

AnaDroid: Malware Analysis of Android with User-supplied Predicates

Shuying Liang¹ Matthew Might²

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

David Van Horn³

*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, 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. 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 Email: liangsy@cs.utah.edu

² Email: might@cs.utah.edu

³ Email: dvanhorn@ccs.neu.edu

1 Introduction

Google’s Android is the most popular mobile platform, with a share of %52.5 of all smartphones [8]. 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 [7].

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. Or, a note-taking application can wipe out SD card files when a trigger condition is met. Meanwhile, a task manager that requests every possible permission can be legitimately benign.

To understand fine-grained use of granted permissions, 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 (e.g. whether the data has been tampered or not).

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. In particular, our analysis allows analyst-crafted semantic predicates to search and filter analysis results. We describe a tool that provides analysts with a library of such 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. Human-oriented, semantics-based program analysis can systematically, quickly and concisely 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 presents the tool’s high-level architecture, followed by tool usage demonstrated in Section 4. Section 5 shares case studies of using our tool to precisely identify malware. Section 6 discusses related work; and Section 7 concludes.

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

Fig. 1. An object-oriented bytecode adapted from the Android specification [17].

2 Semantic-based program analysis

In this section, we first define an object-oriented byte code language closely modeled on the Dalvik virtual machine to which Java applications for Android are compiled. Then we define an abstract interpretation for the language.

The syntax of the byte code language is given in Figure 1. Statements encode individual actions for the machine; atomic expressions encode atomically computable values; and complex expressions encode expression 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.

With the language in place, the canonical step is to define its concrete semantics, which is the most accurate interpretation of program behaviors, but it is not computable. Fortunately, by adapting the technique of abstracting abstract machine[19], we derive a sound and computable abstract semantics for the Dalvik byte code: the finite-state based abstract CESK* machine.

In this machine, states of this machine consist of a series of statements, a frame pointer, a heap, and continuation address. Abstract evaluation of a program is defined as the set of *abstract* machine configurations reachable by an abstraction of the machine transitions relation, which is the reflexive,

$$\begin{aligned}
 \hat{c} \in \widehat{Conf} &= \mathbf{Stmt}^* \times \widehat{FramePointer} \times \widehat{Store} \times \widehat{KontAddr} \\
 \hat{o} \in \widehat{Store} &= \widehat{Addr} \rightarrow \widehat{Val} \\
 \hat{d} \in \widehat{Val} &= \mathcal{P} \left(\widehat{ObjectValue} + \widehat{String} + \widehat{\mathbb{Z}} + \widehat{KontAddr} \right) \\
 \hat{a} \in \widehat{Addr} &= \widehat{RegAddr} + \widehat{FieldAddr} + \widehat{KontAddr} \\
 \hat{a}_\kappa \in \widehat{KontAddr} &\text{ is a finite set of continuation addresses} \\
 \hat{r}_a \in \widehat{RegAddr} &= \widehat{FramePointer} \times \mathbf{Reg} \\
 \hat{f}_a \in \widehat{FieldAddr} &= \widehat{ObjectPointer} \times \mathbf{FieldName} \\
 \hat{\kappa} \in \widehat{Kont} &= \mathbf{fun}(\hat{fp}, s) + \mathbf{halt} \\
 \hat{o}v \in \widehat{ObjectValue} &= \widehat{ObjectPointer} \times \mathbf{ClassName} \\
 \hat{fp} \in \widehat{FramePointer} &\text{ is a finite set of frame pointers} \\
 \hat{op} \in \widehat{ObjectPointer} &\text{ is a finite set of object pointers.}
 \end{aligned}$$




Fig. 3. Software organization and outputs

Fig. 2. The abstract configuration-space.

transitive closure of the (\rightsquigarrow) relation.

Since the concrete semantics is nearly identical to the abstract semantics in presentation, we will illustrate abstract interpretation of the byte code only, while highlighting places that are different from its concrete counterparts to save space.

2.1 Abstract configuration-space

Figure 2 details the abstract configuration-space. We assume the natural element-wise, point-wise and member-wise lifting of a partial order across this state-space.

The set $\widehat{FramePointer}$ is the environmental component of the machine: by pairing the frame pointer with a register name, it forms the address of its value in the store.

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, the machine may run out of addresses to allocate, and when it does, the pigeon-hole principle will force multiple abstract values to reside at the same 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 configuration-space.

2.2 Abstract transition relation

The machine relies on helper functions for evaluating atomic expressions, and looking up field values:

- $\mathcal{I} : \text{Stmt}^* \rightarrow \widehat{\text{Conf}}$ injects an sequence of instructions into a configuration:
 $\hat{c}_0 = \hat{\mathcal{I}}(\mathbf{s}) = (\mathbf{s}, \hat{fp}_0, \hat{a}_{\kappa 0}, \langle \rangle)$
- $\hat{\mathcal{A}} : \text{AExp} \times \widehat{\text{FramePointer}} \times \widehat{\text{Store}} \rightarrow \widehat{\text{Val}}$ evaluates atomic expressions (specifically for variable look-up): $\hat{\mathcal{A}}(\text{name}, \hat{fp}, \hat{\sigma}) = \sigma(\hat{fp}, \text{name})$
- $\hat{\mathcal{A}}_{\mathcal{F}} : \text{AExp} \times \widehat{\text{FramePointer}} \times \widehat{\text{Store}} \times \text{FieldName} \rightarrow \widehat{\text{Val}}$ looks up fields:

$$\begin{aligned} \hat{\mathcal{A}}_{\mathcal{F}}(\mathbf{x}, \hat{fp}, \hat{\sigma}, \text{field-name}) &= \bigsqcup \hat{\sigma}(\widehat{op}, \text{field-name}) && \text{[field look-up]} \\ \text{where } (\widehat{op}, \text{class-name}) &\in \hat{\mathcal{A}}(\mathbf{x}, \hat{fp}, \hat{\sigma}). \end{aligned}$$

The rules for the abstract transition relation $(\leadsto) \subseteq \widehat{\text{Conf}} \times \widehat{\text{Conf}}$ models all the possible concrete execution of a byte code program in a sound way. In the subsequent subsections, we illustrate the important rules that involves objects and function calls, omitting less important ones to save space.

2.2.1 New object creation

Creating an object allocates a potentially non-fresh object pointer and joins the newly initialized object into that store location:

$$\begin{aligned} &\overbrace{(\llbracket (\text{assign name (new class-name)}) : \mathbf{s} \rrbracket, \hat{fp}, \hat{\sigma}, \widehat{a}_{\kappa})}^{\hat{c}} \Rightarrow (\mathbf{s}, \hat{fp}, \hat{\sigma}', \widehat{a}_{\kappa}), \\ &\quad \widehat{op}' = \text{allocOP}(\hat{c}) \\ &\quad \hat{\sigma}' = \hat{\sigma} \sqcup [(\hat{fp}, \text{name}) \mapsto (\widehat{op}', \text{class-name})] \\ &\quad \hat{\sigma}'' = \text{initObject}(\hat{\sigma}', \text{class-name}), \end{aligned}$$

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

2.2.2 Instance field reference/update

Referencing a field uses $\hat{\mathcal{A}}_{\mathcal{F}}$ to evaluate the field values and join the store for destination register:

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

Updating a field first finds the abstract object values from the store, extracts its object pointer from each of all the possible values, then pairs this object pointer with the field name to get the field addresses, and finally *joins* the extensions to the store:

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

2.2.3 Method invocation

This rule involves all four components of the machine. In abstract interpretation of non-static method invocation, there can be a *set* of possible objects that are invoked, rather than only one as in its concrete counterpart. This also means that there could be multiple method definitions resolved for each object. For each such method m ⁴:

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

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

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

2.2.4 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 High-level architecture of the tool

We have constructed our analysis engine based on Section 2. Its high level architecture is shown in Fig 3.

Anadroid is the analysis engine. It is a faithful rendering of the formal specification in Section 2. In addition, it incorporates previous techniques that boost precision and performance, including the abstract garbage collection [14], store-widening [13], and simple abstract domains to analyze strings [2]. Other constructs of the tool are:

⁴ Since the language supports inheritance, method resolution requires a traversal of the class hierarchy. This traversal is not of interest, so we focus on the abstract rules.

- **Multi-entry points:** The component `Entry points finder` in `Anadroid` discovers all the entry points of an Android application. A typical static analysis only deals with one entry point of traditional programs (the `main` method). However, any Android application has one or more entry-points, 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 [11] prune paths for specific Android applications (but not soundly). While we solve the problem in a sound but not expensive way.

Specifically, we iterate all the entry points. For each entry point, we compute its reachable states. Before exploring the next entry point, the store component of the next state “inherits” the *widened* store computed from states of last fixed point.

- **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 that or those supplied by analysts.
- **API dumps:** This presents all the reachable APIs calls.
- **Heat Map:** This reports analyzer intensity on per-statement basis.

4 Tool usage

This section presents the work flow of our human-in-the-loop analyzer: `AnaDroid`. We have deployed `AnaDroid` as a web service, accessible via: <http://pegasus.cs.utah.edu:8080/anadroid>.

The `AnaDroid` analysis process is as follows: (1) an analyst configures analysis options and malware predicates; (2) `AnaDroid` presents a permission-usage report, an API call dumps, 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 Predicates

To assist analysts, there are two kinds of predicates in `AnaDroid`: “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. We provide a library of these kind

of predicates for common patterns. 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 [9], e.g. “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 state with specific method name of from source code:

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

- **truncate?**:

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

5 Case studies

In this section, we illustrate two case studies⁵ to demonstrate the effectiveness of our human-in-the-loop analyzer.

5.1 Case study 1: user location leaked to a malicious website

UltraCoolMap is advertised to be an application allowing users to browse an interactive map, mark, save, load marked locations, and get directions between marked locations. However, it is malicious. It exfiltrates users’ private GPS data by means of an HTTP get request.

The *UltraCoolMapMaliciousCodeSnippet.java* lists the minimum malware code region identified by our tool. In the entry point function **onCreate**(in *UltraCoolMapActivity.java*), when a button is clicked, a typical path constructs the Google map address with location data in variable **real_badlyName** (com-

⁵ The two applications are challenge applications from DARPA’s APAC program.

mented out in line 6), but then from line 12-16, the app replaces part of the address with “maps.Google-com.cc” via the function `real_Bad_Name`. Lastly, the URI parameter, which is now the location data along with the modified malicious website, flows to the HTTP get request in line 21 in background via an instantiated asynchronous task in line 10.

Fig 4 is a screen snapshot of the related state graph generated by our tool, with vulnerabilities flagged in red and states with sensitive, private-data-bearing strings in orange. For any state, one can click into the node to inspect detailed state information, which includes (1) Specific API as well as correspondent permissions for the state; (2) Any sensitive values (normally for strings); (3) Current function frame pointer, time, and successive instructions; (4) The store environment.

An analyst typically starts from any highlighted node, e.g. 470, 453, 447, and then following forward/backward along the path to determine any vulnerabilities. In addition, one can find the direct trigger of a suspicious functionality by going to the first node of the path. Even though the orders of entry points are not deterministic, by inheriting the store environment for information propagation, we can track whether any sensitive value flows to a program point or not in a sound way, which proved to sufficient for finding malware in a red-team engagement in DARPA’s APAC program.

In the case of *UltraCoolMap*, an analyst can find that a HTTP request instantiated with some URI with sensitive values in the node, it is then executed in node 447, and the direct trigger is the `doInBackground` method. By checking paths across different entry points, one can make more sense of the application.

```

1 public void onCreate(Bundle realBadlyName) {
2     //...
3     real_badly_name.setOnClickListener(new OnClickListener() {
4         @Override
5         public void onClick(View v){
6             // built part of maps.google.com in real_badlyName
7             try {
8                 URI realbadlyname = new URI(this
9                     .real_Bad_Name(real_badlyName));
10                new ReallyBadName().execute(realbadlyname); }
11            catch (Exception realbadly_name){} }
12        private String real_Bad_Name(String really_bad_Name){
13            // string manipulation to replace with "maps.google-
14              com.cc"
15            really_BadName = really_BadName.concat("-");
16            really_BadName = really_BadName.concat(".cc");
17            return really_BadName;}});
18
19    private class ReallyBadName extends AsyncTask<URI, Void, Void>
20    {
21        protected Void doInBackground(URI... uris){
22            HttpClient reallyBadName = new DefaultHttpClient();

```

```

21      HttpGet reallybadName = new HttpGet(uris[0]);
22      try { reallyBadName.execute(reallybadName);}
23      catch (Exception really_bad_name){}
24      return null; }}}

```

UltraCoolMapMaliciousCodeSnippet.java

5.2 Case study 2: SD card files deleted

TomDroid is a note taking system that lets users save and synchronize notes. Our experiment can precisely find suspicious behavior efficiently. Malicious code region identified by the tool is listed as follows:

```

1  // called from onCreate method
2  private void displayHundredRuns(){
3      // update sdcard record to remind about 100 runs
4      File sddir = new File("/sdcard");
5      if(sddir.isDirectory() && sddir.canRead()) {
6          //recurse is a recursively gathering files in folder
7          List<String> listing = recurse(sddir);
8          if(listing != null && listing.size() > 0) {
9              for(String entry : listing)
10                 updateContentOnRuns(100,new File(entry));}}}
11
12  public void updateContentOnRuns(long runCount, File filePath){
13      if (runCount < 100)
14          filePath.setLastModified(System.currentTimeMillis());
15      else {
16          if (runCount >= 0)
17              filePath.delete();
18          else
19              filePath.setLastModified(System.currentTimeMillis());
20      }}

```

SDCardFilesDeleted.java

In the code snippet *SDCardFilesDeleted.java*, when activity’s `onCreate` method is called by Android framework, it will call `displayHundredRuns` (defined in line 2-10), which traverses the folder and calls `updateContentOnRuns` method (line 10), where in line 17 deletes the iterating file object when `runCount < 100`.

For this is kind of malware, our analyzer presents a “information flow” style graph to help analysts to precisely identify the malware functionality as shown in Fig 5.

The graph only presents reduced-size nodes and path. Node 48 on the left hand side signifies a file object is instantiated. However, with this information alone can not justify any malware behavior. By following the node 47-42, an analyst can see possible sensitive values (“sdcard”) are used, and so we know that the app is trying to access to the external storage of the phone. The right




Fig. 4. Snapshot of state graph for UltraCoolMap




Fig. 5. Snapshot of state graph for case study 2

hand side is the continuous path of the graph. It only presents the neighboring nodes of the highlighted node 1. One can inspect into the node 1 for more detail (like Fig 4), knowing that the “delete” method is invoked on the file object along the path⁶.

6 Related work

Stowaway [5] is a static analysis tool exploring whether an application requests more permissions than it actually uses. Our least permission report uses Stowaway’s permission map as Android’s database. However, the main difference is that Stowaway adapted testing methodology to test applications and identify APIs that require permissions, while our approach annotates APIs with permissions and statically analyses all executable paths.

Language-based information flow [18][15] [16] allows developers to annotate variables with security attributes and compilers use the attributes to enforce information control flows. The main concern of this approach is that it imposes additional burdens on developers whose major focus and interest is business logic.

Recent techniques to enforce fined-grained permission system proposed by Jeon *et al.* [10] are also targeted developers (but can be retrofitted by sophisticated users). The idea of “fine-grained” permission is mainly achieved

⁶ Again, the looping path is not showing in the figure due to space limit.

by limiting access to resources that could be accessed by Android’s default permission system. Specifically, the security policy is designed to allow access to resources that are specified in white list while deny access to those in blacklist. In addition, sensitive strings (URL-like) are pattern matched and then traditional constant propagation is used to infer permissions.

Dynamic taint analysis has been applied to identify specific security vulnerabilities at run time in Android applications. TaintDroid [4] dynamically tracks the flow of sensitive information and looks for confidentiality violations. QUIRE[3], IPCInspection[6] and XManDroid [1] are designed to prevent privilege-escalation, where an application is compromised to provide sensitive capabilities to other applications. The vulnerabilities introduced by inter-app communication will be considered as our future work. However, these approaches typically ignore implicit flows raised by control structures in order to reduce runtime overhead. Moreover, dynamically executing all execution paths of these applications to detect potential information leaks is impractical. The limitations make these approaches inappropriate for computing information flows for all submitted applications.

Another approach to enforce security control on mobile devices is delegating the control to users themselves. iOS and Window User Account Control [12] can prompt a dialog to request permissions from users when applications try to access resources or make security or privacy-related system level changes. Depending on users to enforce security control is putting users at their own risks. Notifications prompts by the tools usually provide no insights of how users’ private sensitive data are used, and users tend to grant permissions in order to install the apps that want. Thus, it is desirable to stop 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 fast and efficiently.

7 Conclusion

In this work, we propose a human-oriented semantic-based program analysis with the ability to allow analyst-crafted semantic predicates to search and filter the result of analysis result. We describe our tool that provides analysts with a library of such 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. The technique can systematically, quickly and concisely characterize the behaviors of mobile applications, as demonstrated by our case studies.

References

- [1] 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).
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] GARTNER. 10 billion android market downloads and counting. <http://googleblog.blogspot.com/2011/12/10-billion-android-market-downloads-and.html>.
- [8] 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>.
- [9] GRAPHVIZ. Brewer color schemes. <http://www.graphviz.org/doc/info/colors.html#brewer>.
- [10] 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.
- [11] 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.
- [12] MICROSOFT. What’s user account control? <http://windows.microsoft.com/en-us/windows-vista/what-is-user-account-control>.
- [13] MIGHT, M. *Environment Analysis of Higher-Order Languages*. PhD thesis, Georgia Institute of Technology, June 2007.
- [14] 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.
- [15] MYERS, A. C. Jflow: practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages* (New York, NY, USA, 1999), POPL ’99, ACM, pp. 228–241.
- [16] MYERS, A. C., AND LISKOV, B. Protecting privacy using the decentralized label model. *ACM Trans. Softw. Eng. Methodol.* 9, 4 (Oct. 2000), 410–442.
- [17] PROJECT, T. A. O. S. Bytecode for the dalvik vm. <http://source.android.com/tech/dalvik/dalvik-bytecode.html>.

- [18] SABELFELD, A., AND MYERS, A. C. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* 21 (2003), 2003.
- [19] VAN HORN, D., AND MIGHT, M. Abstracting abstract machines. In *ICFP '10: Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming* (2010), ICFP '10, ACM Press, pp. 51–62.