

AnaDroid: Malware Analysis of Android with User-supplied Predicates

Shuying Liang,^{1,2} Matthew Might¹

*School of Computing, University of Utah
Salt Lake City, Utah, USA*

David Van Horn¹

*College of Computer and Information Sciences
Northeastern University
Boston, Massachusetts, USA*

Abstract

Today's mobile platforms provide only coarse-grained permissions to users with regard to how third-party applications use sensitive private data. Unfortunately, it is easy to disguise malware within the boundaries of legitimately-granted permissions. For instance, granting access to "contacts" and "internet" may be necessary for a text-messaging application to function, even though the user does not want contacts transmitted over the internet. To understand fine-grained application use of permissions, we need to statically analyze their behavior. Even then, malware detection faces three hurdles: (1) analyses may be prohibitively expensive, (2) automated analyses can only find behaviors that they are designed to find, and (3) the maliciousness of any given behavior is application-dependent and subject to human judgment. To remedy these issues, we propose semantic-based program analysis, with a human in the loop as an alternative approach to malware detection. In particular, our analysis allows analyst-crafted semantic predicates to search and filter analysis results. Human-oriented semantic-based program analysis can systematically, quickly and concisely characterize the behaviors of mobile applications. We describe a tool that provides analysts with a library of the semantic predicates and the ability to dynamically trade speed and precision. It also provides analysts the ability to statically inspect details of every suspicious state of (abstract) execution in order to make a ruling as to whether or not the behavior is truly malicious with respect to the intent of the application. In addition, permission and profiling reports are generated to aid analysts in identifying common malicious behaviors.

Keywords: static analysis, human analysis, malware detection

¹ Supported by the DARPA Automated Program Analysis for Cybersecurity Program.

² An extended report is available: <http://matt.might.net/a/2013/05/25/anadroid/>

1 Introduction

Google’s Android is the most popular mobile platform, with a share of 52.5% of all smartphones [10]. Due to Android’s open application development community, more than 400,000 apps are available with 10 billion cumulative downloads by the end of 2011 [9].

While most of those third-party applications have legitimate reasons to access private data, the grantable permissions are too coarse: malware can hide in the cracks. For instance, an app that should only be able to read information from a specific site and have access to GPS information must necessarily be granted full read/write access to the entire internet, thereby allowing a possible location leak over the net. Or, a note-taking application can wipe out SD card files when a hidden trigger condition is met. Meanwhile, a task manager that requests every possible permission can be legitimately benign.

To understand fine-grained use of security-critical resources, we need to statically analyze the application with respect to what data is accessed, where the sensitive data flows, and what operations have been performed on the data (*i.e.*, determine whether the data is tampered with). Even then, automated malware detection faces three hurdles: (1) analyses may be prohibitively expensive, (2) automated analyses can only find behaviors that they are designed to find, and (3) the maliciousness of any given behavior is application-dependent and subject to human judgment.

In this work, we propose semantics-based program analysis with a human in the loop as an alternative approach to malware detection. Specifically, we derive an analytic engine, an abstract CESK* machine based on the design methodology of Abstracting Abstract Machines (AAM) [20] to analyze object-oriented bytecode. Then we extend the foundational analysis to analyze specific features: multiple entry points of Android apps and reflection APIs. Finally, we describe a tool that provides analysts with a library of semantic predicates that can be used to search and filter analysis results, and the ability to dynamically trade speed and precision. The tool also provides analysts the ability to statically inspect details of every suspicious state of (abstract) execution in order to make a ruling as to whether or not the behavior is truly malicious with respect to the intent of the application. Human-oriented, semantics-based program analysis can systematically characterize the behaviors of mobile applications.

Overview

The remainder of the paper is organized as follows: Section 2 presents the syntax of an object-oriented byte code, and illustrates a finite-state-space-based abstract interpretation of the byte code. Section 3 discusses analysis

$$\begin{aligned}
& \text{program} ::= \text{class-def } \dots \\
\text{class-def} \in \text{ClassDef} & ::= (\text{attribute } \dots \text{ class } \text{class-name} \text{ extends } \text{class-name} \\
& \quad (\text{field-def } \dots) (\text{method-def } \dots)) \\
\text{field-def} & ::= (\text{field } \text{attribute } \dots \text{ field-name } \text{type}) \\
\text{method-def} \in \text{MethodDef} & ::= (\text{method } \text{attribute } \dots \text{ method-name } (\text{type } \dots) \text{type} \\
& \quad (\text{throws } \text{class-name } \dots) (\text{limit } n) s \dots) \\
s \in \text{Stmt} & ::= (\text{label } \text{label}) \mid (\text{nop}) \mid (\text{line } \text{int}) \mid (\text{goto } \text{label}) \\
& \quad \mid (\text{if } \text{æ} (\text{goto } \text{label})) \mid (\text{assign } \text{name } [\text{æ} \mid \text{ce}]) \mid (\text{return } \text{æ}) \\
& \quad \mid (\text{field-put } \text{æ}_o \text{ field-name } \text{æ}_v) \mid (\text{field-get } \text{name } \text{æ}_o \text{ field-name}) \\
\text{æ} \in \text{AExp} & ::= \text{this} \mid \text{true} \mid \text{false} \mid \text{null} \mid \text{void} \mid \text{name} \mid \text{int} \\
& \quad \mid (\text{atomic-op } \text{æ} \dots \text{æ}) \mid \text{instance-of}(\text{æ}, \text{class-name}) \\
\text{ce} & ::= (\text{new } \text{class-name}) \\
& \quad \mid (\text{invoke-kind } \text{method-name } (\text{æ} \dots \text{æ}) (\text{type}_0 \dots \text{type}_n)) \\
\text{invoke-kind} & ::= \text{invoke-static} \mid \text{invoke-direct} \mid \text{invoke-virtual} \\
& \quad \mid \text{invoke-interafce} \mid \text{invoke-super} \\
\text{type} & ::= \text{class-name} \mid \text{int} \mid \text{byte} \mid \text{char} \mid \text{boolean} \\
\text{attribute} & ::= \text{public} \mid \text{private} \mid \text{protected} \mid \text{final} \mid \text{abstract}.
\end{aligned}$$

Fig. 1. An object-oriented bytecode adapted from the Android specification [18].

techniques to analyze Android-specific issues: multiple entry points and reflection APIs. Section 4 presents the tool implementation with user-supplied predicates. Section 5 discusses related work, and Section 6 concludes.

2 Semantic-based program analysis

Android apps are written in Java, and compiled into Dalvik virtual machine byte code (essentially a register-based version of Java byte code). In this section, we present how to derive an analysis for a core object-oriented (OO) byte code language based on Dalvik. After presenting this foundational analysis, we shall illustrate Android-specific analysis techniques in subsequent sections.

The first step is to define a syntax. Figure 1 presents the syntax of an OO byte code language that is closely modeled on the Dalvik virtual machine. Statements encode individual actions; atomic expressions encode atomically computable values; and complex expressions encode expressions with possible non-termination or side effects. There are four kinds of names: **Reg** for registers, **ClassName** for class names, **FieldName** for field names and **MethodName** for method names. The special register name **ret** holds the return value of the last function called. With respect to a given program, we assume a syntactic meta function $\mathcal{S} : \text{Label} \rightarrow \text{Stmt}^*$, which maps a label to the sequence of statements that start with that label.

Ordinarily, the next step toward an analyzer would be to derive a concrete

machine to interpret the language just defined. The meaning of a program will be defined as the set of machine states reachable from an initial state. The purpose of static analysis is to derive a computable approximation of the concrete machine’s behavior—of these states. We’ll construct an abstract semantics to do that. Since the concrete and the abstract semantics are so close in structure, we will present only the abstract semantics of the byte code, while highlighting places that are different from its concrete counterpart to save space.

2.1 Abstract semantics

We define our abstract interpretation as a direct, structural abstraction of a machine model for the OO bytecode [20]. Because the structural abstraction creates an abstract machine nearly identical to the machine model itself (with exceptions that we explain), we don’t provide the concrete machine model. The analysis of a program is defined as the set of *abstract* machine states reachable by an abstract transition relation (\rightsquigarrow)—the core of the abstract semantics. That is, abstract evaluation is defined by the set of states reached by the reflexive, transitive closure of the (\rightsquigarrow) relation.

Figure 2 details the abstract state-space. We assume the natural element-wise, point-wise and member-wise lifting of a partial order (\sqsubseteq) across this state-space. States of this machine consist of a tuples of of statements, frame pointers, heaps, and stack pointers. To synthesize the abstract state-space, we force frame pointers and object pointers (and thus addresses) to be a finite set. When we compact the set of addresses into a finite set during a structural abstraction, the machine may (and likely will) run out of addresses to allocate, and when it does, the pigeon-hole principle will force multiple abstract values to reside at the same (now abstract) address. As a result, we have no choice but to force the range of the \widehat{Store} to become a power set in the abstract state-space: now each abstract address can hold multiple values.

2.1.1 Abstract transition relation

In this section, we provide major cases for the abstract transition relation. The abstract transition relation delegates to helper functions for injecting programs into states, and for evaluating atomic expressions and looking up field values:

- $\widehat{\mathcal{I}} : \text{Stmt}^* \rightarrow \widehat{State}$ injects an sequence of instructions into an initial state:
 $\hat{c}_0 = \widehat{\mathcal{I}}(\mathbf{s}) = (\mathbf{s}, \hat{fp}_0, [\hat{a}_{\kappa_0} \mapsto \text{halt}], \hat{a}_{\kappa_0})$
- $\widehat{\mathcal{A}} : \text{AExp} \times \widehat{FramePointer} \times \widehat{Store} \rightarrow \widehat{Val}$ evaluates atomic expressions (specifically for variable look-up): $\widehat{\mathcal{A}}(\text{name}, \hat{fp}, \hat{\sigma}) = \sigma(\hat{fp}, \text{name})$

$\hat{c} \in \widehat{State} = \mathbf{Stmt}^* \times \widehat{FramePointer} \times \widehat{Store} \times \widehat{KontAddr}$	[states]
$\hat{\sigma} \in \widehat{Store} = \widehat{Addr} \rightarrow \widehat{Val}$	[stores]
$\hat{a} \in \widehat{Addr} = \widehat{RegAddr} + \widehat{FieldAddr} + \widehat{KontAddr}$	[addresses]
$\widehat{a}_\kappa \in \widehat{KontAddr}$ is a finite set of continuation addresses	
$\hat{r}a \in \widehat{RegAddr} = \widehat{FramePointer} \times \mathbf{Reg}$	
$\hat{f}a \in \widehat{FieldAddr} = \widehat{ObjectPointer} \times \mathbf{FieldName}$	
$\hat{\kappa} \in \widehat{Kont} = \mathbf{fun}(\hat{f}p, \mathbf{s}, \widehat{a}_\kappa) + \mathbf{halt}$	[continuations]
$\hat{d} \in \widehat{Val} = \mathcal{P}(\widehat{ObjectValue} + \widehat{String} + \widehat{\mathcal{Z}} + \widehat{Kont})$	[abstract values]
$\hat{o}v \in \widehat{ObjectValue} = \widehat{ObjectPointer} \times \mathbf{ClassName}$	
$\hat{f}p \in \widehat{FramePointer}$ is a finite set of frame pointers	[frame pointers]
$\hat{o}p \in \widehat{ObjectPointer}$ is a finite set of object pointers	[object pointers].

Fig. 2. The abstract state-space.

- $\hat{\mathcal{A}}_{\mathcal{F}} : \mathbf{AExp} \times \widehat{FramePointer} \times \widehat{Store} \times \mathbf{FieldName} \rightarrow \widehat{Val}$ looks up fields:

$$\hat{\mathcal{A}}_{\mathcal{F}}(\mathfrak{x}_o, \hat{f}p, \hat{\sigma}, \mathit{field-name}) = \bigsqcup \hat{\sigma}(\hat{o}p, \mathit{field-name}), \text{ where}$$

$$(\hat{o}p, \mathit{class-name}) \in \hat{\mathcal{A}}(\mathfrak{x}_o, \hat{f}p, \hat{\sigma}).$$

The rules for the abstract transition relation $(\rightsquigarrow) \subseteq \widehat{State} \times \widehat{State}$ describe how components of state evolve in light of each kind of statement. In subsequent paragraphs, we will illustrate the important rules that involve objects and function calls, omitting less important ones to save space:

- *New object creation* Creating a new object allocates a potentially non-fresh address and joins the newly initialized object to other values residing at this store address.

$$\overbrace{(\llbracket \mathbf{assign\ name\ (new\ class-name)} \rrbracket : \mathbf{s} \rrbracket, \hat{f}p, \hat{\sigma}, \widehat{a}_\kappa)}^{\hat{c}} \Rightarrow (\mathbf{s}, \hat{f}p, \hat{\sigma}'', \widehat{a}_\kappa), \text{ where}$$

$$\hat{o}p' = \mathit{allocOP}(\hat{c}), \hat{\sigma}' = \hat{\sigma} \sqcup [(\hat{f}p, \mathit{name}) \mapsto (\hat{o}p', \mathit{class-name})],$$

$$\hat{\sigma}'' = \mathit{initObject}(\hat{\sigma}', \mathit{class-name}),$$

where the helper $\mathit{initObject} : \widehat{Store} \times \mathbf{ClassName} \rightarrow \widehat{Store}$ initializes fields.

- *Instance field reference/update* Referencing a field uses $\hat{\mathcal{A}}_{\mathcal{F}}$ to lookup the field values and joins these values with the values at the store location for

the destination register:

$$\begin{aligned} & \llbracket (\text{field-get } \mathfrak{x}_o \text{ field-name}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa} \rightsquigarrow (\mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a_\kappa}), \text{ where} \\ & \hat{\sigma}' = \hat{\sigma} \sqcup [(\hat{fp}, \text{field-name}) \mapsto \hat{\mathcal{A}}_{\mathcal{F}}(\mathfrak{x}_o, \hat{fp}, \hat{\sigma}, \text{field-name})]. \end{aligned}$$

Updating a field first determines the abstract object values from the store, extracts the object pointer from all the possible values, then pairs the object pointers with the field name to get the field address, and finally *joins* the new values to those found at this store location:

$$\begin{aligned} & \llbracket (\text{field-put } \mathfrak{x}_o \text{ field-name } \mathfrak{x}_v) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa} \rightsquigarrow (\mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a_\kappa}), \text{ where} \\ & \hat{\sigma}' = \hat{\sigma} \sqcup [(\widehat{op}, \text{field-name}) \mapsto \hat{\mathcal{A}}(\mathfrak{x}_v, \hat{fp}, \hat{\sigma})], (\widehat{op}, \text{class-name}) \in \hat{\mathcal{A}}(\mathfrak{x}_o, \hat{fp}, \hat{\sigma}). \end{aligned}$$

- *Method invocation* This rule involves all four components of the machine. The abstract interpretation of non-static method invocation can result in the method being invoked on a *set* of possible objects, rather than a single object as in the concrete evaluation. Since multiple objects are involved, this can result in different method definitions being resolved for the different objects. The method is resolved³ and then applied as follows:

$$\begin{aligned} & \overbrace{\llbracket (\text{invoke-kind } \text{method-name } (\mathfrak{x}_0 \dots \mathfrak{x}_n) (\text{type}_0 \dots \text{type}_n)) \rrbracket : \mathbf{s}, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}}^{\varepsilon} \\ & \rightsquigarrow \widehat{\text{applyMethod}}(m, \mathfrak{x}, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}) \end{aligned}$$

where the function $\widehat{\text{applyMethod}}$ takes a method definition, arguments, a frame pointer, a store, and a new stack pointer and produces the next states:

$$\begin{aligned} & \widehat{\text{applyMethod}}(m, \mathfrak{x}, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}) = (\mathbf{s}, \hat{fp}', \hat{\sigma}', \widehat{a_\kappa}'), \text{ where} \\ & \hat{fp}' = \widehat{\text{allocFP}}(\hat{c}), \quad \widehat{a_\kappa}' = \widehat{\text{allocK}}(\hat{c}), \\ & \hat{\sigma}' = \hat{\sigma} \sqcup [\widehat{a_\kappa}' \mapsto \{\mathbf{fun}(\hat{fp}, \mathbf{s}, \widehat{a_\kappa})\}], \quad \hat{\sigma}'' = \hat{\sigma}' \sqcup [(\hat{fp}', \text{name}_i) \mapsto \hat{\mathcal{A}}(\mathfrak{x}_i, \hat{fp}', \hat{\sigma}')]. \end{aligned}$$

- *Procedure return* Procedure return restores the caller's context and *extends* the return value in the dedicated return register, `ret`.

$$\begin{aligned} & \llbracket (\text{return } \mathfrak{x}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa} \rightsquigarrow (\mathbf{s}', \hat{fp}', \hat{\sigma}', \widehat{a_\kappa}'), \text{ where} \\ & \mathbf{fun}(\hat{fp}', \mathbf{s}', \widehat{a_\kappa}') \in \sigma(\widehat{a_\kappa}) \text{ and } \hat{\sigma}' = \hat{\sigma} \sqcup [(\hat{fp}', \text{ret}) \mapsto \hat{\mathcal{A}}(\mathfrak{x}, \hat{fp}', \hat{\sigma}')]. \end{aligned}$$

³ Since the language supports inheritance, method resolution requires a traversal of the class hierarchy. This traversal follows the expected method and is omitted here so we can focus on the abstract rules.

3 Analysis of reflection and multiple entry points in abstract CESK* machine

While the crux of the analysis for Android apps has been presented in the previous section, our abstract CESK* machine has to be extended to analyze multiple entry points in Android apps and some intricate APIs like reflection.

3.1 Fixed-point computation for multiple entry points

A typical static analysis only deals with one entry point of traditional programs (the **main** method). However, any Android application has more than one entry point, due to the event-driven nature of the Android platform. Intuitively, to explore the reachable states for all the entry points seems to require the exploration for all the permutation of entry points. But this can easily lead to state-space explosion. Related works like [13] prune paths for specific Android apps (but not soundly). We solve the problem in a sound but inexpensive way. Specifically, we iterate over all entry points that have been found. For each entry point, we compute its reachable states via the abstract CESK* machine. Then we compute a single widened store from those states using the widening techniques similar to the ones presented in [15]. The store then forms part of next state to continue the next entry point fixed-point computation. Obviously, the store is monotonic, which ensures a sound approximation of the effects introduced from all entry points. This diminishes precision slightly, but the gains in speed are considerable. The effects of similar technique are also noted in [15][19].

3.2 Reflection

This section presents how to extend the abstract CESK* machine to analyze one of the most commonly used dynamic features in Android—reflection.

Reflection enables programs to access class information to create objects and invoke methods at runtime. Type information involved is dynamically retrieved from strings. The strings can come from user input, files, network or hard-coded, literal strings. Literal strings are not infrequent in reflection. The following code snippet demonstrates a common case of reflection in Java:

```

1 Class<?> aeco = Class.forName("android.os.Environment");
2 Method externalDir = aeco.getMethod("getExternalStorageDirectory", (Class[])null);
3 (File)externalDir.invoke(null);

```

A class object is created in Ln.1 and the method object for *getExternalStorageDirectory*⁴ is created in Ln.2. Finally, the method is invoked in Ln.3

⁴ Since the method is a static method, so no instantiated object is needed, which is null.

via the method object *externalDir*. Since it is a static method with no arguments, the receiver object being invoked is null. Otherwise, the argument *aeco.newInstance()* needs to be supplied in Ln.3.

To analyze such reflection, we can integrate string analysis into abstract interpretation of Java API calls. In the abstract CESK*, we need five additional transition rules, mainly for simple string analysis and the APIs involving creation of class object, method object, class instantiation and method invocation:

- *String*: Strings are objects in Java, and so string instantiation is a special case for the new rule (see Section 2.1)⁵:

$$\overbrace{\llbracket (\text{const-string } name \text{ } str) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}}^{\hat{c}} \rightsquigarrow (\mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a_\kappa}), \text{ where } \widehat{op} = \widehat{allocOP}(\hat{c}).$$

$\hat{\sigma}' = \hat{\sigma} \sqcup [(\hat{fp}, name) \mapsto \{(\widehat{op}, \text{java/lang/String})\}], \hat{\sigma}'' = \hat{\sigma}' \sqcup [(\widehat{op}, \text{value}) \mapsto \alpha(str)]$, Unlike the usual case for new rule, there is a field value paired with the string object pointer as field offset to store abstract string values. α is the abstraction function for string values. The simplest form is to construct a flat lattice for strings. Other string analysis such as Costantini *et al.* [4], Christensen *et al.* [3], *etc.* can be directly incorporated.

- *Class objects*: In byte code, the creation of a class object using `Class.forName` is an `invoke-static` statement, with the first argument referencing to string values. The rule will allocate a new class object on the heap, with the field offset `class-name` points to the string reference looked up by the address (\hat{fp}, \mathfrak{a}) . In addition, the class object reference is stored into the `ret` address:

$$\overbrace{\llbracket (\text{invoke-static java/lang/Class/forName } \mathfrak{a} \text{ java/lang/String) : } \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}}^{\hat{c}} \rightsquigarrow (\mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a_\kappa}), \text{ where } \widehat{op}_{\text{Cls}} = \widehat{allocOP}(\hat{c}),$$

$\hat{\sigma}' = \hat{\sigma} \sqcup [(\widehat{op}_{\text{Cls}}, \text{class-name}) \mapsto \hat{\sigma}(\hat{fp}, \mathfrak{a})]$,
 $\hat{\sigma}'' = \hat{\sigma}' \sqcup [(\hat{fp}, \text{ret}) \mapsto (\widehat{op}_{\text{Cls}}, \text{java/lang/Class})]$

- *Method objects*: Method objects are represented as method headers, including function name, arguments and their types, return values and exceptions that the method can throw.⁶ A method object is resolved from a class object, whose class name can be obtained from the first argument \mathfrak{a}_0 . The second argument will be resolved as the method name. Arrays of argument types of the method object are stored in the third register \mathfrak{a}_3 .⁷

$$\overbrace{\llbracket (\text{invoke-virtual java/lang/Class/getMethod } (\mathfrak{a}_0 \ \mathfrak{a}_1 \ \mathfrak{a}_2) \ \text{types}_{args}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}}^{\hat{c}} \rightsquigarrow \text{newMethodObject}(\widehat{op}_{\text{Method}}, \mathbf{m}, \mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a_\kappa}), \text{ where}$$

⁵ `java/lang/StringBuilder` is interpreted in similar way.

⁶ Exceptions handling is omitted in the semantics.

⁷ We don't interpret the arrays of the types explicitly.

$$\begin{aligned}
 (\widehat{op}_{\text{Cls}}, \text{java/lang/Class}) &\in \hat{\sigma}(\hat{fp}, \mathfrak{x}_0), (\widehat{op}_0, \text{java/lang/String}) \in \hat{\sigma}(\widehat{op}_{\text{Cls}}, \text{class-name}), \\
 \text{class-name} &\in \hat{\sigma}(\widehat{op}_0, \text{value}), (\widehat{op}_1, \text{java/lang/String}) \in \hat{\sigma}(\hat{fp}, \mathfrak{x}_1) \\
 \text{method-name} &\in \hat{\sigma}(\widehat{op}_1, \text{value}), \widehat{op}_{\text{Method}} = \widehat{allocOP}(\hat{c}) \\
 \hat{\sigma}' &= \hat{\sigma} \sqcup [(\hat{fp}, \text{ret}) \mapsto (\widehat{op}_{\text{Method}}, \text{java/lang/Reflect/Method})].
 \end{aligned}$$

Similarly like the transition rule for function call, the method resolution process is omitted here. The resolution process needs the information `class-name` and `method-name`. Also note that the resolution result is a set of public methods \mathbf{m} , rather than one. The helper function `newMethodObject` takes the newly allocated method object pointer, the set of method definitions in the domain `MethodDef`, the rest statements, the frame pointer, store, and the stack pointer and returns the successor states. Again, the method object value will be stored into the `ret` address.

- *Class instantiation:* The API call `java/lang/Class/newInstance` is used to instantiate a new object of a concrete class type (not an abstract class nor interface). The class type name can be resolved from the first argument \mathfrak{x} of the instruction. Unlike the normal `new` statement, the class instantiation requires the invocation of default class constructor. Therefore, we first resolve class definitions by using a helper function $\mathcal{C} : \mathcal{P}(\text{ClassName}) \rightarrow \mathcal{P}(\text{ClassDef})$, and then use `getDefaultConstructor` : `ClassDef` \rightarrow `MethodDef` to get the a constructor method. After that, the control is transferred to the constructor invocation via `invoke-direct` statement, which is inserted in front of the rest states \mathbf{s} .

$$\overbrace{\llbracket (\text{invoke-virtual java/lang/Class/newInstance } \mathfrak{x} \text{ java/lang/Class}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a}_\kappa}^{\hat{c}} \\
 \rightsquigarrow (s' : \mathbf{s}, \hat{fp}, \hat{\sigma}'', \widehat{a}_\kappa), \text{ where}$$

$$\begin{aligned}
 (\widehat{op}_{\text{Cls}}, \text{java/lang/Class}) &\in \hat{\sigma}(\hat{fp}, \mathfrak{x}), (\widehat{op}, \text{java/lang/String}) \in \hat{\sigma}(\widehat{op}_{\text{Cls}}, \text{class-name}), \\
 \text{class-def} &\in \mathcal{C}(\hat{\sigma}(\widehat{op}, \text{value})), \widehat{op}_{\text{Cls}} = \widehat{allocOP}(\hat{c}), \\
 m &= \text{getDefaultConstructor}(\text{class-def}), \\
 s' &= (\text{invoke-direct}, m.\text{method-name } (\mathfrak{x}_0 \dots \mathfrak{x}_n) (\text{type}_0 \dots \text{type}_n)), \\
 \hat{\sigma}' &= \hat{\sigma} \sqcup [(\hat{fp}, \mathfrak{x}_0) \mapsto (\widehat{op}_{\text{Cls}}, \text{java/lang/Class})], \\
 \hat{\sigma}'' &= \hat{\sigma}' \sqcup [(\hat{fp}, \text{ret}) \mapsto (\widehat{op}_{\text{Cls}}, \text{java/lang/Class})].
 \end{aligned}$$

- *Method invocation:* Reflection method invocation in byte code is achieved by invoking the API `java/lang/reflect/Method/invoke` via `invoke-virtual` on a method object, which can be obtained from the first argument. The second argument is the receiver object that the method will be invoked on.⁸ The

⁸ It will be null if the method to be invoked is a static method. Its transition rule can be easily adapted from the rule presented.

third argument is an array of arguments.

$$\overbrace{\llbracket (\text{invoke-virtual java/lang/reflect/Method/invoke } (\mathfrak{x}_0 \ \mathfrak{x}_1 \ \mathfrak{x}_2) \ \text{types}_{args}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}}^{\hat{\epsilon}}$$

$$\rightsquigarrow \widehat{\text{applyMethod}}(m, \mathfrak{x}, \hat{fp}, \hat{\sigma}, \widehat{a_\kappa}).$$

Like the general rule for function call, we have to resolve the method m , and then by using $\widehat{\text{applyMethod}}$ we can get successor states.

4 The tool with user-supplied predicates

In this section, we first briefly presents the implemented analyzer, since the core of the analysis has been specified in Section 2 and 3. Then we illustrate the tool usage, particularly with respect to user-supplied predicates.

The analysis engine is a faithful rendering of the formal specification in Section 2 and 3. In addition, it incorporates previous techniques that boost precision and performance, including the abstract garbage collection [17], store-widening [16], and simple abstract domains to analyze strings [4]. Other constructs of the tool are:

- **Entry points finder:** This component discovers all the entry points of an Android application. Then the engine will explore reachable states based on the algorithm presented in Section 3.
- **Permission violation report:** It reports whether an application asks for more permissions than it actually uses or vice versa.
- **State graph:** This presents all reachable states with states-of-interest highlighted according to default predicates or those supplied by analysts.
- **API dumps:** This presents all the reachable API calls.
- **Heat map:** This reports analyzer intensity on per-statement basis.

The work flow of our human-in-the-loop analyzer—AnaDroid—is as follows: (1) an analyst configures analysis options and malware predicates; (2) AnaDroid presents a permission-usage report, an API call dump, a state graph and a heat map. The major parameters of the analyzer include call-site context-sensitivity— k , widening, abstract garbage collection, cutoffs, and predicates. An analyst can make the trade-off between runtime and precision of the analyzer with these parameters. In addition, the predicates enable analysts to inspect states of interests to detect malware.

4.1 Semantic predicates

To assist analysts, we provide a library of predicates for common patterns. The two major kinds of predicates in AnaDroid are: “State color predicate” renders matching states in a customized color; “State truncate predicate” optimizes the analysis exploration by allowing analysts to manually prune

paths beginning at matching states.

Examples of usage of the predicates are listed as follows:

`uses-API?`: It is used to specify what color to render the state that uses the specified API call. The color is a string representing a SVG color scheme [11], *i.e.* “red, colorscheme=set312”:

```
(lambda (state)
  (if (uses-API? state "org/apache/http/client/HttpClient/execute" st-attr)
      "red,colorscheme=set312"
      #f))
```

Note that `st-attr` is a specialized keyword used by our analyzer. `state` is the parameter of the predicate.

`uses-name?`: It is a variant of `uses-API?`, used to identify a state with the specific method name in code:

```
(lambda (state)
  (if (uses-name? state "org/ucomb/android/testinterface/RectanglePlus/getArea")
      "red,colorscheme=set312"
      #f))
```

An analyst can also use `cond` to specify multiple colors:

```
(lambda (state)
  (cond
    [(uses-API? state "org/apache/http/client/HttpClient/execute" st-attr )
     "red,colorscheme=set312"]
    [(uses-name? state "org/ucomb/android/testinterface /RectanglePlus/getArea")
     "8,colorscheme=set312"]
    [else #f]))
```

`truncate?`:

```
(lambda (state)
  (if (truncate? state "org/apache/http/client/HttpClient/execute")
      "12,colorscheme=set312"
      #f))
```

5 Related work

Stowaway [7] is a static analysis tool identifying whether an application requests more permissions than it actually uses. PScout [1] aims for a similar goal, but produces more precise and fine-grained mapping from APIs to permissions. Our least permission report uses the Stowaway permission map as AnaDroid’s database.⁹

⁹ Our new version analyzer is upgraded to PScout data set.

Jeon *et al.* [12] proposes enforcing a fine-grained permission system. It limits access to resources that could be accessed by Android’s default permissions. Specifically, the security policy uses a white list to determine which resources an app can use and a black list to deny access to resources. In addition, strings potentially containing URLs are identified by pattern matching and constant propagation is used to infer more specific Internet permissions.

Dynamic taint analysis has been applied to identify security vulnerabilities at run time in Android apps. TaintDroid [6] dynamically tracks the flow of sensitive information and looks for confidentiality violations. QUIRE [5], IPCInspection [8], and XManDroid [2] are designed to prevent privilege escalation, where an application is compromised to provide sensitive capabilities to other applications. The vulnerabilities introduced by interapp communication is considered future work. However, these approaches typically ignore implicit flows raised by control structures in order to reduce run-time overhead.

The other approach to enforce security on mobile devices is delegating the control to users. iOS and Window User Account Control [14] can prompt a dialog to request permissions from users when applications try to access resources or make security or privacy-related system level changes. However, we advocate stopping potential malware from floating in the market beforehand via strict inspections. Our tool has designed with analysts in mind and can help them identify malicious behaviors of submitted applications.

6 Conclusion

In this work, we propose a human-oriented semantic-based program analysis for Android apps. We derive an abstract CESK* machine to analyze object-oriented bytecode. Then the foundational analysis is extended to analyze specific features: multiple entry points of Android apps and reflection APIs. We also describe a tool that provides analysts with a library of semantic predicates that can be used to search and filter analysis results, and the ability to dynamically trade speed and precision. It also provides analysts the ability to statically inspect details of every suspicious state of (abstract) execution in order to make a ruling as to whether or not the behavior is truly malicious with respect to the intent of the application. In addition, permission and profiling reports are generated to aid analysts in identifying common malicious behaviors. The technique can systematically, quickly and concisely characterize the behaviors of mobile applications, as demonstrated by case studies in the extended report.

References

- [1] AU, K. W. Y., ZHOU, Y. F., HUANG, Z., AND LIE, D. Pscout: analyzing the android permission specification. In *Proceedings of the 2012 ACM conference on Computer and communications security* (New York, NY, USA, 2012), CCS '12, ACM, pp. 217–228.
- [2] BUGIEL, S., DAVI, L., DMITRIENKO, A., FISCHER, T., SADEGHI, A.-R., AND SHASTRY, B. Towards taming privilege-escalation attacks on Android. In *Proceedings of the 19th Annual Network & Distributed System Security Symposium* (Feb. 2012).
- [3] CHRISTENSEN, A. S., MØLLER, A., AND SCHWARTZBACH, M. I. Precise analysis of string expressions. In *Proceedings of the 10th international conference on Static analysis* (Berlin, Heidelberg, 2003), SAS'03, Springer-Verlag, pp. 1–18.
- [4] COSTANTINI, G., FERRARA, P., AND CORTESI, A. Static Analysis of String Values. In *Proceedings of the 13th International Conference on Formal Engineering Methods (ICFEM 2011)* (2011), S. Qin, Z. Qiu, S. Qin, and Z. Qiu, Eds., vol. 6991 of *Lecture Notes in Computer Science*, Springer, pp. 505–521.
- [5] DIETZ, M., SHEKHAR, S., PISETSKY, Y., SHU, A., AND WALLACH, D. S. Quire: lightweight provenance for smart phone operating systems. In *Proceedings of the 20th USENIX conference on Security* (Berkeley, CA, USA, 2011), SEC'11, USENIX Association, pp. 23–23.
- [6] ENCK, W., GILBERT, P., CHUN, B.-G., COX, L. P., JUNG, J., MCDANIEL, P., AND SHETH, A. N. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX conference on Operating systems design and implementation* (Berkeley, CA, USA, 2010), OSDI'10, USENIX Association, pp. 1–6.
- [7] FELT, A. P., CHIN, E., HANNA, S., SONG, D., AND WAGNER, D. Android permissions demystified. In *Proceedings of the 18th ACM conference on Computer and communications security* (New York, NY, USA, 2011), CCS '11, ACM, pp. 627–638.
- [8] FELT, A. P., WANG, H. J., MOSHCHUK, A., HANNA, S., AND CHIN, E. Permission Re-Delegation: Attacks and defenses. In *Security 2011, 20st USENIX Security Symposium* (Aug. 2011), D. Wagner, Ed., USENIX Association.
- [9] GARTNER. 10 billion android market downloads and counting. <http://googleblog.blogspot.com/2011/12/10-billion-android-market-downloads-and.html>.
- [10] GARTNER. Gartner says sales of mobile devices grew 5.6 percent in third quarter of 2011; smartphone sales increased 42 percent. <http://www.gartner.com/newsroom/id/1848514>.
- [11] GRAPHVIZ. Brewer color schemes. <http://www.graphviz.org/doc/info/colors.html#brewer>.
- [12] JEON, J., MICINSKI, K. K., VAUGHAN, J. A., FOGEL, A., REDDY, N., FOSTER, J. S., AND MILLSTEIN, T. Dr. android and mr. hide: fine-grained permissions in android applications. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices* (New York, NY, USA, 2012), SPSM '12, ACM, pp. 3–14.
- [13] LU, L., LI, Z., WU, Z., LEE, W., AND JIANG, G. Chex: statically vetting android apps for component hijacking vulnerabilities. In *Proceedings of the 2012 ACM conference on Computer and communications security* (New York, NY, USA, 2012), CCS '12, ACM, pp. 229–240.
- [14] MICROSOFT. What's user account control? <http://windows.microsoft.com/en-us/windows-vista/what-is-user-account-control>.
- [15] MIGHT, M. *Environment Analysis of Higher-Order Languages*. PhD thesis, Georgia Institute of Technology, June 2007.
- [16] MIGHT, M. *Environment Analysis of Higher-Order Languages*. PhD thesis, Georgia Institute of Technology, June 2007.
- [17] MIGHT, M., AND SHIVERS, O. Improving flow analyses via Gamma-CFA: Abstract garbage collection and counting. In *ICFP '06: Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming* (New York, NY, USA, 2006), ACM, pp. 13–25.

- [18] PROJECT, T. A. O. S. Bytecode for the dalvik vm. <http://source.android.com/tech/dalvik/dalvik-bytecode.html>.
- [19] SHIVERS, O. G. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 1991.
- [20] VAN HORN, D., AND MIGHT, M. Abstracting abstract machines. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming* (New York, NY, USA, 2010), ICFP '10, ACM, pp. 51–62.