`→`SWAN`. Figure 3

¹We are developing a frontend to support Xcode projects. In the past, our frontend has successfully analyzed Xcode projects. However, due to recent changes to Xcode, Swift, and macOS versioning, our Xcode frontend is no longer functional.

```

1 // SWAN:sources: "source() -> Swift.String"
2 // SWAN:sinks: "sink(sunk: Swift.String) -> ()"
3 func source() -> String { return "I'm bad"; }
4 func sink(sunk: String) { print(sunk); }
5 func random() -> String { return "whatever"; }
6 let whatever = random();
7 let src = source();
8 let combined = whatever + src;
9 sink(sunk: combined);

```

Figure 1: A Swift program with tainted dataflow.

```

10 $./utils/run-swan-single -sdk $SDK_PATH -path ./ca.
    maple.swan.swift.test/tests/taint-single-source/
    StringConcat.swift
11 [...]
12 ===== RESULTS =====
13 -- PATH
14 -- SOURCE
15 7:11 in [...]/StringConcat.swift
16 let src = source();
17 ^
18 -- INTERMEDIATE
19 8:16 in [...]/StringConcat.swift
20 let combined = whatever + src;
21 ^
22 -- SINK
23 9:1 in [...]/StringConcat.swift
24 sink(sunk: combined);
25 ^
26 ===== END OF RESULTS =====

```

Figure 2: An example illustrating how to use the SWAN CLI frontend to analyze a single Swift file.

```

27 {
28 "swan.SDKPath": "/Applications/Xcode.app/Contents/
    Developer/Platforms/MacOSX.platform/Developer/
    SDKs/MacOSX.sdk/",
29 "swan.CustomSSS": {
30 "swan.Sources": ["source() -> Swift.String"],
31 "swan.Sinks": ["sink(sunk: Swift.String) -> ()"],
32 "swan.Sanitizers": []
33 },
34 "swan.SingleFilePath": "<user>/Documents/swan/ca.
    maple.swan.swift.test/tests/taint-single-source/
    StringConcat.swift",
35 "swan.TaintAnalysisMode": "Refined"
36 }

```

Figure 3: An example illustrating the contents of the configuration file (settings.json) for the SWAN VSCode extension.

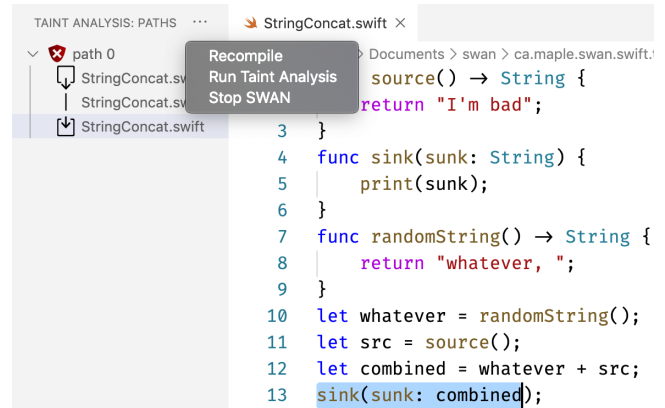


Figure 4: The main GUI of the SWAN VSCode extension.

shows an example configuration for the SWAN settings.json file. Figure 4 shows the main GUI elements of the extension. To easily edit the configuration file, SWAN provides function name autocompletion. If the configuration file changes, the user may still quickly re-run the client analysis, without recompiling the code. To recompile the Swift program, the user must press *Recompile*.

After configuring the extension, the user may start SWAN by selecting the SWAN tab and pressing *Run Taint Analysis* in the sidebar. The extension then automatically attaches to an existing SWAN Java Virtual Machine (JVM), if one is running, or starts a new JVM if one is not already running. We recommend to first launch the JVM separately to view the console output, especially for debugging purposes. To enable bidirectional communication with the running SWAN JVM, the extension uses sockets.

For our example taint analysis, SWAN displays the results in the sidebar in filetree-like form. Each vulnerable path is an element in the tree with a red cross beside it. Its children are the nodes in the path. The first child is the source, the last is the sink, and any nodes in between are intermediates. The user may select any path node, and the extension will open the file at the corresponding source location. In Figure 4, the user has selected the last node. Therefore, SWAN highlights Line 13 in the source file, showing the user that tainted data reaches the function `sink()` as a parameter.

3 THE MAIN WORKFLOW OF SWAN

SWAN has a linear workflow where each component produces the data requested by its parent component. Figure 5 represents the workflow components and numerically labels them according to their execution order. At the beginning of the workflow ①, a SWAN frontend instantiates a JVM using its corresponding *SWAN Driver* ② with build and analysis options. The user may configure these options in the frontend. The build options contain arguments needed to call the Swift compiler properly. Therefore, SWAN propagates them through its components all the way to *SWAN Hook* ⑥, where SWAN eventually calls the Swift compiler.

To represent the program dataflow, SWAN uses the System Dependence Graph (SDG), an internal WALA data structure. To generate the SDG, *SWAN Driver* first calls the frontend-agnostic *WALA Driver* ③, which then calls *WALA Analysis Engine* ④ to construct




Figure 5: The main workflow of SWAN.

a call graph (CG) for the input program. Using WALA libraries, *WALA Driver* then generates the SDG from that CG and returns it back to *SWAN Driver*. Besides storing call information, the CG data structure stores the class hierarchy and the Intermediate Representation (IR) of the input program. To construct the CG, *WALA Analysis Engine* first translates the input Swift program to WALA IR, and then generates the CG from that IR. Internally, WALA uses the Common Abstract Syntax Tree (CAst) as an IR between the input language and WALA IR itself. To translate the input Swift code to WALA CAst, the *WALA Analysis Engine* uses *SWAN Translator* (5).

After extensive discussions with the Apple Swift team, we determined that it is best for SWAN to consume Swift Intermediate Language (SIL), an internal IR within the Swift compiler pipeline, as input instead of the Swift Abstract Syntax Tree (AST). To maintain fidelity to the original source code, SWAN operates on *raw* SIL before the Swift compiler applies any code transformations or optimizations. Through a Java Native Interface (JNI) call, *SWAN Hook* (6) invokes the Swift compiler and provides it with a callback handler that receives the SIL during compilation. The callback handler uses a custom visitor to process all compilation contexts: *SILModule*, *SILFunction*, *SILBlock*, and *SILInstruction*. This visitor packages all information needed for translating it into JNI objects (i.e., JVM objects in C++ form). To translate the code to WALA CAst, *SWAN Hook* sends back the packaged information to *SWAN Translator* using JNI.

Once SWAN translates the input Swift code to WALA CAst, SWAN may run a client analysis (e.g., taint analysis) on the program, represented in SDG form (7). SWAN then returns the results to the frontend to be displayed in the appropriate format.

4 TRANSLATING SWIFT SOURCE CODE

SWAN leverages the existing infrastructure in WALA to translate the SIL representation of an input Swift program into WALA IR. To achieve that, SWAN first translates the input SIL to CAst, which requires an intermediate level of abstraction due to the significant differences between the two IRs. In particular, SIL uses pointers whereas CAst uses references. To serve as this abstraction, we have developed SWAN IR as a simple, hybrid IR between SIL and WALA IR. Table 1 lists the main SWAN IR instructions, which currently support all but two SIL instructions: *partial_apply* and *assign_by_wrapper*. Similar to SIL, SWAN IR is only composed of functions, basic blocks, and instructions.

Table 1: A list of the main SWAN IR instructions.

Instruction	Notation
new	$v_0 := \text{new } \$String$
assign	$v_0 := v_1$
literal	$v_0 := \#foo, v_0 := \#12$
goto	goto bb0
conditional goto	cond_br v0 true: bb1, false: bb2
throw	throw v0
conditional throw	throw if v0
return	return v0
field read	$v_0 := v_1.foo$
field write	$v_0.foo := v_1$
function ref	$v_0 := \text{func_ref test.foo() } \rightarrow \text{Swift.Int}$
builtin	$v_0 := \text{builtin Swift.Int.init[...]}$
apply	$v_0 := v_1(v_2, v_3)$
try apply	try v0(v1) normal: bb1, error: bb2
binary op	$v_0 := v_1 + v_2$
unary op	$v_0 := !v_1$

4.1 Field Aliases

SWAN Translator converts a SIL pointer to an object with the field value representing its underlying value. This conversion is nontrivial. For example, the SIL instruction *ref_element_addr* derives the address of a class field, and writes it to a value that may be accessed later. Therefore, treating the instruction as a regular field read is not sufficient. To address this problem, the SWAN IR symbol table has a special type called *Field Alias* that aliases a field access path. To handle instructions such as *ref_element_addr*, SWAN uses this type to lazily defer the field access until the field value is accessed. Figure 6 illustrates a SIL example where *ref_element_addr* derives the address of the class field *A.foo* (Line 45) and stores it to %2. This class field is later read by the local variable %4 (Line 47) through the alias %3 (Line 46). To translate this example correctly, SWAN uses the *Field Alias* type.

4.2 SIL Coroutines and Instructions

SWAN Translator simplifies complex components of SIL. In particular, SWAN handles asymmetric coroutines by inlining them into their caller. Moreover, SWAN ignores instructions that we have determined do not mutate or move data in a manner that is relevant to dataflow analysis such as low-level memory management instructions.

4.3 SIL Builtin Functions

SIL handles many operations through builtin functions. For instance, the three data structures that Swift provides (i.e., Array, Set, and Dictionary) are entirely accessed via builtin function calls. Simple tasks such as literal manipulation and string operations are also handled using builtins. To support builtin functions in SWAN, we have developed a SWAN IR parser that reads handwritten summaries. Currently, SWAN supports numerous builtins such as literal and array operations. To fully support the main Swift data structures, we are continuously adding support for more builtins, updating the SWAN IR instructions whenever necessary.

```

37 ---- Swift ----
38 class A {
39     var foo = "bar";
40 }
41 ---- SIL of A.foo.getter ----
42 sil hidden [transparent] [ossa] @$s4temp1AC3fooSSvg :
      $@convention(method) (@guaranteed A) -> @owned
      String {
43 bb0(%0 : @guaranteed $A):
44     debug_value %0 : $A, let, name "self", argno 1
45     %2 = ref_element_addr %0 : $A, #A.foo
46     %3 = begin_access [read] [dynamic] %2 : $*String //
          copy %2 to %3
47     %4 = load [copy] %3 : $*String // realize alias %3
          as a field read to %4
48     end_access %3 : $*String
49     return %4 : $String
50 }
51 ---- SWANIR of A.foo.getter ----
52 func $String`temp.A.foo.getter` : Swift.String`(v0 :
      $A) {
53     bb0(v0 : $A) :
54         v1 := v0.foo
55         return v1
56 }

```

Figure 6: An example illustrating how `ref_element_addr` may cause field aliasing.

To enable translation to other languages, we have designed SWAN to provide a modular workflow. Other framework designers may use the SWAN IR visitor to easily translate other input languages to WALA IR.

5 RELATED WORK

5.1 Android Analysis Frameworks

The Android platform has seen an abundance of analysis frameworks over the past decade. FlowDroid [1] is a lifecycle-aware and context-sensitive, flow-sensitive, field-sensitive, and object-sensitive taint analysis tool for Android apps. The design of FlowDroid inspired us in the design of SWAN. Similar to SWAN, SCanDroid [3] also uses WALA, but for the purpose of matching Android app manifests to dataflow analyses to ensure that apps do not overreach their permissions. DroidInfer [14] uses context-free language reachability to perform type-based and context-sensitive taint analyses for Android apps.

While all these frameworks work well for the Android platform, there is no openly-available equivalent counterpart for the Swift platform. SWAN bridges this gap by providing the first open-source static analysis framework for Swift.

5.2 LLVM-Based Analyses

While LLVM [11] and Clang [12] support some low-level analyses, they are not suitable for deeper analyses of Swift applications such

as precise taint tracking. This is because most Swift-specific structures and information are typically lost during the compilation of Swift source code to low-level LLVM IR. Moreover, the most useful analyses that Clang provides (i.e., memory sanitizer and thread sanitizer) are primarily dynamic analyses. Unlike static analyses, dynamic analyses require running the Swift program under analysis multiple times with various inputs to ensure enough coverage of the program behaviour. SWAN overcomes this limitation by providing a framework for static analysis of Swift programs.

The Phasar framework [19] provides call graph construction and dataflow analyses on LLVM IR, enabling it to analyze Swift applications. However, similar to other LLVM-based frameworks, Phasar focuses more on low-level LLVM constructs, whereas SWAN analyzes the SIL representation [5], preserving important information from the Swift source code such as the file and line number of the corresponding originating code, which is important for notifying developers of the locations of identified issues.

5.3 Swift Analysis Tools

Most publicly available analysis tools for Swift are linters such as SwiftLint [18] and Tailor [20]. Those tools only help enforce Swift code standards and best practices.

SonarSwift [22] is a static Swift code analyzer which allows users to define rules for bugs, code smells, and vulnerabilities to find in their codebase. Some existing rules listed on the website [21] include not using identical expressions on both sides of a binary operator (i.e., bug), not duplicating string literals (i.e., code smell), and avoiding using DES (i.e., vulnerability). However, at the time of this writing, Swift is not supported in the free version of the software and requires a paid license. As a result, we are unable to verify its correctness and effectiveness at Swift static analysis.

6 CONCLUSION

We presented SWAN, a static analysis framework for Swift that we built on top of the WALA analysis framework. SWAN provides various analyses to its users including call graph analysis, pointer analysis, and inter-procedural dataflow analysis. To enable wider use of SWAN, we have developed two user interfaces for it: a command-line interface and a VSCode extension. We have also designed SWAN to be a modular framework where some of its components (e.g., SWAN IR and SDG-based dataflow analysis) may be used for other WALA-based analysis frameworks. In the future, we plan to build more support for analyzing iOS apps in SWAN, which will put it on par with its Android counterparts with respect to analyzing mobile applications. SWAN is open source [16], and we welcome contributions under the Eclipse Public License 2.0 [8].

ACKNOWLEDGMENTS

We would like to thank all the following contributors to the SWAN project: Noah Weninger, Leo Li, Mark Mroz, Yaser Alkayale, Lydia Wu, Chen Song, Bryan Tam, and Anthony Hill. We would also like to thank Dillon Pratt for his help with video-editing our screencast. This material is based upon work supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick D. McDaniel. 2014. FlowDroid: precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for Android apps. In *Programming Language Design and Implementation (PLDI)*, 259–269. <https://doi.org/10.1145/2594291.2594299>
- [2] Nathaniel Ayewah, David Hovemeyer, J. David Morgenthaler, John Penix, and William Pugh. 2008. Using Static Analysis to Find Bugs. *IEEE Software* 25, 5 (2008), 22–29. <https://doi.org/10.1109/MS.2008.130>
- [3] Tanzirul Azim and Iulian Neamtiu. 2013. Targeted and depth-first exploration for systematic testing of android apps. In *Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes (Eds.), 641–660. <https://doi.org/10.1145/2509136.2509549>
- [4] Swift Community. 2015. *The Swift Programming Language*. Retrieved Jan 24, 2020 from <https://swift.org/>
- [5] Swift Community. 2020. *Swift Intermediate Language (SIL)*. Retrieved May 28, 2020 from <https://github.com/apple/swift/blob/master/docs/SIL.rst>
- [6] WALA Community. 2006. *T.J. Watson Libraries for Analysis*. Retrieved May 28, 2020 from <https://github.com/wala/WALA/>
- [7] Lisa Nguyen Quang Do, Karim Ali, Benjamin Livshits, Eric Bodden, Justin Smith, and Emerson R. Murphy-Hill. 2017. Just-in-time Static Analysis. In *International Symposium on Software Testing and Analysis (ISSTA)*, 307–317. <https://doi.org/10.1145/3092703.3092705>
- [8] Eclipse Foundation. 2017. *Eclipse Public License 2.0*. Retrieved May 28, 2020 from <https://www.eclipse.org/legal/epl-2.0/>
- [9] StatCounter GlobalStats. 2019. *Desktop Operating System Market Share Worldwide*. Retrieved Mar 17, 2020 from <https://gs.statcounter.com/os-market-share/desktop/worldwide/#monthly-201901-201912>
- [10] StatCounter GlobalStats. 2019. *Mobile Operating System Market Share Worldwide*. Retrieved Mar 17, 2020 from <https://gs.statcounter.com/os-market-share/mobile/worldwide/#monthly-201901-201912>
- [11] LLVM Developer Group. 2003. *The LLVM Compiler Infrastructure*. Retrieved May 28, 2020 from <https://llvm.org/>
- [12] LLVM Developer Group. 2007. *Clang: a C language family frontend for LLVM*. Retrieved May 28, 2020 from <https://clang.llvm.org/>
- [13] The Sable Group. 1999. *Soot - A Java optimization framework*. McGill University, Montréal, QC, Canada. Retrieved June 4, 2020 from <https://github.com/Sable/soot>
- [14] Wei Huang, Yao Dong, Ana Milanova, and Julian Dolby. 2015. Scalable and precise taint analysis for Android. In *International Symposium on Software Testing and Analysis (ISSTA)*, Michal Young and Tao Xie (Eds.), 106–117. <https://doi.org/10.1145/2771783.2771803>
- [15] iOS Team. 2007. *iOS 13 - Apple (CA)*. Apple, Cupertino, California, USA. Retrieved Jan 24, 2020 from <https://www.apple.com/ca/ios/>
- [16] The Maple Lab. 2017. *A Swift Static Analysis Framework*. University of Alberta, Edmonton, Alberta, Canada. Retrieved May 28, 2020 from <https://github.com/themaplelab/swan>
- [17] macOS Team. 2001. *macOS Mojave - Apple (CA)*. Apple, Cupertino, California, USA. Retrieved Jan 24, 2020 from <https://www.apple.com/ca/macOS/mojave/>
- [18] Realm. 2015. *SwiftLint - A tool to enforce Swift style and conventions*. Retrieved May 28, 2020 from <https://github.com/realm/SwiftLint>
- [19] Philipp Dominik Schubert, Ben Hermann, and Eric Bodden. 2019. PhASAR: An Inter-procedural Static Analysis Framework for C/C++. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 393–410. https://doi.org/10.1007/978-3-030-17465-1_22
- [20] Sleekbyte. 2015. *Tailor - Cross-platform static analyzer and linter for Swift*. Retrieved May 28, 2020 from <https://tailor.sh/>
- [21] SonarSource. 2017. *Swift Static Code Analysis Rules*. Retrieved May 28, 2020 from <https://rules.sonarsource.com/swift>
- [22] SonarSwift. 2015. *Code Quality and Security for Swift*. SonarSource, Geneva, Switzerland. Retrieved May 28, 2020 from <https://www.sonarsource.com/swift/>
- [23] TAJs Team. 2009. *Type Analyzer for JavaScript*. Århus University, Århus, Denmark. Retrieved June 4, 2020 from <https://github.com/cs-au-dk/TAJS>
- [24] Visual Studio Team. 2015. *Visual Studio Code - Code Editing. Redefined*. Microsoft, Redmond, Washington, USA. Retrieved Feb 19, 2020 from <https://code.visualstudio.com>