



# Networking and Security Research Center

## *Technical Report*

---

NAS-TR-0140-2010

## The ded Decompiler

Damien Oceau and William Enck and Patrick McDaniel

23 September 2010 (updated 10 May 2011)

The Pennsylvania State University  
344 IST Building  
University Park, PA 16802 USA  
<http://nsrc.cse.psu.edu>

©2011 by the authors. All rights reserved.

# The ded Decompiler

Damien Oceau, William Enck and Patrick McDaniel  
Systems and Internet Infrastructure Security (SIIS) Laboratory  
Department of Computer Science and Engineering  
The Pennsylvania State University  
{oceau, enck, mcdaniel}@cse.psu.edu

## ABSTRACT

Smartphone applications are frequently incompletely vetted, poorly isolated, and installed by users without restraint. Such behavior is fraught with peril: applications containing malicious logic or critical vulnerabilities are likely to be identified only after substantial damage has already occurred. Unfortunately, the limitations of application markets make them a poor agent for certifying that applications are secure. This paper presents a certification process that allows the consumers of applications to validate applications' security directly. Built for the Android mobile phone platform, we reverse engineer downloaded application images into application source code and use static analysis to detect vulnerabilities. We develop and document a multi-stage process for VM retargeting and code recovery. A study of the top 1,100 free Android market applications recovers source code for over 95% of the 143 thousand class files containing over 12 million lines of code. A preliminary analysis of the recovered source code identified over 3,100 potential vulnerabilities involving a broad range of program features.

## 1. INTRODUCTION

The explosion of smartphones has led to a cottage industry in phone applications. Developers use phone APIs to build applications sold for as little as 99¢ through *application markets*. For example, the Apple iPhone [1] and Android [24] platforms provide easy to use on-phone interfaces for browsing literally thousands of applications. Users simply click a button to download and install an application. For technical and logistical reasons, existing application markets do little to detect malicious or vulnerable applications. Thus, “*caveat emptor*”—users have little information on which to judge the trustworthiness of downloaded applications. Mispaced trust in these applications has left users' phone data and interfaces vulnerable to malicious code [26, 22].

In response to the limitations of current markets, recent experimental systems have explored other means of vetting smart phone applications. The Kirin system [23] inspects

the permissions associated with each application as it is installed. Based on a formal analysis of a program's manifest, applications that govern sensitive interfaces poorly or attempt or request unsafe sets of permissions are rejected. In contrast, the Saint system [34] seeks to regulate how smartphone applications share information by enforcing safe workflows between coordinated applications through mandatory system policies. These techniques are limited in that focus solely on the analysis of policy associated with a particular phone and application. Identifying vulnerabilities or malicious misuse of permissions is not within the scope of their analysis. Thus, otherwise valid and correctly regulated applications may, for example, leak sensitive data [26].

This paper explores an alternative to current smart phone application market analysis. In this model, platform developers, users, cellular providers, consumer protection agencies, enterprises, professional software certification consultants, and anyone else with an interest in the security of the phones are free to certify the applications meet whatever security criteria they choose. This model changes the nature of the market by altering the oft-broken incentives—developers and markets will be encouraged to create secure code and consumers will have the tools to detect and ultimately punish those who fail to do so. Users will “vote with their feet” by avoiding developers with poor security practices and gravitate towards those with good ones.

Our certification approach automates the reverse engineering process by extracting the application source *solely from its installation image*. Thereafter, static analysis tools are used to detect potential vulnerabilities based on the evaluator's criteria. Our novel *ded* tool retargets the Dalvik virtual machine bytecode of candidate Android applications into Java `.class` files. The application's original Java source code is then recovered from the retargeted bytecode using existing Java tools (i.e., Soot [38]). Finally, commercial off-the-shelf (COTS) static analysis tools are applied to validate security related invariants defining safe phone behavior. Applications adhering to these invariants are deemed safe and can be installed, and those that fail are rejected.

We make the following contributions in this paper:

- *We develop an automated system for Android application certification* - We develop techniques for extracting source code from installation images and leverage existing analysis tools to measure the code's adherence to desired security policies. While we explore the use of tools for analysis, the majority of the technical research contributions focus on algorithms and tools for retargeting Dalvik bytecode into valid class files.

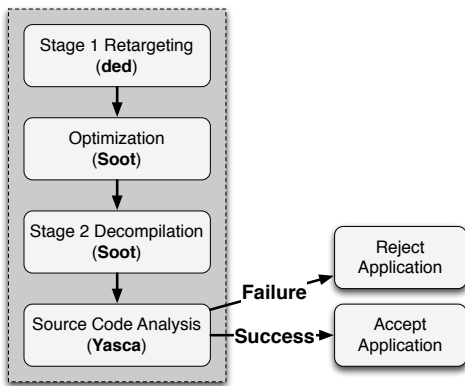


Figure 1: Install-time Application Certification

- We demonstrate the viability of the approach on existing market applications - Our study of the recent top 1,100 free Android applications recovered 12 million SLOC in 143 thousand Java classes—a 95% success rate. The analysis tools flag over 3,100 locations within code as potential vulnerabilities. A preliminary analysis of the vulnerabilities shows that the flags represent both common security failures and false-positives.

Note that the certification process is not a replacement for policy analysis techniques such as Kirin or Saint, but is complementary to them. Such tools inspect or regulate how applications can exercise system permissions. We aim to detect vulnerable interfaces and code, as well as malicious behaviors within the applications themselves. Thus, prior tools sought to analyze different aspects of the system than those discussed here, the former relating to the application behaviors within the phone and the latter about the structure and behaviors of the application itself.

The remainder of this paper proceeds as follows. Section 2 overviews our certification approach. Section 3 describes the differences between Dalvik and Java bytecode, highlighting the challenges that make retargeting difficult. Section 4 describes our retargeting algorithms. Section 5 uses existing source code certification tools to analyze decompiled applications downloaded from the Android Market. Section 6 discusses related work. Section 7 concludes.

## 2. CERTIFICATION APPROACH

Illustrated in Figure 1, certification consists of two phases; *a)* source code recovery via multi-stage decompilation, and *b)* analysis of recovered code against safety-preserving invariants. Java bytecode decompilation has been studied since its introduction in the 1990s. Tools such as Mocha [10] date back over a decade, with many other techniques being developed [35, 30, 29, 8, 7, 5] (see Section 6). Unfortunately, there exists no functional tool for the Dalvik bytecode<sup>1</sup>. As described in the next section, the JVM and DVM are vastly different architectures, thus simple modification of existing decompilers was not possible. Discussed in detail in the following sections, we developed *ded*<sup>2</sup> for this purpose. Rather

<sup>1</sup>The undx and dex2jar tools attempt to decompile .dex files. As discussed in Section 6 the tools are currently are incapable of reliably recovering source code.

<sup>2</sup>While the dex decompiler (*ded*) only performs retargeting,

than directly recover Java source code from Dalvik bytecode, *ded* retargets Dalvik .dex files to Java .class files. The *ded* retargeting process yields complex, unoptimized Java bytecode. We use the Soot framework [38] to optimize the *ded* output, rendering Java bytecode from which the application source code can be recovered. The .class files are then decompiled using the optimized bytecode also using Soot.

The output of the decompilation process is the application source code. At this stage, a certifying party can use any code analysis tool at their disposal. In this paper, we use the Yasca (Yet Another Source Code Analyzer [15]) tool with the PMD plugin [11] to evaluate the recovered source code. These tools search the code for poor coding practices, errors, and exploitable code. In the study presented in Section 5, with one exception, we use default tests that identify simple security issues. However, the tools are designed to allow the user to create custom analyses testing—thus allowing them to enforce whatever security criteria they deem necessary (within the limitations of the analysis process).

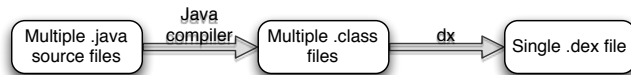
Note that the substance of the “security-preserving invariants” validated via analysis is critical to certification; it defines the security policy being enforced. This paper does not attempt to identify what a good application policy is, as it is a complex context-sensitive issue. We only seek to demonstrate that the code recovery and subsequent analysis process is viable. In practice, we expect that the kinds of analysis needed to vet application code will be specific to the platform, the users, and the security goals of the evaluator.

## 3. THE DALVIK VIRTUAL MACHINE

The Dalvik Virtual Machine (DVM) was designed for resource constrained Android mobile phones. While the software run in the DVM is compiled from Java source code, its bytecode and run-time environment differ substantially from that for existing JVMs. We highlight several key differences between the bytecode formats used by the JVM and DVM as they relate to code recovery. We simplify the description of the two formats for ease of exposition.

*Application Structure* - Java applications are composed of one or more .class files, one file for each class. The JVM loads the bytecode in the .class file associated with each Java class as it is referenced using the Java class loader at run time. Conversely, a Dalvik application consists of a single .dex file containing all classes composing the application. The entirety of the application is loaded when it is launched by the DVM.

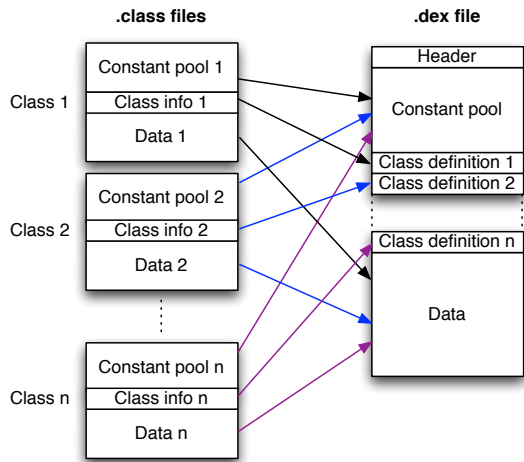
To support this, Android extends the compilation process:



Here, the Java compiler operates normally to produce a collection of .class files. The Dalvik dx compiler then consumes the classes, recompiles them to Dalvik bytecode, and writes the resulting application into a single .dex file.

Figure 2 provides a conceptual view of the dx .class to .dex compilation process. This process consists of the translation, reconstruction, and interpretation of three basic elements of the application: the constant pools, the class def-

it is named such for historical reasons. Together, *ded* and Soot provide decompilation of Android applications.



**Figure 2: dx compilation of Java classes to a DVM application (simplified view).**

initions, and the data segment. A constant pool describes, not surprisingly, the constants used by a class. This includes among other items references to other classes, method names and numerical constants. The class definitions consist in the basic information such as access flags and class names. The data element contains the method code executed by the target VM, as well as other information related to methods (e.g., number of DVM registers used, local variable table and operand stack sizes) and to class and instance variables.

*Register architecture* - The DVM is register-based, whereas existing JVMs are stack-based. Java bytecode can assign local variables to a local variable table before pushing them onto an operand stack for manipulation by opcodes, but it can also just work on the stack without explicitly storing variables in the table. Dalvik bytecode assigns local variables to any of the  $2^{16}$  available registers. The Dalvik opcodes directly manipulate registers, rather than accessing elements on a program stack. For example, Figure 3 shows how a typical operation (“add” function Figure 3(a)) is implemented in Java (stack) versus Dalvik (register) virtual machine. The Java bytecode shown in Figure 3(b) pushes the local variables `a` and `b` onto the operand stack using the `iload` opcode, and returns the result from the stack via the `ireturn` opcode. By comparison, the Dalvik bytecode shown in Figure 3(c) simply references the registers directly.

*Instruction set* - The Dalvik bytecode instruction set is substantially different than that of Java: 218 opcodes vs. 200, respectively. The nature of the opcodes is very different: for example, Java has tens of opcodes dedicated to pushing and pulling elements between the stack and local variable table. Obviously the Dalvik bytecode instruction set does not require any such opcodes. Moreover, as illustrated in the example in Figure 3, Dalvik instructions tend to be longer than Java instructions as they include the source and destination registers (when needed). As a result, Dalvik applications require fewer instructions. Applications encoded in Dalvik bytecode have on average 30% fewer instructions than in Java, but have a 35% larger code size (bytes) [17]. This increased code size has limited impact on performance, as the DVM reads instructions by units of two bytes.

```
public int add(int a, int b)
{
    return a + b;
}
```

(a) Source Code

```
public int add(int, int)
0:  iload_1
1:  iload_2
2:  iadd
3:  ireturn
```

(b) Java (stack) bytecode

```
public int add(int, int)
0:  add-int v0,v2,v3
2:  return v0
```

(c) Dalvik (register) bytecode

**Figure 3: Register vs. Stack Opcodes**

*Constant pool structure* - Java applications necessarily replicate elements in constant pools within the multiple `.class` files, e.g., referrer and referent method names. The `dx` compiler attempts to reduce application size by eliminating much of this replication. Dalvik uses a single large constant pool that all classes simultaneously reference. Additionally, `dx` eliminates some constants by inlining their values directly into the bytecode. In practice, pool elements for integers, long integers, and single and double precision floating-point constants simply disappear as bytecode constants during the transformation process.

*Control flow Structure* - Programmatic control flow elements such as loops, switch statements and exception handlers are structured very differently by Dalvik and Java bytecode. To simplify, Java bytecode structure loosely mirrors the source code, whereas Dalvik bytecode does not. The restructuring is likely performed to increase performance, reduce code size, or address changes in the way the underlying architecture handles variable types and registers.

*Ambiguous typing of primitive types* - Java bytecode variable assignments distinguish between integer (`int`) and single-precision floating-point (`float`) constants and between long integer (`long`) and double-precision floating-point (`double`) constants. However, Dalvik constant assignments (`int/float` and `long/double`) use the same opcodes for integers and floats, e.g., the opcodes are untyped beyond specifying precision. This complicates decompilation of Dalvik bytecode because the variable type is not indicated by its declaration. Thus, the decompilation process must observe a variable creation and inspect its subsequent use to infer its type to create accurate correct Java bytecode and constant pools. This is a specific instance of a broad class of widely-studied *type inference* problems. Note that an incorrect inference (and thus type) may result in incorrect behavior (and analysis) of the decompiled program.

*Null references* - The Dalvik bytecode does not specify a `null` type, instead opting to use a zero value constant. Thus, constant zero values present in the Dalvik bytecode have ambiguous typing that must be recovered. The decompilation process must recover the `null` type by inspecting the variable’s use *in situ*. If the `null` type is not correctly recovered, the resulting bytecode can have illegal integer zero

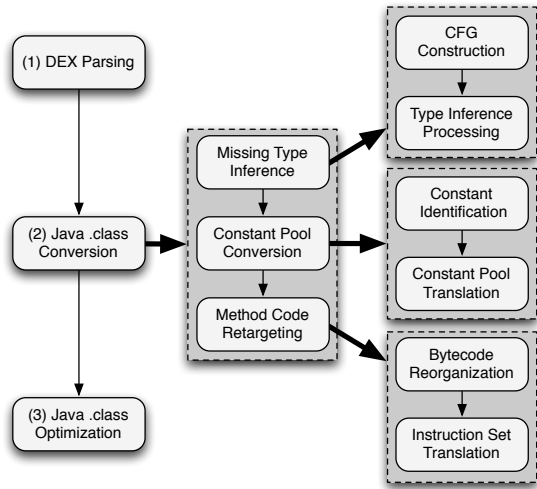


Figure 4: Dalvik bytecode retargeting

assignments to object references, and vice-versa.

*Comparison of object references* - The Java bytecode uses typed opcodes for the comparison of object references (`if_acmpeq` and `if_acmpne`) and for null comparison of object references (`ifnull` and `ifnonnull`). The Dalvik bytecode uses a more simplistic integer comparison for these purposes—respectively the comparison between two integers and the comparison of integer to zero. This requires the decompilation process to recover types for integer comparisons. Again, incorrect or missed inferences could result in illegal bytecode, typing violations, and ultimately inaccurate decompiled source code.

## 4. APPLICATION RETARGETING

The initial stage of decompilation is to retarget the application `.dex` file to Java classes. Figure 4 shows the three central research challenges we faced in this process: recovering typing information, translating the constant pool, and retargeting the Dalvik bytecode. For brevity, we focus on only the most challenging issues that relate to `ded`.

### 4.1 Type Inference

The first step in retargeting is to identify class and method constants and variables. However, the Dalvik bytecode does not always provide enough information to determine the type of a variable or constant from its register declaration. There are two generalized cases where variable types are ambiguous: 1) constant and variable declaration only specifies the variable width (e.g., 32 or 64 bits), but not whether it is a float, integer, or null reference; and 2) comparison operators do not distinguish between integer and object reference comparison (i.e., null reference checks).

Type inference has been widely studied [37]. The seminal Hindley-Milner [31] algorithm provides the basis for type inference algorithms used by many languages such as Haskell and ML. Broadly speaking, the approach is to determine unknown types by observing how variables are used in operations with known type operands. Other languages such as OCAML [18] perform strong type inference and others like Perl [39] employ weaker algorithms using similar techniques.

`ded` adopts the accepted approach: it infers register types by observing how they are used in subsequent operations with known type operands. Dalvik registers loosely correspond to Java variables. Because Dalvik bytecode reuses registers whose variables are no longer in scope, we must evaluate the register type within its context of the method control flow, i.e., inference must be *path-sensitive*. Note further that `ded` type inference is also *method-local*. Because the types of passed parameters and return values are identified by method signatures, there is no need to search outside the declaring method.

There are three ways `ded` infers a register’s type. First, any comparison of a variable or constant with a known type exposes the type. Comparison of dissimilar types requires type coercion in the original Java code, which is propagated to the Dalvik bytecode. Hence any legal Dalvik comparison must involve registers of the same type. Second, instructions such as `add-int` only operate on specific types, and manifestly expose typing information. Third, any instruction that passes a register another method or assigns it to a return value exposes the type via the method signature.

The `ded` type inference algorithm proceeds as follows. After reconstructing the control flow graph, `ded` identifies any ambiguous register declaration. For each such register, `ded` walks the instructions in the control flow graph starting from its declaration. Each branch of the control flow encountered is pushed onto a inference stack, e.g., `ded` performs a depth-first search of the control flow graph looking for type-exposing instructions. If a type-exposing instruction is encountered, the variable is labeled and the process is complete for that variable. There are three events that cause a branch search to terminate: a) when the register is re-assigned to another variable (e.g., a new declaration is encountered), b) when a return function is encountered (all DVM methods terminate with a return), and c) when an exception is thrown. After a branch is abandoned, another is popped off the stack and the search continues. Lastly, type information is forward propagated, modulo register re-assignment, through the control flow graph from each register declaration to all subsequent ambiguous uses. This algorithm resolves all ambiguous primitive types, except for one isolated case when all paths leading to a type ambiguous instruction originate with ambiguous constant instructions (e.g., all paths leading to an integer comparison originate with registers assigned a constant zero). However, in this case, the exact type does not impact decompilation, and we can safely assign a default type (e.g., integer).

Note that it is sufficient to find *any* type-exposing instruction for a given register assignment. Any two instructions that expose different types for the same register would represent illegal bytecode. If this were to occur, the primitive type would be dependent on the path taken at run time, a clear violation of Java’s type system.

### 4.2 Constant Pool Conversion

There are two central differences between `.dex` and `.class` file constant pools: 1) Dalvik maintains a single constant pool for the application and Java maintains one for each class, and 2) Dalvik bytecode places primitive type constants directly in the bytecode, whereas Java bytecode uses the constant pool for most references. We convert constant pool information in two steps, as shown in Figure 4.

The first step is to identify which constants are needed

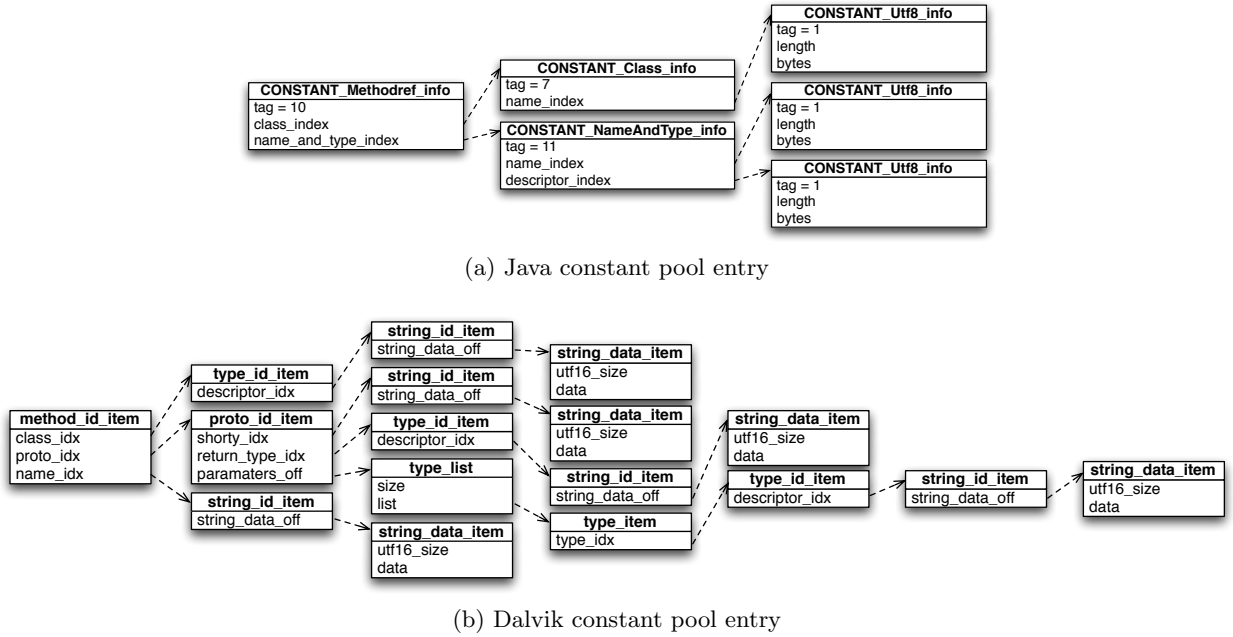


Figure 5: Constant pool entry structure for a method reference

Table 1: Example Dalvik to Java bytecode translation rules

Dalvik Bytecode Instructions	Java Bytecode Instructions	Description
<code>add-int</code> $d_0, s_0, s_1$	<code>iload</code> $s_0$ <code>iload</code> $s_1$ <code>iadd</code> <code>istore</code> $d_0$	Integer addition
<code>invoke-virtual</code> $s_0, \dots, s_k, m_r$ <code>move-result</code> $d_0$	<code>iload</code> $s_0$ ... <code>iload</code> $s'_i$ <code>invokevirtual</code> $m'_r$ <code>istore</code> $d_0$	Virtual method invocation with assigned return value.

for a `.class` file. Constants include references to classes, methods, and instance variables. `ded` traverses the bytecode for each method in a class, noting such references. `ded` also identifies all constant primitives specified in the Dalvik bytecode. Here, `ded` notes the types determined during type inference, described in Section 4.1.

Once `ded` identifies the constants required by a class, it adds them to the target `.class` file. For primitive type constants, new entries are created. For class, method, and instance variable references, we must populate the Java constant pool entry based on information available in the Dalvik constant pool. The two constant pool entries differ in complexity. Specifically, Dalvik constant pool entries use significantly more references to reduce memory overhead.

Figure 5 depicts the method entry constant in both Java and Dalvik formats. Other constant pool entry types have similar structures. Each box is a data structure. Index entries (denoted as “idx” for the Dalvik format) are pointers to a data structure. The Java method constant pool entry, Figure 5(a), provides three strings: 1) the class name, 2) the method name, and 3) a descriptor string representing the argument and return types. The Dalvik method constant pool entry, Figure 5(b), also contains these strings, but uses

more indirection. `ded` ignores Dalvik-specific entries such as “shorty” strings used as simplified method descriptors.

### 4.3 Method Code Retargeting

The final stage of the retargeting process is the translation of the method code. This is a two stage process, as shown in Figure 4. First, we preprocess the bytecode to reorganize structures that cannot be directly retargeted. Second, we linearly translate DVM bytecode to the JVM.

The preprocessing phase considers multidimensional arrays. Both Dalvik and Java use blocks of bytecode instructions to create multidimensional arrays; however, the instructions have different semantics and layout. `ded` reorders and annotates the bytecode with array size and type information to allow linear instruction translation.

The bytecode translation linearly processes each Dalvik instruction. First, `ded` maps each referenced register to a Java local variable table index. Second, `ded` performs an instruction translation for each encountered Dalvik instruction. As Dalvik bytecode is more compact and takes more arguments, one Dalvik instruction frequently expands to multiple Java instructions. Third, `ded` patches the relative offsets used for branches based on preprocessing anno-

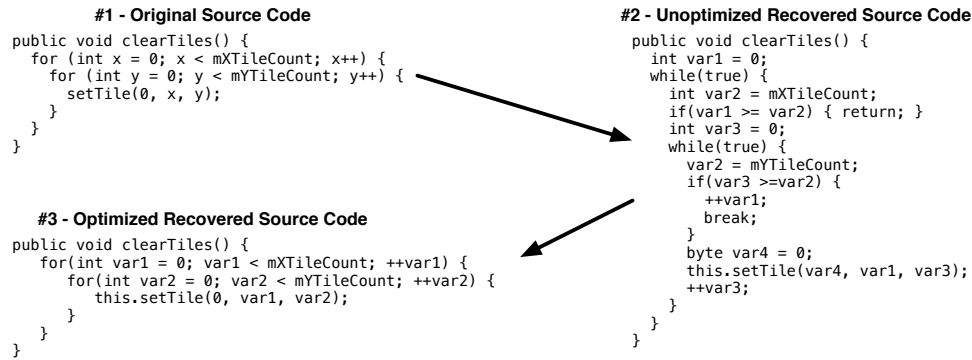


Figure 6: Original, recovered unoptimized, and recovered optimized source code.

tations. Finally, `ded` defines exception tables that describe `try/catch/finally` blocks. The result when combined with the constant pool is a Java `.class` file.

Table 1 provides two example translation rules. The first example illustrates the translation of an `add` instruction. For each register, `ded` creates a corresponding Java local variable, i.e.,  $d_0 \rightarrow d'_0$ ,  $s_0 \rightarrow s'_0$ , etc. The translation creates four Java instructions: two instructions to push the variables onto the stack, one instruction to add, and one to pop the result. The second example, `invoke-virtual`, shows the translation of a virtual method invocation. This example demonstrates two intricate details for which `ded` must account. First, Dalvik’s register method arguments do not always correspond to Java’s variable method arguments. In Dalvik, registers are always 32-bits, and 64-bit variables are stored in two adjacent registers. Hence, `ded` must inspect the method descriptor when translating registers to variables for method invocation. Second, Dalvik bytecode commonly omits the `move-result` instruction if the return value is not used. However, the Java bytecode must always pop the return value off the stack, regardless of its use. `ded` includes related logic to ensure stack integrity.

#### 4.4 Optimization and Decompilation

At this stage, the retargeted `.class` files can be decompiled using existing tools, e.g., Fernflower [5] or Soot [38]. However, `ded`’s bytecode translation process yields unoptimized Java code. For example, Java tools often optimize out unnecessary assignments to the local variable table, e.g., unneeded return values. The lack of optimization yields complex decompiled code and frustrates its subsequent analysis. Furthermore, artifacts of the retargeting process lead to potential decompilation errors in some decompilers. To remedy these issues, we post process `.class` files using a bytecode optimizer such as Soot [38].

The need for bytecode optimization can be demonstrated by looking at decompiled loops. Most decompilers convert `for` loops into infinite loops with `break` instructions. While the resulting source code is functionally equivalent to the original, it is significantly more difficult to understand and analyze, especially for nested loops. Figure 6 shows example retargeted source with and without optimization; (#1) shows the original source code, (#2) shows unoptimized decompiled code, and (#3) shows optimized decompiled code. Note that the optimized decompiled code is, modulo variable names, virtually identical to the original source, whereas the unoptimized code is nearly indecipherable.

The Soot tool [38] performs a myriad of optimization, data

and control analysis, and decompilation on each program. We use Soot with default optimizations for the next two stages of application processing; optimization and decompilation (of Figure 1). Note that these processes can fail due to the limitations of our retargeting. We revisit the frequency and source of these failures in our evaluation of Android applications in the next section. Also note that other decompilers such as Fernflower may perform better than Soot. For example, an informal comparison of Soot and Fernflower decompilation suggests that the latter is more robust. We defer investigation of other tools to future work.

## 5. EVALUATION

This section studies the accuracy and results of the proposed certification process. We begin by validating the accuracy of the recovery process, then analyze the results of a preliminary certification of 1,100 applications downloaded from the Android market.

### 5.1 Recovery Validation

We first sought to validate the recovered code was faithful to the original source. Although tools exist to check the similarity of Java source code, we desired a much stronger comparison. The only reliable way to perform validation was, sadly, through manual inspection. Five open source applications ranging in size and origin were selected for validation. A percentage of the recovered methods (selected at random) were compared against the original source to ensure it was functionally identical. All aspects of the source code were compared, e.g., control flow, calls, arithmetic operators, etc. The validated applications included; Snake (421 LOC, 100% of methods validated), Radar (581 LOC, 20% checked), Translate (1868 LOC, 20% checked), Countdown by OpenIntents (4450 LOC, 20% checked), and WordPress (11,727 LOC, 5% checked).

The validation process found no errors in the recovery process; every method checked was functionally identical to the original code, with three exceptions. The exceptions occurred due to the limitations of Soot, and included a) incorrect `super` references, b) improper structure in `try/catch` blocks, and c) improper `super` constructor calls. These translation error are not present in tools such as FernFlower. Note that these errors did not effect our vulnerability detection described below in any meaningful way.

### 5.2 Certification Evaluation

Our second analysis focuses on the results of the certification itself. The experiments described here use the top 50

free applications available from the Android Market in each of the 22 application categories (API level 4). These 1,100 applications were downloaded using an automated tool from Android Market on December 11th, 2009. Our decision to use only free applications does not meaningfully impact the experimental results (beyond a potential for the paid applications to have undergone a more thorough development process). Android makes no distinction between paid or free applications, and no addition of DRM is applied. For completeness, we successfully recovered and analyzed source code from sample paid applications without error. We omit further discussion of paid applications for brevity.

ded was compiled using gcc version 4.4.1. The code recovery experiments were run on 64-bit, 8 core, 2.33 GHz (8 GB RAM) Dell blades running Linux Ubuntu version 9.10 (kernel version 2.6.31-17). We used a customized version of Soot with default optimization options. We relaxed superfluous (for our purposes) checks on internal string representations and deactivated the elimination of `store/load/load` trios that can yield illegal code on legal input. We further deactivated the code transformations that add unnecessary (non-Android) system support code. Certification tests use Yasca 2.1 with the PMD plugin version 4.2.5 4-core 2.33 using PHP version 5.2.10 on a 3.20 GHz Dell desktop. SLOC counts were generated using SLOCCount version 2.26 [13].

Note that imperfect analysis is subject to false-positives; we may incorrectly flag that source code violates a tested security property. How this is handled by the certifier is subject to policy; the developer may refactor the code to remove the violation or provide evidence/explanation that the reported error is a false positive. As with iPhone, the likely answer is probably a combination of both.

### 5.2.1 Source Code Recovery

The first set of experiments measured the code recovery process. Table 2 shows the number of classes that were successfully *retargeted* (converted from Dalvik to Java bytecode) and *recovered* (source code retrieved). Notably, over 99.6% of the `.class` files were successfully retargeted, and the source code for 95% of the classes was recovered.

The computational costs of recovering the source code were not onerous. The total time to recover the source code for all 1,100 applications was about 272 computation-hours, or a little over 14 minutes of processor time per application. Interestingly, the retargeting process was completely dominated by the optimization and decompilation process. The retargeting process took on average about 17 msec per class, whereas the source code recovery took about 6.8 seconds per class. The per-application recovery costs were largely linear in the number of classes (standard deviation 6.694 sec).

The 534 retargeting failures fell into three classes. The first class of failures was the result of unresolved class references. We confirmed that these errors can be fixed by investigating and supplying Soot precise versions of support classes needed by the failed application. We are investigating ways to automate the labor-intensive process of finding correct support libraries. The second class of errors occurred due to invalid input Dalvik bytecode. We observed a small number of applications whose bytecode had type system violations. This occurred when specific exceptional control flow structures (i.e., specific placement of `try/catch` blocks) were present in a method. Such bytecode can be repaired by reducing the scope of a `try` block to only instructions that are

capable of throwing the target exception. While we were able confirm the fix by directly editing the bytecode using a binary editor and retargeting, we have yet to identify an algorithm to do this automatically.<sup>3</sup> The last very infrequent class of errors remains unexplained. We are continuing to explore pathological boundary cases or errors in ded that produce invalid bytecode.

The source code recovery failures also were equally diverse. Soot is an optimization tool with the ability to recover source code in most cases, but does not process certain legal program idioms (code structures) generated by ded. We encountered and refactored many of these idioms while developing ded, but were unable to resolve all of them. In particular, two central problems we have encountered involve interactions between synchronized blocks and exception handling, and complex control flows caused by break statements. The difference between the success rates of retargeted and recovered rates is largely due to Soot's inability to extract source code from these otherwise legal idioms. Note that other tools such as Fernflower [5] can be more resilient in the face of these idioms. We chose not to use Fernflower in this paper because it is an online service that requires direct access to the bytecode. Sending retargeted bytecode over the Internet to a third party presents both logistical and potential legal challenges.

### 5.2.2 Application Vulnerability Analysis

The second battery of tests evaluated application vulnerability detection. Yasca was configured to detect *a) weak crypto*: the use of poor crypto algorithms, e.g., DES, unsafe PRNGs; *b) bad practices*: general bad programming practices, e.g., empty exception catch blocks, required but uncalled superclass constructors; *c) AJAX*: the use of AJAX (which Yasca/PMD deem dangerous); *d) file use*: files being created or modified in unsafe ways, e.g., poor log file handling; *e) DoS*: exposing applications to denial of service, e.g., reading from arbitrary files; *f) process control*: the use of external programs such as "su" or other low-level programs, or loading external Java libraries; and *g) hidden permissions*: permissions asserted within code, rather than through policy configuration manifest (see below).

Table 3 shows the results of the analysis of the Android applications. Yasca flagged over 3,000 different locations in the recovered source code. An investigation of the flagged code yields some interesting results. One surprising discovery was that there was widespread use of weak cryptographic algorithms. In particular, 56-bit DES, RC4 and MD5 were widely used. While it is unclear what these are used for, their use at all is enough for Yasca to consider the code suspect. In another vein, the vast majority of flagged code for AJAX existed in a single sports-score application. The application was tightly integrated with an online Web system, and thus these flags are likely false positives. The DoS vulnerabilities represented about 32% of the 3,119 vulnerabilities. Most of these flags were due to file reads; Yasca flags blocking I/O read calls as potential vulnerabilities requiring closer inspection to ensure an adversary cannot supply an input file that consumes unexpected resources or blocks indefinitely. In many cases the adversary will not be able to control the input file, therefore it is reasonable to expect that many errors are false positives. We were also surprised

<sup>3</sup>Soot has a documented feature to perform this code transformation, but it failed on all our sample bytecode.

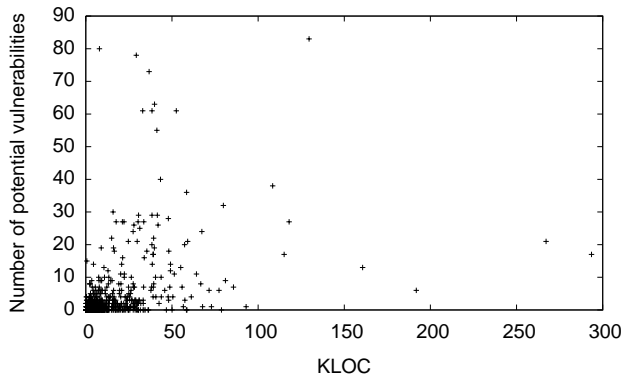


Table 2: Application source code recovery for the top 1,100 free Android Market applications (Dec. 2009).

Category	Size (bytes)	App classes	Retargeted classes	%	Recovered classes	%	SLOC
Comics	3,675,436	2,767	2,764	99.89%	2,656	95.99%	158,526
Communication	17,984,508	12,976	12,925	99.61%	12,010	92.56%	991,205
Demo	4,849,012	4,236	4,235	99.98%	4,061	95.87%	374,433
Entertainment	5,103,224	5,423	5,421	99.96%	5,239	96.61%	253,033
Finance	7,116,100	5,141	5,132	99.82%	4,972	96.71%	411,413
Arcade Games	10,244,184	5,401	5,361	99.26%	5,114	94.69%	610,042
Puzzle Games	4,995,604	3,308	3,290	99.46%	3,141	94.95%	257,554
Casino Games	9,586,980	4,947	4,920	99.45%	4,688	94.76%	572,903
Casual Games	5,662,640	3,492	3,483	99.74%	3,316	94.96%	303,673
Health	7,905,040	5,260	5,251	99.83%	5,099	96.94%	463,285
Lifestyle	9,547,124	6,692	6,676	99.76%	6,289	93.98%	544,005
Multimedia	12,889,520	9,902	9,844	99.41%	9,336	94.28%	784,082
News/Weather	9,565,676	6,610	6,574	99.46%	6,160	93.19%	466,465
Productivity	18,512,912	12,376	12,324	99.58%	11,745	94.90%	1,147,235
Reference	7,159,300	4,614	4,599	99.67%	4,403	95.43%	441,536
Shopping	11,263,896	8,318	8,300	99.78%	8,009	96.29%	738,213
Social	18,090,512	14,342	14,311	99.78%	13,744	95.83%	1,176,859
Libraries	4,757,036	3,653	3,646	99.81%	3,527	96.55%	301,248
Sports	13,580,008	9,589	9,550	99.59%	9,288	96.86%	1,650,778
Themes	411,568	508	506	99.61%	491	96.65%	13,079
Tools	5,331,940	3,690	3,670	99.46%	3,509	95.09%	304,347
Travel	13,906,124	9,845	9,774	99.28%	9,231	93.76%	877,031
<b>TOTAL</b>	<b>202,138,344</b>	<b>143,090</b>	<b>142,556</b>	<b>99.63%</b>	<b>136,028</b>	<b>95.06%</b>	<b>12,840,945</b>

Table 3: Application Analysis - potential vulnerabilities flagged by Yasca analysis.

Category	Weak crypto	Bad practice	AJAX	File use	DoS	Process control	Hidden permission	Total
Comics	0	0	0	0	48	1	0	49
Communication	22	1	0	3	83	131	235	475
Demo	3	0	0	0	6	14	9	32
Entertainment	0	0	79	0	42	0	19	140
Finance	1	0	0	0	44	23	8	76
Arcade Games	0	3	0	0	34	28	0	65
Puzzle Games	0	0	0	0	28	24	0	52
Casino Games	1	4	0	0	27	8	2	42
Casual Games	0	1	0	0	34	3	0	38
Health	5	0	0	1	43	64	12	125
Lifestyle	24	0	2	0	67	48	28	169
Multimedia	15	0	1	0	100	72	156	344
News/Weather	5	1	0	1	46	2	50	105
Productivity	31	7	6	0	58	19	66	187
Reference	5	1	0	0	36	13	4	59
Shopping	9	0	0	0	53	49	28	139
Social	19	0	0	0	74	166	47	306
Libraries	60	5	0	0	30	135	25	255
Sports	3	2	0	0	65	22	13	105
Themes	8	0	0	0	0	0	0	8
Tools	14	3	0	4	28	7	79	135
Travel	10	0	0	3	63	119	18	213
<b>TOTAL</b>	<b>235</b>	<b>28</b>	<b>88</b>	<b>12</b>	<b>1,009</b>	<b>948</b>	<b>799</b>	<b>3,119</b>



**Figure 7: Number of potential vulnerabilities based on lines of code in an application.**

at the pervasiveness of the process control flags; it seems that several programs were using the underlying embedded Linux tools. e.g., “su.” Such uses are dangerous under normal circumstances, but in Android they serve to bypass the permission system. Therefore they should be used carefully if at all.

Lastly, our Android-specific test searched for “hidden permissions” used for IPC. We characterize such instances as potential vulnerabilities because it requires developers to correctly enforce policy through code, rather than relying on the platform permission system. Thus, poorly implemented applications may silently violate the system policy. Moreover, because they are embedded in source code, policy analysis tools are not able to evaluate their correctness (e.g., Kirin [23]). Note that our custom permission check is only one (simple) heuristic for identifying potentially vulnerable security policy specification. However, even simple heuristic demonstrates that these intents were used in a small number of applications, but that those that did use them did so frequently. Future work will perform a more sophisticated analysis of security policy embedded within code.

Figure 7 plots the number of identified potential vulnerabilities in an application based on number of lines of code in that application. Note that both the application with the most vulnerabilities and the application with the most LOC were removed for presentation purposes. The plot does not indicate any obvious trend of potential vulnerabilities based on the size of the application. We experimented with several different representations (e.g., number of classes, binary size) with similar results.

Finally, we stress here that these are preliminary tests. Beyond the brief notes presented here, we have not systematically evaluated the flagged source code to see which represent real vulnerabilities and which are false positives. The focus of this work was to validate the certification process was feasible in practice. The identification of over 3,000 potential problems within the tested application serves as validation of this. The real certification of 1,100 applications, however, is an orthogonal and substantial effort in its own right.

## 6. RELATED WORK

Java decompilation has been around almost as long as the language itself. Proebsting invented a comprehensive java

decompilation method [35] in 1997 shortly after Mocha [10] decompiler was released. This remained the state of the art until 2001 when Dava [30, 29, 33] was introduced. Dava is built on the Soot tool used throughout. The Soot framework uses sophisticated type inference [25, 16, 32] and code analysis techniques to accurately recover the original code. Closed source decompilers such as Jad [8], JD [7] and Fernflower [5] are extremely effective at recovering source code, but their associated algorithms are not made public.

Ours is not the first work to use decompilation to enable security analysis [21]. Tools such as Yasca [15] and PMD [11], FindBugs [6], JLint [2], QJ-Pro [12], Lint4j [9], Checkstyle [3] verify best coding practices and the absence or presence of vulnerabilities or bugs. Interested readers are referred to surveys on code level [40] and bytecode-level [27] security analysis. These and similar tools have applied in many domains of security [20]. In particular, [28] presents a technique based on a points-to analysis for the detection of vulnerabilities in Java applications. Also, significant work has been done in large scale static analysis of C code [19, 36]. Our preliminary evaluation relies on simple lexical analysis tools [41, 20, 28]. Thus, future work will focus on more powerful analysis tools that will enable stronger certification.

Two parallel efforts are also attempting to recover source code from Android application code files. Like ded, the undx [14] and dex2jar [4] tools attempts to recover Dalvik application class files from the installation image. Our experiments using these tools on a array of applications demonstrated that they only recover about 14% (undx) and 6% (dex2jar) of classes. Further investigation of the tools showed that they lacked many of the key algorithms and structures needed to accurately retarget Android applications.

## 7. CONCLUSION

Smartphone application development is one of the fastest growing technology markets in the world. However, apart from weak assurances of market providers, users currently have no way of evaluating whether these applications are safe to use. This paper has presented a new user-centric certification method for this evaluation; organizations and concerned users extract the application source code from the installation image and perform code-level analysis of its structure. We have illustrated the design and challenges of the ded decompiler and performed a large but preliminary analysis of 1,100 mobile phone applications. The experiment shows that the certification process is feasible in practice.

This paper has described tools for the execution of application certification. However, this is only half of the story. Defining security policies appropriate for cell phones and the data they manage is a complex and context-sensitive process. Moreover, to enforce those policies, we need to perform deeper validation of the source code recovered using stronger analysis techniques. Our future work will focus centrally on these critical activities.

## 8. REFERENCES

- [1] Apple iphone - mobile phone, ipod, and internet device. <http://www.apple.com/iphone/>.
- [2] Artho software - jlint. <http://artho.com/jlint/>.
- [3] Checkstyle. <http://checkstyle.sourceforge.net/>.
- [4] dex2jar. <https://code.google.com/p/dex2jar/>.

- [5] Fernflower - java decompiler. <http://www.reversed-java.com/fernflower/>.
- [6] Findbugs - find bugs in java programs. <http://findbugs.sourceforge.net/>.
- [7] Jad - the fast java decompiler. <http://www.kpdus.com/jad.html>.
- [8] Jd java decompiler. <http://java.decompiler.free.fr/>.
- [9] Lint4j. <http://www.jutils.com/>.
- [10] Mocha, the java decompiler. <http://www.brouhaha.com/eric/software/mocha/>.
- [11] Pmd. <http://pmd.sourceforge.net/>.
- [12] Qj-pro - code analyzer for java. <http://qjpro.sourceforge.net/>.
- [13] Sloccount. <http://www.dwheeler.com/sloccount/>.
- [14] Undx, a reconstructor for dalvik bytecode. <http://www.illegalaccess.org/undx/>.
- [15] Yasca - yet another source code analyzer. <http://www.yasca.org/>.
- [16] B. Bellamy, P. Avgustinov, O. de Moor, and D. Sereni. Efficient local type inference. In G. Kiczales, editor, *OOPSLA '07: 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*. ACM, 2008.
- [17] D. Bornstein. Google i/o 2008 - dalvik virtual machine internals. <http://www.youtube.com/watch?v=ptjed0ZEXPM>.
- [18] E. Chaillou, P. Manoury, and B. Pagano. *Developing Applications with Objective Caml*. O'Reilly, France, 2000. Available from <http://caml.inria.fr/pub/docs/oreilly-book/index.html>.
- [19] H. Chen, D. Dean, and D. Wagner. Model checking one million lines of c code. In *Proc. of the 11th Annual Network and Distributed System Security Symposium (NDSS)*, 2004.
- [20] B. Chess and G. McGraw. Static analysis for security. *IEEE Security and Privacy*, 2(6):76–79, 2004.
- [21] C. Cifuentes, T. Waddington, and M. V. Emmerik. Computer security analysis through decompilation and high-level debugging. In *Proc. of the Workshop on Decompilation Techniques*, pages 375–380. IEEE Press, 2001.
- [22] G. Cluley. First iphone worm discovered - ikee changes wallpaper to rick astley photo, November 2009. <http://www.sophos.com/blogs/gc/g/2009/11/08/iphone-worm-discovered-wallpaper-rick-astley-photo/>.
- [23] W. Enck, M. Ongtang, and P. McDaniel. On Lightweight Mobile Phone App Certification. In *Proc. of the 16th ACM Conference on Computer and Communications Security (CCS)*, November 2009.
- [24] W. Enck, M. Ongtang, and P. McDaniel. Understanding android security. *IEEE Security and Privacy*, 7(1):50–57, 2009.
- [25] E. Gagnon, L. J. Hendren, and G. Marceau. Efficient inference of static types for java bytecode. In *SAS '00: Proc. of the 7th International Symposium on Static Analysis*, pages 199–219. Springer-Verlag, 2000.
- [26] D. Goodin. Backdoor in top iphone games stole user data, suit claims. The Register, November 2009. [http://www.theregister.co.uk/2009/11/06/iphone\\_games\\_storm8\\_lawsuit/](http://www.theregister.co.uk/2009/11/06/iphone_games_storm8_lawsuit/).
- [27] X. Leroy. Java bytecode verification: Algorithms and formalizations. *J. Autom. Reason.*, 30(3-4):235–269, 2003.
- [28] V. B. Livshits and M. S. Lam. Finding security vulnerabilities in java applications with static analysis. In *Proc. of the 14th conference on USENIX Security Symposium*. USENIX Association, 2005.
- [29] J. Miecznikowski and L. Hendren. Decompiling java using staged encapsulation. In *WCRE '01: Proc. of the Eighth Working Conference on Reverse Engineering*, page 368. IEEE Computer Society, 2001.
- [30] J. Miecznikowski and L. J. Hendren. Decompiling java bytecode: Problems, traps and pitfalls. In *CC '02: Proc. of the 11th International Conference on Compiler Construction*, pages 111–127. Springer-Verlag, 2002.
- [31] R. Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, August 1978.
- [32] A. Mycroft. Type-based decompilation (or program reconstruction via type reconstruction). In *Proc. of the 8th European Symposium on Programming Languages and Systems*, pages 208–223. Springer-Verlag, 1999.
- [33] N. A. Naeem and L. Hendren. Programmer-friendly decompiled java. In *ICPC '06: Proc. of the 14th IEEE International Conference on Program Comprehension*, pages 327–336. IEEE Computer Society, 2006.
- [34] M. Ongtang, S. McLaughlin, W. Enck, and P. McDaniel. Semantically Rich Application-Centric Security in Android. In *Proc. of the 25th Annual Computer Security Applications Conference (ACSAC)*, December 2009.
- [35] T. A. Proebsting and S. A. Watterson. Krakatoa: Decompilation in java (does bytecode reveal source?). In *Third USENIX Conference on Object-Oriented Technologies and Systems (COOTS)*, 1997.
- [36] B. Schwarz, H. Chen, D. Wagner, J. Lin, W. Tu, G. Morrison, and J. West. Model checking an entire linux distribution for security violations. In *ACSAC '05: Proc. of the 21st Annual Computer Security Applications Conference*, pages 13–22. IEEE Computer Society, 2005.
- [37] J. Tiuryn. Type inference problems: A survey. In *MFCS '90: Proc. of the Mathematical Foundations of Computer Science 1990*, pages 105–120. Springer-Verlag, 1990.
- [38] R. Vallee-Rai, E. Gagnon, L. Hendren, P. Lam, P. Pominville, and V. Sundaresan. Optimizing java bytecode using the soot framework: Is it feasible? In *International Conference on Compiler Construction, LNCS 1781*, pages 18–34, 2000.
- [39] L. Wall, T. Christiansen, and J. Orwant. *Programming Perl*. O'Reilly & Associates, Inc., 2000.
- [40] M. S. Ware and C. J. Fox. Securing java code: heuristics and an evaluation of static analysis tools. In *SAW '08: Proc. of the 2008 workshop on Static analysis*, pages 12–21. ACM, 2008.
- [41] M. Zitser, R. Lippmann, and T. Leek. Testing static analysis tools using exploitable buffer overflows from open source code. *SIGSOFT Softw. Eng. Notes*, 29(6):97–106, 2004.