Reflection Analysis for Java

Benjamin Livshits John Whaley
Monica S. Lam

Computer Science Department
Stanford University
Stanford, CA 94305, USA

{livshits, jwhaley, lam}@cs.stanford.edu

Technical Report

October 29, 2005

Contents

Abstract

Reflection has always been a thorn in the side of Java static analysis tools. Without a full treatment of reflection, static analysis tools are both *incomplete* because some parts of the program may not be included in the application call graph, and *unsound* because the static analysis does not take into account reflective features of Java that allow writes to object fields and method invocations. However, accurately analyzing reflection has always been difficult, leading to most static analysis tools treating reflection in an unsound manner or just ignoring it entirely. This is unsatisfactory as many modern Java applications make significant use of reflection.

In this paper we propose a static analysis algorithm that uses points-to information to approximate the targets of reflective calls as part of call graph construction. Because reflective calls may rely on input to the application, in addition to performing reflection resolution, our algorithm also discovers all places in the program where user-provided specifications are necessary to fully resolve reflective targets. As an alternative to user-provided specifications, we also propose a reflection resolution approach based on type cast information that reduces the need for user input, but typically results in a less precise call graph.

We have implemented the reflection resolution algorithms described in this paper and applied them to a set of six large, widely-used benchmark applications consisting of more than 600,000 lines of code combined. Experiments show that our technique is effective for resolving most reflective calls without any user input. Certain reflective calls, however, cannot be resolved at compile time precisely. Relying on a user-provided specification to obtain a conservative call graph results in graphs that contain 1.43 to 6.58 times more methods that the original. In one case, a conservative call graph has 7,047 more methods than a call graph that does not interpret reflective calls. In contrast, ignoring reflection leads to missing substantial portions of the application call graph.

SECTION 1

Introduction

Whole-program static analysis requires knowing the targets of function or method calls. The task of computing a program's call graph is complicated for a language like Java because of virtual method invocations and reflection. Past research has addressed the analysis of function pointers in C [EGH94, MRR01, MRR04] as well as virtual method calls in C++ [AH96, BS96, CG94, PR96] and Java [GC01, GDDC97, RRHK00, SHR+00, TP00]. Reflection, however, has mostly been neglected.

Reflection in Java allows the developer to perform runtime actions given the descriptions of the objects involved: one can create objects given their class names, call methods by their name, and access object fields given their name [FF04]. Because names of methods to be invoked can be supplied by the user, especially in the presence of dynamic class loading, precise static construction of a call graph is generally undecidable. Even if we assume that all classes that may be used are available for analysis, without placing *any restrictions* of the targets of reflective calls, a sound (or conservative) call graph would be prohibitively large.

Many projects that use static analysis for optimization, error detection, and other purposes ignore the use of reflection, which makes static analysis tools *incomplete* because some parts of the program may not be included in the call graph and potentially *unsound*, because some operations, such as reflectively invoking a method or setting an object field, are ignored. Others require the user to specify the methods invoked reflectively [TLSS99]. Completeness is not the only problem with ignoring reflection: the results may in fact also be *unsound* if the static analysis does not take into account reflective features of Java that allow writes to object fields and reflective method invocations. Our research is motivated by the practical need to improve the coverage of static error detection tools [KPK02, LL05a, RSS+04, WN04]. The success of such tools in Java is predicated upon having a call graph available to the error detection tool. Unless reflective calls are interpreted, the tools run the danger of only analyzing a small portion of the available code and giving the developer a false sense of security when no bugs are reported. Moreover, when static results are used to reduce runtime instrumentation, all parts of the application that are used at runtime *must* be statically analyzed.

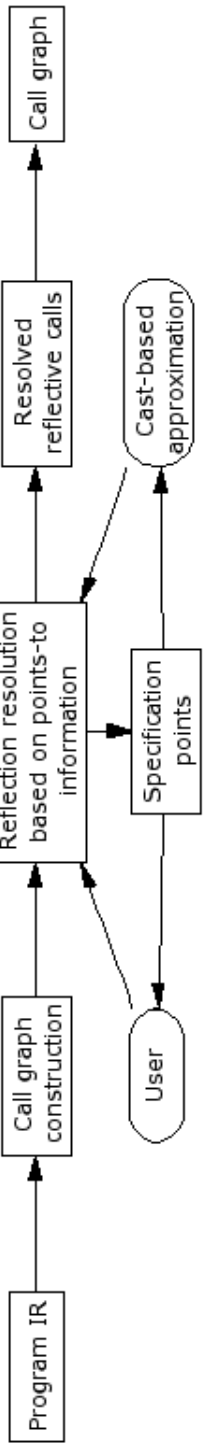

Figure 1: Architecture of our static analysis framework.

Retrieving `Class` or `Constructor` objects:

`Class`	`java.lang.Class.forName(String className)`
`Class`	`java.lang.Class.forName(String name, boolean initialize, ClassLoader loader)`
`Class`	`java.lang.Object.getClass()`
`Constructor[]`	`java.lang.Class.getConstructors()`
`Constructor`	`java.lang.Class.getConstructor(Class[] args)`

Creating new objects:

`Object`	`java.lang.Class.newInstance()`
`Object`	`java.lang.reflect.Constructor.newInstance(Object[] initargs)`

Figure 2: Java API methods for reflective object creation.

While finding *some* bugs is valuable, a tool that claims to find *all* possible bugs of a particular kind provides a much stronger guarantee: a software system for which no errors are statically reported in known to be error-free.

A recent paper by Hirzel, Diwan, and Hind proposes the use of dynamic instrumentation to collect the reflection targets discovered at run time [HDH04]. They use this information to extend Andersen's context-insensitive, inclusion-based pointer analysis for Java into an online algorithm [And94]. Reflective calls are generally used to offer a choice in the application control flow, and a dynamic application run typically includes only several of all the possibilities. However, analyses used for static error detection and optimization often require a *full* call graph of the program in order to achieve complete coverage.

In this paper we present a static analysis algorithm that uses points-to information to determine the targets of reflective calls. Often the targets of

Retrieving `Method` objects:

`Method`	`java.lang.Class.getDeclaredMethod(String name, Class[] parameterTypes)`
`Method[]`	`java.lang.Class.getDeclaredMethods()`
`Method[]`	`getMethods()`

Calling methods:

`Object`	`java.lang.reflect.Method.invoke(Object obj, Object[] args)`

Figure 3: Java API methods for method invocation.

Retrieving `Field` objects:

```
Field    getField(String name)
Field[]  getFields()
Field    getDeclaredField(String name)
Field[]  getDeclaredFields()
```

Accessing field values: getter and setter methods for objects and primitive types:

```
Object   java.lang.Reflect.Field.get(Object obj)
byte     java.lang.Reflect.Field.getByte(Object obj)
char     java.lang.Reflect.Field.getChar(Object obj)
int      java.lang.Reflect.Field.getInt(Object obj)
                                ...
void     java.lang.Reflect.Field.set(Object obj, Object value)
void     java.lang.Reflect.Field.setByte(Object obj, byte b)
void     java.lang.Reflect.Field.setChar(Object obj, char c)
void     java.lang.Reflect.Field.setInt(Object obj, int i)
```

Figure 4: Java API methods for field access and manipulation.

reflective calls can be determined precisely by analyzing the flow of strings that represent class names throughout the program. This allows us to precisely resolve many reflective calls and add them to the call graph. However, in some cases reflective call targets may depend on user input and require user-provided specifications for the call graph to be determined. Our algorithm determines all *specification points* — places in the program where user-provided specification is needed to determine reflective targets. The user is given the option to provide a specification and our call graph is complete with respect to the specifications provided [TP00].

Because providing reflection specifications can be time-consuming and error-prone, we also provide a conservative, albeit sometimes imprecise, approximation of targets of reflective calls by analyzing how type casts are used in the program. A common coding idiom consists of casting the result of a call to `Class.newInstance` used to create new objects to a more specific type before the returned object can be used. Relying on cast information allows us to produce a conservative call graph approximation without requiring user-provided reflection specifications in most cases. A flow diagram summarizing the stages of our analysis is shown in Figure 1.

Our reflection resolution approach hinges on three assumptions about the use of reflection: (a) all the class files that may be accessed at run-time are available for analysis; (b) the behavior of `Class.forName` is consistent with its API definition in that it returns a class whose name is specified by the first parameter, and (c) cast operations that operate on the results of `Class.newInstance` calls are correct. In rare cases when no cast information is available to aid with reflection resolution, we report this back to the user as a situation requiring specification.

1.1 Contributions

This paper makes the following contributions:

- We present a case study of common uses of reflection in large modern Java systems. This study shows the importance of handling reflection in call graph construction.
- We formulate a set of natural assumptions that hold in most Java applications and make the use of reflection amenable to static analysis.
- We propose a call graph construction algorithm that uses points-to information about strings used in reflective calls to statically find potential call targets. When reflective calls cannot be fully "resolved" at compile time, our algorithms determines a set of specification points — places in the program that require user-provided specification to resolve reflective calls.
- As an alternative to having to provide a reflection specification, we propose an algorithm that uses information about type casts in the program to statically approximate potential targets of reflective calls.
- We provide an extensive experimental evaluation of our analysis approach based on points-to results by applying it to a suite of six large open-source Java applications consisting of more than 600,000 lines of code combined. We evaluate how the points-to and cast-based analyses of reflective calls compare to a local intra-method approach. While all these analyses find at least one constant target for most `Class.forName` call sites, they only moderately increase the call graph size. However, the conservative call graph obtained with the help of a user-provided specification results is a call graph than is almost 7 times as big as the

original. We assess the amount of effort required to come up with a specification and how cast-based information can significantly reduce the specification burden placed on the user.

1.2 Paper Organization

The rest of the paper is organized as follows. In Section 2, we provide background information about the use of reflection in Java. In Section 3 we show some of the common usage idioms that justify the need for reflection analysis. In Section 4, we lay out the simplifying assumptions made by our static analysis. In Sections 5 we describe our analysis approach. Section 6 provides a comprehensive experimental evaluation. In Sections 7 and 8 we describe related work and conclude. Finally, Section 9 provides the acknowledgements.

SECTION 2

Overview of Reflection in Java

In this section we first informally introduce the reflection APIs in Java and then show some characteristic ways in which they are used in large Java applications.

2.1 Reflection APIs in Java

The most typical use of reflection by far is for creating new objects given the object class name. The most common usage idiom for reflectively creating an object is shown in Figure 5. In the rest of this section, we fully describe reflective APIs that Java provides for creating objects, invoking methods, and reading and writing to data structures at runtime. There are also read-only API methods that are used for *runtime discovery*; for example, an application can check if a certain method exists before trying to invoke it.

2.1.1 Obtaining `Class` Objects

Obtaining a class given its name is most typically done using a call to one of the static functions `Class.forName(String,...)` and passing the class name as the first parameter. We should point out that while `Class.forName` is the most common way to obtain a class given its name, it may not be the only method for doing so. An application may define a native method that implements the same functionality. The same observation applies to other standard reflective API methods.

The `.class` construct is syntactic sugar that is translated by the compiler down into basic calls to `Class.forName`. The translation is somewhat different depending on the version of the compiler. For example, when `T.class` is

```
1.  String className = ...;
2.  Class c  = Class.forName(className);
3.  Object o = c.newInstance();
4.  T      t = (T) o;
```

Figure 5: Typical use of reflection to create new objects.

```
static java.lang.Class class$test$T;

   ...
   0:   getstatic       #7; //Field class$test$T:Ljava/lang/Class;
   3:   ifnonnull       18
   6:   ldc             #8; //String test$T
   8:   invokestatic    #9; //class$:(Ljava/lang/String;)Ljava/lang/Class;
   11:  dup
   12:  putstatic       #7; //Field class$test$T:Ljava/lang/Class;
   15:  goto            21
   18:  getstatic       #7; //Field class$test$T:Ljava/lang/Class;
   21:  astore_1
   22:  getstatic       #10; //Field java/lang/System.out:Ljava/io/PrintStream;
   25:  new             #11; //class StringBuffer
   28:  dup
   29:  invokespecial   #12; //StringBuffer."<init>":()V
   32:  ldc             #13; //String c:
   34:  invokevirtual   #14; //StringBuffer.append:(Ljava/lang/String;)Ljava/lang/StringBuffer;
   37:  aload_1
   38:  invokevirtual   #15; //StringBuffer.append:(Ljava/lang/Object;)Ljava/lang/StringBuffer;
   41:  invokevirtual   #16; //StringBuffer.toString:()Ljava/lang/String;
   44:  invokevirtual   #17; //PrintStream.println:(Ljava/lang/String;)V
   47:  return

static java.lang.Class class$(java.lang.String);
  Code:
   0:   aload_0
   1:   invokestatic    #1; //Class.forName:(Ljava/lang/String;)Ljava/lang/Class;
   4:   areturn
   5:   astore_1
   6:   new             #3; //class NoClassDefFoundError
   9:   dup
   10:  aload_1
   11:  invokevirtual   #4; //ClassNotFoundException.getMessage:()Ljava/lang/String;
   14:  invokespecial   #5; //NoClassDefFoundError."<init>":(Ljava/lang/String;)V
   17:  athrow
  Exception table:
   from   to  target type
     0     4     5   Class java/lang/ClassNotFoundException
```

Figure 6: Interpretation of .`class` in JDK version 1.4 and below.

translated, Sun's version of `javac` in JDK 1.4 produces bytecode shown in Figure 6. In this case, method `class$` that takes a class name and returns the class returned by `Class.forName` is generated by the compiler. The result of the call is stored in field `class$test$T`. The same compiler in JDK 1.5 takes a more efficient approach that results in a much shorter bytecode sequence:

```
ldc_w   #2; //class test$T
astore_1
```

In this case, the `Class` object is loaded from a constant pool. The analysis

described here handles the JDK 1.4 interpretation; supporting the JDK 1.5 interpretation requires a simple extension of our algorithm to reflect the creation of the contact pool.

2.1.2 Reflective Object Creation

Object creation APIs in Java provide a way to programmatically create objects of a class, whose name is provided at runtime; parameters of the object constructor can be passed in as necessary. Relevant Java API methods are summarized in Figure 2. Creating an object with an empty constructor is achieved through a call to `newInstance` on the appropriate `java.lang.Class` object, which provides a runtime representation of a class.

While methods `Class.forName` and `Class.newInstance` represent the majority of uses of reflection in real-life software systems, Java also provides ways to reflectively invoke a method given its name and to access the value of an object field at runtime, as described below [FF04].

2.1.3 Reflective Method Invocation

Methods are obtained from a `Class` object by supplying the method signature or by iterating through the array of `Method`s returned by one of `Class` functions. `Method`s are subsequently invoked by calling `Method.invoke`. The complete list of relevant API functions is summarized in Figure 3.

2.1.4 Reflective Field Access

Fields of Java runtime objects can be read and written at runtime. Calls to `Field.get` and `Field.set` can be used to get and set fields containing objects. Additional methods are provided for fields of primitive types. (All Java primitive types are supported, we limit the list in Figure 4 to several representative ones only.)

SECTION 3

Use of Reflection: Case Studies

In this Section, we describe some of the common usage patterns for reflection found in large Java systems. We identified these patterns by studying large Java applications downloaded from SourceForge; more details on these applications can be found in Section 6. In addition to describing each use case, we show why statically resolving reflection is important.

```
public void addHandlers(String path) {
        XmlIO xmlFile = new XmlIO(DiskIO.getResourceURL(path));
        xmlFile.load();

        XmlElement list = xmlFile.getRoot().getElement("handlerlist");
        Iterator it = list.getElements().iterator();
        while (it.hasNext()) {
            XmlElement child = (XmlElement) it.next();
            String id = child.getAttribute("id");
            String clazz = child.getAttribute("class");

            AbstractPluginHandler handler = null;
            try {
                Class c = Class.forName(clazz);
                handler = (AbstractPluginHandler) c.newInstance();
                registerHandler(handler);
            } catch (ClassNotFoundException e) {
                if (Main.DEBUG) e.printStackTrace();
            } catch (InstantiationException e1) {
                if (Main.DEBUG) e1.printStackTrace();
            } catch (IllegalAccessException e1) {
                if (Main.DEBUG) e1.printStackTrace();
            }
        }
    }
```

Figure 7: Creating objects reflectively based on an XML specification.

```
1.   String geneClassName = thisGeneElement.
2.               getAttribute(CLASS_ATTRIBUTE);
3.
4.   Gene thisGeneObject = (Gene) Class.forName(
5.                  geneClassName).newInstance();
```

Figure 8: Creating objects reflectively based on an XML specification.

3.1 Specifying Application Extensions

Many large applications support plugins, which are usually detected by the application upon startup by either reading a specification file or looking for files in a specific directory. For example, `columba`, an open-source email client parses an XML specification file to determine which plugins to instantiate, as shown in Figure 7. A similar scheme is supported by `jedit`, which also supports high levels of customization and downloadable plugins.

Application servers such as Apache Tomcat, use similar schemes where they retrieve plugin descriptions from files or by traversing a predefined directory `WEB – INF` [BD03]. Clearly, static analysis needs to be aware of these application extensions. If reflective calls are not properly resolved, most of the plugins or Web applications in the case of Tomcat would be completely missing from the call graph. However, this represents a case where without "hints" from the user static analysis cannot determine which plugins to analyze.

```
try {
    Class macOS  = Class.forName("gruntspud.standalone.os.MacOSX");
    Class argC[] = {ViewManager.class};
    Object arg[] = {context.getViewManager()};
    Method init = macOS.getMethod("init", argC);
    Object obj  = macOS.newInstance();
    init.invoke(obj, arg);
} catch (Throwable t) {
    // not on macos
}
```

Figure 9: Calling a method if it is present on the runtime platform.

```
Method m = c.getMethod("clone", null);
if (Modifier.isPublic(m.getModifiers())) {
    try {
        result = m.invoke(object, null);
    }
    catch (Exception e) {
        e.printStackTrace();
    }
}
```

Figure 10: Checking if a method is present before calling it.

3.2 Custom-made Object Serialization Scheme

Often objects are reflectively created based on a specification that is passed to the application.

Example 1. Reflection is used by an open-source genetic algorithm library, `jgap`, to implement a customized serialization scheme. Information on genes is saved in an XML file and then later loaded to create runtime data structures. The names of the classes to be created are read from an XML element on lines 1–2 in Figure 8 and the objects are created on lines 4–5. □

Like the plugin examples above, this example demonstrates the need for user-provided specifications of reflective targets [TP00] when the strings on which reflective calls depend are not constant.

```
try {
    //  Test for being run under JDK 1.4+
    Class.forName("javax.imageio.ImageIO");
    //  Test for JFreeChart being compiled
    //  under JDK 1.4+
    Class.forName("org.jfree.chart.encoders.SunPNGEncoderAdapter");
} catch (ClassNotFoundException e) {
    ...
}
```

Figure 11: Using reflection to circumventing JDK inconsistencies.

```
Method getVersionMethod =
    Class.forName("org.columba.core.main.ColumbaVersionInfo").
        getMethod("getVersion", new Class[0]);

return (String) getVersionMethod.invoke(null,new Object[0]);
```

Figure 12: Using a method that is not available at compile time.

3.3 Improving Portability Using Reflection

While many Java applications are fully platform-independent, there are often subtle reasons to use platform-specific code, especially in large systems.

Example 2. Reflection is used in `gruntspud`, an open-source graphical CVS client, to improve code portability across different platforms. As shown in the code excerpt in Figure 9, call of method `init` is executed *only* if the call to `Class.forName` in the `try` clause succeeds when the Mac-OS-specific class is available. □

Similarly, sometimes applications check if certain methods are available before calling them:

Example 3. A generic cloning routine in `jfreechart` checks that `clone` is available and declared to be `public` before attempting to call it, as shown in Figure 10. The call is attempted only if the method is `public`. □

Platform-dependent features are not the only reason to use reflection for the purpose of introspection. The behavior of the program can differ depending on the JDK version being used as well:

Example 4. The code in Figure 11 illustrates another use of reflection to get around incompatibilities in the JDK implementations across different distri-

```
public JDBCCategoryDataset(String url, String driverName,
                           String user, String passwd)
    throws ClassNotFoundException, SQLException
{
    Class.forName(driverName);
    this.connection = DriverManager.getConnection(url, user, passwd);
}
```

Figure 13: Using `Class.forName` for its side-effects.

```
1.      fieldSysPath = ClassLoader.class.getDeclaredField("sys_paths");
2.      fieldSysPath.setAccessible(true);
3.
4.      if (fieldSysPath != null) {
5.          fieldSysPath.set(System.class.getClassLoader(), null);
6.      }
```

Figure 14: Circumventing static type checking to set a field.

butions. The code conditionally creates an instance of `SunPNGEncoderAdapter` if `jfreechar` is used with a JDK version 1.4 and above. □

Examples 2—4 illustrate an inherent weakness of dynamic analysis that manifests itself when it comes to platform-specific code. Only a subset of the code in an application is executed on any particular platform. Static techniques, on the other hand, can analyze parts of code intended to be executed on different platforms all at once. If we want to detect subtle platform-specific errors that are hard to reproduce at runtime, obtaining a full call graph of the application requires reflection resolution.

3.4 Code Unavailable Until Deployment

Reflection is also used to examine information that does not exist at compile time and only becomes available after the application is deployed.

Example 5. The code from `columba` in Figure 12 invokes method `getVersion` of class `org.columba.core.main.ColumbaVersionInfo`. Upon examining the code, we found that this class is not created until after the application is deployed, which is when the version information becomes available. □

As described in Section 5.4.1, to make classes generated at deployment time available to our analysis, our techniques collect information on all classes *after* application deployment. Example 5 is a specific case of a more general Java design pattern, in which interface types are used and their implementations are substituted in a manner that is deployment-specific. Unless reflection is resolved, all objects of the interface types will lack an implementation that can be statically analyzed.

3.5 Using `Class.forName` for its Side-effects

The call to `Class.forName` has the additional effect of calling the class constructor of the class being referenced. Occasionally, the result of the call is ignored and the call is used as a convenient way to invoke the class constructor. This coding idiom is commonly used to initialize database drivers as shown in a code excerpt extracted from `jfreechart` in Figure 13.

3.6 Getting Around Static Type Checking

While this is relatively uncommon, reflection makes it possible to circumvent the standard Java type system.

Example 6. As shown in the code snippet in Figure 14 extracted from `columba`, reflection is used to reset the system library paths of the default class loader by setting field `sys_paths` to `null`. Since the field is non-public, the accessibility flag of the field is first reset on line 2 before assigning to the field on lines 5. □

Without taking into account methods that assign to fields of objects when constructing the program representation, the representation will be simply incomplete.

3.7 Providing a Built-in Interpreter

On occasion, a very wide set of classes may be returned by a reflective calls, as shown below.

Example 7. One of our benchmark applications, `jedit`, contains an embedded BeanShell, a Java source interpreter used to write editor macros [Nie]. Within the BeanShell interpreted, one of the calls to `Class.forName` takes type parameters extracted from the Bean shell macros. □

Clearly, ignoring the targets of the `Class.forName` call in this case leads to the code within the macro file not being analyzed. But this example also reveals a key difficulty: the set of macros is hardly static. New macros can be downloaded or written, so the approximation of the reflective targets is only valid with respect to a specific application configuration.

SECTION 4

Assumptions About Reflection

This section presents assumptions we make in our static analysis for resolving reflection in Java programs. We believe that these assumptions are quite reasonable and hold for many real-life Java applications.

The problem of precisely determining the classes that an application may access is undecidable. Furthermore, for applications that access the network, the set of classes that may be accessed is *unbounded*: we cannot possibly hope to analyze all classes that the application may conceivably download from the net and load at runtime. Programs can also dynamically generate classes to be subsequently loaded. Our analysis assumes a closed world, as defined below.

Assumption 4.1 Closed world.
We assume that only classes reachable from the class path at analysis time can be used by the application at runtime.

In the presence of user-defined class loaders, it is impossible to statically determine the behavior of function `Class.forName`. If custom class loaders are used, the behavior of `Class.forName` can change; it is even possible for a malicious class loader to return completely unrelated classes in response to a `Class.forName` call. The following assumption allows us to interpret calls to `Class.forName`. We assume that the behavior of `Class.forName` is consistent with the API declaration even when custom class loaders are used, which postulates that:

> *"Given the fully qualified name for a class or interface (in the same format returned by* `getName`*) this method attempts to locate, load, and link the class or interface."*

Assumption 4.2 Well-behaved class loaders.
The name of the class returned by a call to `Class.forName(className)` *equals* `className`*.*

To check the validity of this assumption, we have instrumented large applications to observe the behavior of `Class.forName`; we have never encountered a violation of this assumption. Finally, we introduce the following assumption that allows us to leverage type cast information contained in the program to constrain the targets of reflective calls.

Assumption 4.3 Correct casts.
Type cast operations that always operate on the result of a call to `newInstance` *are correct; they will always succeed without throwing a* `ClassCastException`.

We believe this to be a valid practical assumption: while it is possible to have casts that fail, causing an exception that is caught so that the instantiated object can be used afterwards, we have not seen such cases in practice. Typical `catch` blocks around such casts lead to the program terminating with an error message.

SECTION 5

Analysis of Reflection

In this section, we present techniques for resolving reflective calls in a program. Our analysis consists of the following three steps:

1. We use a sound points-to analysis to determine all the possible sources of strings that are used as class names. Such sources can either be constant strings or derived from external sources. The pointer analysis-based approach *fully resolves* the targets of a reflective call if constant strings account for all the possible sources. We say that a call is *partially resolved* if the sources can be either constants or inputs and *unresolved* if the sources can only be inputs. Knowing which external sources may be used as class names is useful because users can potentially specify all the possible values; typical examples are return results of file read operations. We refer to program points where the input strings are defined as *specification points*.

2. Unfortunately the number of specification points in a program can be large. Instead of asking users to specify the values of every possible input string, our second technique takes advantage of casts, whenever available, to determine a conservative approximation of targets of reflective calls *that are not fully resolved*. For example, as shown in Figure 5, the call to `Class.newInstance`, which returns an `Object`, is always followed by a cast to the appropriate type before the newly created object can be used. Assuming no exception is raised, we can conclude that the new object must be a subtype of the type used in the cast, thus restricting the set of objects that may be instantiated.

3. Finally, we rely on user-provided specification for the remaining set of calls — namely calls whose source strings are not all constants — in order to obtain a conservative approximation of the call graph.

We start by describing the call graph discovery algorithm in Section 5.1 as well as how reflection resolution fits in with call graph discovery. Section 5.2 presents a reflection resolution algorithm based on pointer analysis results. Finally, Section 5.4 describes our algorithm that leverages type cast information for conservative call graph construction without relying on user-provided specifications.

5.1 Call Graph Discovery

Our static techniques to discover reflective targets are integrated into a context-insensitive points-to analysis that discovers the call graph on the fly [WL04]. As the points-to analysis finds the pointees of variables, type information of these pointees is used to resolve the targets of virtual method invocations, increasing the size of the call graph, which in turn is used to find more pointees. Our analysis of reflective calls further expands the call graph, which is used in the analysis to generate more points-to relations, leading to bigger call graphs. The discovery algorithm terminates when a fixpoint is reached and no more call targets or points-to relations can be found.

By using a points-to analysis to discover the call graph, we can obtain a more accurate call graph than by using a less precise technique such as class hierarchy analysis CHA [DGC95] or rapid type analysis RTA [Bac98]. We use a context-insensitive version of the analysis because context sensitivity does not seem to substantially improve the accuracy of the call graph [WL04, GC01] and the context-insensitive version is substantially faster.

5.2 Pointer Analysis for Reflection

This section describes how we leverage pointer analysis results to resolve calls to `Class.forName` and track `Class` objects. This can be used to discover the types of objects that can be created at calls to `Class.newInstance`, along with resolving reflective method invocations and field access operations. Pointer analysis is also used to find specification points: external sources that propagate string values to the first argument of `Class.forName`.

5.2.1 Reflection and Points-to Information

The programming idiom that motivated the use of points-to analysis for resolving reflection was first presented in Figure 5. This idiom consists of the following steps:

1. Obtain the name of the class for the object that needs to be created.
2. Create a `Class` object by calling the static method `Class.forName`.
3. Create the new object with a call to `Class.newInstance`.

4. Cast the result of the call to `Class.newInstance` to the necessary type in order to use the newly created object.

When interpreting this idiom statically, we would like to "resolve" the call to `Class.newInstance` in step 3 as a call to the default constructor `T()`. However, analyzing even this relatively simple idiom is nontrivial.

The four steps shown above can be widely separated in the code and reside in different methods, classes, or jar libraries. The `Class` object obtained in step 2 may be passed through several levels of function calls before being used in step 3. Furthermore, the `Class` object can be deposited in a collection to be later retrieved in step 3. The same is true for the name of the class created in step 1 and used later in step 2. To determine how variables `className`, `c`, `o`, and `t` defined and used in steps 1–4 may be related, we need to know what runtime objects they may be referring to: a problem addressed by *points-to* analysis. Point-to analysis computes which objects each program variable may refer to.

Resolution of `Class.newInstance` of `Class.forName` calls is not the only thing made possible with points-to results: using points-to analysis, we also track `Method`, `Field`, and `Constructor` objects. This allows us to correctly resolve reflective method invocations and field accesses. Reflection is also commonly used to invoke the class constructor of a given class via calling `Class.forName` with the class name as the first argument. We use points-to information to determine potential targets of `Class.forName` calls and add calls to class constructors of the appropriate classes to the call graph.

5.2.2 The bddbddb Program Database

In the remainder of this section we describe how pointer information is used for reflection resolution. We start by describing how the input program can be represented as a set of relations in `bddbddb`, a BDD-based program database [LWL+05, WL04]. The program database and the associated constraint resolution tool allows program analyses to be expressed in a succinct and natural fashion as a set of rules in Datalog, a logic programming language. Points-to information is compactly represented in `bddbddb` with binary decision diagrams (BDDs), and can be accessed and manipulated efficiently with Datalog queries. The program representation as well as pointer analysis results are stored as relations in the `bddbddb` database. The domains in the database include invocation sites I, variables V, methods M, heap objects named by their allocation site H, types T, and integers Z.

actual: $I \times Z \times V$.	*actual*(i, z, v) means that variable v is zth argument of the method call at i.
ret: $I \times V$.	*ret*(i, v), means that variable v is the return result of the method call at i.
mret: $M \times V$.	*ret*(m, v), means that variable v is the return result of method m.
assign: $V \times V$.	*assign*(v_1, v_2) means that there is an implicit or explicit assignment statement $v_1 = v_2$ in the program.
load: $V \times F \times V$.	*load*(v_1, f, v_2) means that there is a load statement $v_2 = v_1.f$ in the program.
store: $V \times F \times V$.	*store*(v_1, f, v_2) means that there is a store statement $v_1.f = v_2$ in the program.
string2class: $H \times T$.	*string2class*(s, t) means that string constant s is the string representation of the name of type t.
string2method: $H \times M$.	*string2method*(s, m) means that string constant s is the string representation of the name of method m.
string2field: $H \times F$.	*string2field*(s, f) means that string constant s is the string representation of the name of field f.
calls: $I \times M$.	*calls*(i, m) means that invocation site i may invoke method m.

Figure 15: Datalog relations used to represent the input program.

The source program is represented as a number of input relations. For instance, relations *actual* and *ret* represent parameter passing and method returns, respectively. In the following, we say that predicate $A(x_1, \ldots, x_n)$ is true if tuple $(x_1, \ldots, x_n)$ is in relation A.

The definitions of Datalog relations used to represent the input program are presented in Figure 15. All of these input relations are in lower case to make them stand our from relations defined for the purpose of reflection resolution. Points-to results are represented with relations *points-to* and *hpoints-to*:

points-to: $V \times H$ is the variable points-to relation. *points-to*(v, h) means that variable v may point to heap object h.

hpoints-to: $H \times F \times H$ is the heap points-to relation. *hpoints-to*(h_1, f, h_2) means that field f of heap object h_1 may point to heap object h_2.

Finally, an auxiliary relation *freshi2h* is used for reflection resolution to make "fresh" heap allocation sites:

freshi2h: $I \times H$. *freshi2h*(i, h) means that a *freshly created* allocation site h corresponds to the result of the method call at i.

A Datalog query consists of a set of rules, written in a Prolog-style notation, where a predicate is defined as a conjunction of other predicates. For example, the Datalog rule $D(w, z) :- A(w, x), B(x, y), C(y, z).$ says that "$D(w, z)$ is true if $A(w, x)$, $B(x, y)$, and $C(y, z)$ are all true."

5.2.3 Basics of Reflection Resolution Using Points-To Information

The algorithm for computing targets of reflective calls is naturally expressed in terms of Datalog queries. Below we define Datalog rules to resolve targets of `Class.newInstance` and `Class.forName` calls. Handling of constructors, methods, and fields proceed similarly, as described in Sections 5.2.4–5.2.6. To disambiguate relations introduced for reflection resolution from input relations, we use Java identified naming conventions for the former.

To compute reflective targets of calls to `Class.newInstance`, we define two Datalog relations. Relation *classObjects* contains pairs $\langle i, t \rangle$ of invocations sites $i \in I$ calling `Class.forName` and types $t \in T$ that may be returned from

the call. We define *classObjects* using the following Datalog rule:

$$\begin{aligned} \mathit{classObjects}(i,t) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Class.forName}\text{”}), \\ & \mathit{actual}(i,1,v), \mathit{points\text{-}to}(v,s), \\ & \mathit{string2class}(s,t). \end{aligned}$$

The Datalog rule for *classObjects* reads as follows. Invocation site i returns an object of type t if the call graph relation *calls* contains an edge from i to “`Class.forName`”, parameter 1 of i is v, v points to s, and s is a string that represents the name of type t.

Relation *newInstanceTargets* contains pairs $\langle i,t \rangle$ of invocation sites $i \in I$ calling `Class.newInstance` and classes $t \in T$ that may be reflectively invoked by the call. The Datalog rule to compute *newInstanceTargets* is:

$$\begin{aligned} \mathit{newInstanceTargets}(i,t) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Class.newInstance}\text{”}), \\ & \mathit{actual}(i,0,v), \mathit{points\text{-}to}(v,c), \\ & \mathit{points\text{-}to}(v_c,c), \mathit{ret}(i_c,v_c), \\ & \mathit{classObjects}(i_c,t). \end{aligned}$$

The rule reads as follows. Invocation site i returns a new object of type t if the call graph relation *calls* contains an edge from i to `Class.newInstance`, parameter 0 of i is v, v is aliased to a variable v_c that is the return value of invocation site i_c, and i_c returns type t. Targets of `Class.forName` calls are resolved and calls to the appropriate class constructors are added to the invocation relation *calls*:

$$\mathit{calls}(i,m) \ :- \ \mathit{classObjects}(i,t), m = t + \text{“}.<\texttt{clinit}>\text{”}.$$

(The “+” sign indicates string concatenation.) Similarly, having computed relation $\mathit{newInstanceTargets}(i,t)$, we add these reflective call targets invoking the appropriate type constructor to the call graph relation *calls* with the rule below:

$$\mathit{calls}(i,m) \ :- \ \mathit{newInstanceTargets}(i,t), m = t + \text{“}.<\texttt{init}>\text{”}.$$

In Sections 5.2.4–5.2.6 we cover other ways to perform reflective operations.

5.2.4 Handling Reflective Constructor Calls: `Constructor` Objects

Another technique of reflective object creation is to use `Class.getConstructor` to get a `Constructor` object, and then calling `newInstance` on that. We define

a relation *constructorTypes* that contains pairs $\langle i, t\rangle$ of invocations sites $i \in I$ calling `Class.getConstructor` and types $t \in T$ of the type of the constructor:

$$\begin{aligned} \mathit{constructorTypes}(i,t) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Class.getConstructor}\text{”}),\\ & \mathit{actual}(i,0,v), \mathit{points\text{-}to}(v,h),\\ & \mathit{classObjects}(h,t). \end{aligned}$$

Once we have computed *constructorTypes*, we can compute more *newInstanceTargets* as follows:

$$\begin{aligned} \mathit{newInstanceTargets}(i,t) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Class.newInstance}\text{”}),\\ & \mathit{actual}(i,0,v), \mathit{points\text{-}to}(v,c), \mathit{points\text{-}to}(v_c,c),\\ & \mathit{ret}(i_c,v_c), \mathit{constructorTypes}(i_c,t). \end{aligned}$$

This rule says that invocation site i calling method “`Class.newInstance`” returns an object of type t if parameter 0 of i is v, v is aliased to the return value of invocation i_c which calls “`Class.getConstructor`”, and the call to i_c is on type t.

In a similar manner, we add support for `Class.getConstructors`, along with support for reflective field, and method accesses. The specification of these are straightforward and we do not describe them here.

5.2.5 Handling Reflective Invocations: `Method` Objects

Auxiliary relation *getMethod* defines two standard ways reflection API ways to reflectively invoke a method.

$$\begin{aligned} & \mathit{getMethod}(\text{“}\texttt{Class.getMethod(String, Class}\text{”}).\\ & \mathit{getMethod}(\text{“}\texttt{Class.getDeclaredMethod(String, Class}\text{”}). \end{aligned}$$

Relation $\mathit{methodObject}(h,m)$ defines when a heap-allocated object $h \in H$ represents a method $m \in M$. Relation $\mathit{resolvedInvoke}(i,m)$ defines when a method $m \in M$ can be called from an invocation site $i \in I$.

$$\begin{aligned} \mathit{methodObject}(h,m) \ :- \ & \mathit{getMethod}(m_i), \mathit{calls}(i,m_i), \mathit{freshi2h}(i,h),\\ & \mathit{actual}(i,1,v), \mathit{points\text{-}to}(v,h_m),\\ & \mathit{string2method}(h_m,m).\\ \mathit{resolvedInvoke}(i,m) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Method.invoke(Object, Object[]}\text{”}),\\ & \mathit{actual}(i,0,v), \mathit{points\text{-}to}(v,h), \mathit{methodObject}(h,m). \end{aligned}$$

Finally, input relations are updated with the results of method invocation resolution to represent the flow of data through parameters and return values:

$$\begin{aligned}
\mathit{assign}(v_1, v_2) \quad &:- \quad \mathit{resolvedInvoke}(i, m), \\
&\quad\; \mathit{formal}(m, 0, v_1), \mathit{actual}(i, 1, v_2). \\
\mathit{assign}(v_1, v_2) \quad &:- \quad \mathit{resolvedInvoke}(i, m), \mathit{ret}(i, v_1), \mathit{mret}(m, v_2). \\
\mathit{points\text{-}to}(v, h) \quad &:- \quad \mathit{resolvedInvoke}(i, m), \mathit{formal}(m, z, v), z > 0, \\
&\quad\; \mathit{actual}(i, 2, v_2), \mathit{points\text{-}to}(v_2, h_2), \\
&\quad\; \mathit{hpoints\text{-}to}(h_2, \text{“}\texttt{null}\text{”}, h).
\end{aligned}$$

5.2.6 Handling Reflective Field Accesses: `Field` Objects

Resolution of `Field` objects is similar to `Method`s. Auxiliary relation *getField* defines two standard ways reflection API ways to reflectively access a field.

$$\begin{aligned}
&\mathit{getField}(\text{“}\texttt{Class.getField(String)}\text{”}). \\
&\mathit{getField}(\text{“}\texttt{Class.getDeclaredField(String)}\text{”}).
\end{aligned}$$

Relation $\mathit{fieldObject}(h, f)$ defines when a heap-allocated object $h \in H$ represents a field $f \in F$.

$$\begin{aligned}
\mathit{fieldObject}(h, f) \quad &:- \quad \mathit{getField}(m_f), \mathit{calls}(i, m_f), \\
&\quad\; \mathit{freshi2h}(i, h), \mathit{actual}(i, 1, v), \mathit{points\text{-}to}(v, h_f), \\
&\quad\; \mathit{string2field}(h_f, f).
\end{aligned}$$

Finally, *load* and *store* relations are updated to represent the newly discovered reflective field accesses:

$$\begin{aligned}
\mathit{store}(v_1, f, v_2) \quad &:- \quad \mathit{calls}(i, \text{“}\texttt{Field.set(Object,Object)}\text{”}), \\
&\quad\; \mathit{actual}(i, 0, v), \mathit{points\text{-}to}(v, h), \mathit{fieldObject}(h, f), \\
&\quad\; \mathit{actual}(i, 1, v_1), \mathit{actual}(i, 2, v_2). \\
\mathit{load}(v_1, f, v_2) \quad &:- \quad \mathit{calls}(i, \text{“}\texttt{Field.get(Object,Object)}\text{”}), \\
&\quad\; \mathit{actual}(i, 0, v), \mathit{points\text{-}to}(v, h), \mathit{fieldObject}(h, f), \\
&\quad\; \mathit{actual}(i, 1, v_1), \mathit{ret}(i, v_2).
\end{aligned}$$

Notice that because our relations describing the input program only represent *objects* and effectively ignore primitive values, we do not need to model other field accessors such as `Field.getByte`/`Field.setByte`, etc. listed in Figure 4.

```
loadImpl() @ 43 InetAddress.java:1231       => java.net.Inet4AddressImpl
loadImpl() @ 43 InetAddress.java:1231       => java.net.Inet6AddressImpl
...
lookup() @ 86 AbstractCharsetProvider.java:126 => sun.nio.cs.ISO_8859_15
lookup() @ 86 AbstractCharsetProvider.java:126 => sun.nio.cs.MS1251
...
tryToLoadClass() @ 29 DataFlavor.java:64       => java.io.InputStream
...
```

Figure 16: A fragment of a specification file accepted by our system. A string identifying a call site to `Class.forName` is mapped to a class name that that call may resolve to.

5.2.7 Specification Points and User-Provided Specifications

Using a points-to analysis also allows us to determine, when a non-constant string is passed to a call to `Class.forName`, the *provenance* of that string. The *provenance* of a string is in essence a backward data slice showing the flow of data to that string. Provenance allows us to compute *specification points*—places in the program where external sources are read by the program from a configuration file, system properties, etc. For each specification point, the user can provide values that may be passed into the application.

Our implementation accepts specification files that contain a simple textual map of a specification point to the constant strings it can generate. A specification point is represented by a method name, bytecode offset, and the relevant line number. An example of a specification file is shown in Figure 16.

We compute the provenance by propagating through the assignment relation *assign*, aliased loads and stores, and string operations. To make the specification points as close to external sources as possible, we perform a simple analysis of strings to do backward propagation through string concatenation operations. For brevity, we only list the `StringBuffer.append` method used by the Java compiler to expand string concatenation operations here; other string operations work in a similar manner. The following rules for relation *leadsToForName* detail provenance propagation:

$$
\begin{array}{lll}
\mathit{leadsToForName}(v,i) & :- & \mathit{calls}(i, \text{“}\texttt{Class.forName}\text{”}), \mathit{actual}(i,1,v).\\
\mathit{leadsToForName}(v_2,i) & :- & \mathit{leadsToForName}(v_1,i), \mathit{assign}(v_1,v_2).\\
\mathit{leadsToForName}(v_2,i) & :- & \mathit{leadsToForName}(v_1,i),\\
 & & \mathit{load}(v_3,f,v_1), \mathit{points\text{-}to}(v_3,h_3),\\
 & & \mathit{points\text{-}to}(v_4,h_3), \mathit{store}(v_4,f,v_2).\\
\mathit{leadsToForName}(v_2,i) & :- & \mathit{leadsToForName}(v_1,i), \mathit{ret}(i_2,v_1),\\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2,0,v_2).\\
\mathit{leadsToForName}(v_2,i) & :- & \mathit{leadsToForName}(v_1,i), \mathit{ret}(i_2,v_1),\\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2,1,v_2).\\
\mathit{leadsToForName}(v_2,i) & :- & \mathit{leadsToForName}(v_1,i), \mathit{actual}(i_2,0,v_1),\\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2,1,v_2).
\end{array}
$$

To compute the specification points necessary to resolve `Class.forName` calls, we find endpoints of the *leadsToForName* propagation chains that are *not* string constants that represent class names. Relation $\mathit{lhs}(v)$ defines when variable v is ever assigned to in the program.

$$
\begin{array}{lll}
\mathit{lhs}(v) & :- & \mathit{assign}(v,\_).\\
\mathit{lhs}(v) & :- & \mathit{calls}(i, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{ret}(i,v).\\
\mathit{lhs}(v) & :- & \mathit{calls}(i, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i,0,v).\\
\mathit{lhs}(v) & :- & \mathit{calls}(i, \text{“}\texttt{StringBuffer.toString}\text{”}), \mathit{ret}(i,v).\\
\mathit{lhs}(v) & :- & \mathit{calls}(i, \text{“}\texttt{new StringBuffer}\text{”}), \mathit{actual}(i,0,v).\\
\mathit{lhs}(v) & :- & \mathit{load}(v_3,f,v), \mathit{points\text{-}to}(v_3,h), \mathit{points\text{-}to}(v_2,h), \mathit{store}(v_2,f,\_).
\end{array}
$$

Relation $\mathit{isTypeString}(v)$ below holds for variables v that refer to class name constants:

$$\mathit{isTypeString}(v) \quad :- \quad \mathit{points\text{-}to}(v,h), \mathit{string2class}(h,\_).$$

Finally, relation $\mathit{specPts}(v,i)$ defines when variable v is a specification point for a call to `Class.forName` at call site i:

$$\mathit{specPts}(v,i) \quad :- \quad \mathit{leadsToForName}(v,i), \neg\mathit{lhs}(v), \neg\mathit{isTypeString}(v).$$

Variables in *specPts* are often return results of calls to `System.getProperty` in the case of reading from a system property or `BufferedReader.readLine` in

the case of reading from a file. By specifying the possible values at that point that are appropriate for the application being analyzed, the user can construct a complete call graph.

5.2.8 Dealing with Other Reflective Calls

Unfortunately, `Class.forName` calls discussed in detail in the previous section are not the only ones whose results cannot sometimes be resolved without a user-provided specification. According to the definition of *newInstanceTargets* in Section 5.2.3, calls to `Class.newInstance` should all be fully resolved as long as the underlying `Class` and `Constructor` objects are fully resolved.

Similarly, according to the rules in Sections 5.2.5 and 5.2.6, method invocations and field accesses may not be fully resolved either because the underlying field or method objects are not fully resolved or because the method or field names are not fully resolved. Since the underlying objects come from `Class.forName` calls, they will be dealt with when we consider `Class.forName` resolution. However, method and field names can be not fully resolved and we need to address these cases separately by adding to the *specPts* relation.

$$
\begin{array}{lcl}
\mathit{leadsToInvoke}(v, i) & :- & \mathit{calls}(i, \mathit{getMethod}(m)), \mathit{actual}(i, 1, v). \\
\mathit{leadsToInvoke}(v_2, i) & :- & \mathit{leadsToInvoke}(v_1, i), \mathit{assign}(v_1, v_2). \\
\mathit{leadsToInvoke}(v_2, i) & :- & \mathit{leadsToInvoke}(v_1, i), \\
 & & \mathit{load}(v_3, f, v_1), \mathit{points\text{-}to}(v_3, h_3), \\
 & & \mathit{points\text{-}to}(v_4, h_3), \mathit{store}(v_4, f, v_2). \\
\mathit{leadsToInvoke}(v_2, i) & :- & \mathit{leadsToInvoke}(v_1, i), \mathit{ret}(i_2, v_1), \\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2, 0, v_2). \\
\mathit{leadsToInvoke}(v_2, i) & :- & \mathit{leadsToInvoke}(v_1, i), \mathit{ret}(i_2, v_1), \\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2, 1, v_2). \\
\mathit{leadsToInvoke}(v_2, i) & :- & \mathit{leadsToInvoke}(v_1, i), \mathit{actual}(i_2, 0, v_1), \\
 & & \mathit{calls}(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), \mathit{actual}(i_2, 1, v_2).
\end{array}
$$

Relation *isMethodString* below defines all string variables that represent method names:

$$\mathit{isMethodString}(v) \quad :- \quad \mathit{points\text{-}to}(v, h), \mathit{string2method}(h, \_).$$

Finally, we all the necessary $\langle v \in V, i \in I \rangle$ pairs to relation *specPts*:

$$specPts(v, i) :- leadsToInvoke(v, i), \neg lhs(v), \neg isMethodString(v).$$

Finding additional specification points caused by unresolved field accesses proceeds similarly. The relevant rules are shown below:

$$\begin{array}{lll}
leadsToField(v, i) & :- & calls(i, getField(m)), actual(i, 1, v). \\
leadsToField(v, i) & :- & calls(i, \text{“}\texttt{Class.getDeclaredField}\text{”}), actual(i, 1, v). \\
leadsToField(v_2, i) & :- & leadsToField(v_1, i), assign(v_1, v_2). \\
leadsToField(v_2, i) & :- & leadsToField(v_1, i), \\
 & & load(v_3, f, v_1), points\text{-}to(v_3, h_3), \\
 & & points\text{-}to(v_4, h_3), store(v_4, f, v_2). \\
leadsToField(v_2, i) & :- & leadsToField(v_1, i), ret(i_2, v_1), \\
 & & calls(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), actual(i_2, 0, v_2). \\
leadsToField(v_2, i) & :- & leadsToField(v_1, i), ret(i_2, v_1), \\
 & & calls(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), actual(i_2, 1, v_2). \\
leadsToField(v_2, i) & :- & leadsToField(v_1, i), actual(i_2, 0, v_1), \\
 & & calls(i_2, \text{“}\texttt{StringBuffer.append}\text{”}), actual(i_2, 1, v_2). \\
 & & \\
isFieldString(v) & :- & points\text{-}to(v, h), field2method(h, _). \\
 & & \\
specPts(v, i) & :- & leadsToField(v, i), \neg lhs(v), \neg isFieldString(v).
\end{array}$$

5.3 Precision of Points-to Results

Accurate interprocedural pointer alias analysis is critical to the precision of our reflection analysis, because `Class` objects can be passed around the program, placed in collections, etc. to be used later for object creation. The pointer analysis our approach implements is a variation of the inclusion-based pointer analysis by Whaley and Lam [WL04]. Being an inclusion-based analysis, it is more precise than equivalence-based ones as it allows two aliased pointers to point to overlapping but different sets of locations.

The default approach to object naming in this analysis uses object allocation sites as an approximation of object identity. This approach yields insufficient precision in the case of collection classes; all instances of a collection class typically use the same allocation site to store the objects

put into the collection. Introducing new names for these internal data structures of the default collection classes greatly improves analysis precision [LL05a, LL05b, MRR02].

5.4 Reflection Resolution Using Casts

For some applications, the task of providing reflection specifications may be too heavy a burden. Fortunately, we can leverage the type cast information present in the program to automatically determine a conservative approximation of possible reflective targets. Consider, for instance, the following typical code snippet:

```
1.   Object o = c.newInstance();
2.   String s = (String) o;
```

The cast in statement 2 *post-dominates* the call to `Class.newInstance` in statement 1. This implies that all execution paths that pass through the call to `Class.newInstance` must also go through the cast in statement 2 [ASU86]. For statement 2 not to produce a runtime exception, `o` must be a subclass of `String`. Thus, only subtypes of `String` can be created as a result of the call to `newInstance`. More generally, if the result of a `newInstance` call is *always* cast to type t, we say that only subtypes of t can be instantiated at the call to `newInstance`.

Relying on cast operations can possibly be unsound as the cast may fail, in which case, the code will throw a `ClassCastException`. Thus, in order to work, our cast-based technique relies on Assumption 4.3, the correctness of cast operations.

5.4.1 Preparing Subtype Information

We rely on the closed world Assumption 4.2 described in Section 4 to find the set of all classes possibly used by the application. The classes available at analysis time are generally distributed with the application. However, occasionally, there are classes that are generated when the application is compiled or deployed, typically with the help of an Ant script. Therefore, we generate the set of possible classes *after* deploying the application.

We pre-process all resulting classes to compute the subtyping relation $subtype(t_1, t_2)$ that determines when t_1 is a subtype of t_2. Preprocessing even the smallest applications involved looking at many thousands of classes

because we consider all the default jars that the Java runtime system has access to. We run this preprocessing step off-line and store the results for easy access.

5.4.2 Using Cast Information

We integrate the information about cast operations directly into the system of constraints expressed in Datalog. We use a Datalog relation *subtype* described above, a relation *cast* that holds the cast operations, and a relation *unresolved* that holds the unresolved calls to `Class.forName`. The following Datalog rule uses cast operations applied to the return result v_{ret} of a call i to `Class.newInstance` to constrain the possible types t_c of `Class` objects c returned from calls sites i_c of `Class.forName`:

$$\begin{aligned} \mathit{classObjects}(i_c, t) \ :- \ & \mathit{calls}(i, \text{“}\texttt{Class.newInstance}\text{”}), \mathit{actual}(i, 0, v), \\ & \mathit{points\text{-}to}(v, c), \mathit{ret}(i, v_{ret}), \\ & \mathit{cast}(\_, t_c, v_{ret}), \mathit{subtype}(t, t_c), \\ & \mathit{unresolved}(i_c), \mathit{points\text{-}to}(v_c, c), \mathit{ret}(i_c, v_c). \end{aligned}$$

Information propagates both forward and backward—for example, casting the result of a call to `Class.newInstance` constrains the `Class` object it is called upon. If the same `Class` object is used in another part of the program, the type constraint derived from the cast will be obeyed.

5.4.3 Problems with Using Casts

Casts are sometimes inadequate for resolving calls to `Class.newInstance` for the following reasons. First, the cast-based approach is inherently imprecise because programs often cast the result of `Class.newInstance` to a very wide type such as `java.io.Serializable`. This produces a lot of *potential* subclasses, only some of which are relevant in practice. Second, as our experiments show, not all calls to `Class.newInstance` have post-dominating casts, as illustrated by the following example.

Example 8. As shown in Figure 17, one of our benchmark applications, `freetts`, places the object returned by `Class.newInstance` into a vector `voiceDirectories` (line 5). Despite the fact that the objects are subsequently cast to type `VoiceDirectory[]` on line 8, intraprocedural post-dominance is not powerful enough to take this cast into account. □

```
1.      UniqueVector voiceDirectories = new UniqueVector();
2.      for (int i = 0; i < voiceDirectoryNames.size(); i++) {
3.          Class c = Class.forName((String) voiceDirectoryNames.get(i),
4.                                                  true, classLoader);
5.          voiceDirectories.add(c.newInstance());
6.      }
7.
8.      return (VoiceDirectory[]) voiceDirectories.toArray(new
9.                      VoiceDirectory[voiceDirectories.size()]);
```

Figure 17: A case in `freetts` where our analysis is unable to determine the type of objects instantiated on line 5 using casts.

Using cast information significantly reduces the need for user-provided specification in practice. While the version of the analysis that does not use cast information can be made fully sound with user specification as well, we chose to only provide a specification for the cast-based version.

SECTION 6

Experimental Results

In this section we present a comprehensive experimental evaluation of the static analysis approaches presented in Section 5. In Section 6.1 we describe our experimental setup. Section 6.2 presents an overview our experimental results. Section 6.3 presents our baseline local reflection analysis. In Sections 6.4 and 6.5 we discuss the effectiveness of using the points-to and cast-based reflection resolution approaches, respectively. Section 6.6 describes the specifications needed to obtain a sound call graph approximation. Section 6.7 compares the overall sizes of the call graph for the different analysis versions presented in this section.

6.1 Experimental Setup

We performed our experiments on a suite of six large, widely-used open-source Java benchmark applications. These applications were selected among the most popular Java projects available on SourceForge. We believe that real-life applications like these are more representative of how programmers use reflection than synthetically created test suites, or SPEC JVM benchmarks, most of which avoid reflection altogether.

Summary of information about the applications is provided in Figure 18. Notice that the traditional lines of code size metric is somewhat misleading in

Benchmark	Description	Line count	File count	Jars	Available classes
`jgap`	genetic algorithms package	32,961	172	9	62,727
`freetts`	speech synthesis system	42,993	167	19	62,821
`gruntspud`	graphical CVS client	80,138	378	10	63,847
`jedit`	graphical text editor	144,496	427	1	62,910
`columba`	graphical email client	149,044	1,170	35	53,689
`jfreechart`	chart drawing library	193,396	707	6	62,885
Total		**643,028**	**3,021**	**80**	**368,879**

Figure 18: Summary of information about our benchmarks. Applications are sorted by the number of lines of code in column 3.

	None	Local			Points-to				Casts				Sound		
Benchmark	**T**	**T**	**FR**	**UR**	**T**	**FR**	**PR**	**UR**	**T**	**FR**	**PR**	**UR**	**T**	**FR**	**UR**
`jgap`	27	27	19	8	28	20	1	7	28	20	4	4	89	85	4
`freetts`	30	30	21	9	30	21	0	9	34	25	4	5	81	75	6
`gruntspud`	139	139	112	27	142	115	5	22	232	191	19	22	220	208	12
`jedit`	156	156	137	19	161	142	3	16	178	159	12	7	210	197	12
`columba`	104	105	89	16	105	89	2	14	118	101	10	7	173	167	6
`jfreechart`	104	104	91	13	104	91	1	12	149	124	10	15	169	165	4

Figure 19: Results of resolving `Class.forName` calls for different analysis versions.

the case of applications that rely on large libraries. Many of these benchmarks depend of massive libraries, so, while the application code may be small, the full size of the application executed at runtime is quite large. The last column of the table in Figure 18 lists the number of classes available by the time each application is deployed, including those in the JDK.

We ran all of our experiments on an Opteron 150 machine equipped with 4GB or memory running Linux. JDK version 1.4.2_08 was used. All of the running times for our preliminary implementation were in tens of minutes, which, although a little high, is acceptable for programs of this size. Creating subtype information for use with cast-based analysis took well under a minute.

6.2 Evaluation Approach

We have implemented five different variations of our algorithms: None, Local, Points-to, Casts, and Sound and applied them to the benchmarks described above. None is the base version that performs no reflection resolution; Local performs a simple local analysis, as described in Section 6.3. Points-to and Casts are described in Sections 5.2 and 5.4, respectively.

Version Sound is augmented with a user-provided specification to make the answer conservative. We should point out that only the Sound version provides results that are fully sound: None essentially assumes that reflective calls have no targets. Local only handles reflective calls that can be fully resolved within a single method. Points-to and Casts only provide targets for reflective calls for which either string or cast information constraining the possible targets is available and unsoundly assumes that the rest of the calls have no targets.

Figure 19 summarizes the results of resolving `Class.forName` using all five analysis versions. `Class.forName` calls represent by far the most common kind

of reflective operations and we focus on them in our experimental evaluation. To reiterate the definitions in Section 5, we distinguish between:

- *fully resolved calls* to `Class.forName` for which all potential targets are class name constants,
- *partially resolved calls*, which have at least one class name string constant propagating to them, and
- *unresolved calls*, which have no class name string constants propagating to them, only non-constant external sources requiring a specification.

The columns subdivide the total number of calls (T) into fully resolved calls (FR), partially resolved (PR), and unresolved (UR) calls. In the case of LOCAL analysis, there are no partially resolved calls — calls are either fully resolved to constant strings or unresolved. Similarly, in the case of SOUND analysis, all calls are either fully resolved or unresolved, as further explained in Section 6.5.

6.3 Local Analysis for Reflection Resolution (LOCAL)

To provide a baseline for comparison, we implemented a local intra-method analysis that identifies string constants passed to `Class.forName`. This analysis catches only those reflective calls that can be resolved completely within a single method. Because this technique does not use interprocedural points-to results, it cannot be used for identification of specification points. Furthermore, because for method invocations and field accesses the names of the method or field are typically *not* locally defined constants, we do not perform resolution of method calls and field accesses in LOCAL.

A significant percentage of `Class.forName` calls can be fully resolved by local analysis, as demonstrated by the numbers in column 4, Figure 19. This is partly due to the fact that it is actually quite common to call `Class.forName` with a constant string parameter for side-effects of the call, because doing so invokes the class constructor. Another common idiom contributing the number of calls resolved by local analysis is `T.class`, which is converted to a call to `Class.forName` and is *always* statically resolved.

6.4 Points-to & Reflection Resolution (POINTS-TO)

Points-to information is used to find targets of reflective calls to `Class.forName`, `Class.newInstance`, `Method.invoke`, etc. As can be seen from Figure 19, for all of the benchmarks, POINTS-TO information results in more resolved `Class.forName` calls and fewer unresolved ones compared to LOCAL.

6.4.1 Specification Points

Quite frequently, some sort of specification is required for reflective calls to be fully resolved. Points-to information allows us to provide the user with a list of specification points where inputs needs to be specified for a conservative answer to be obtained. Among the specification points we have encountered in our experiments, calls to `System.getProperty` to retrieve a system variable and calls to `BufferedReader.readLine` to read a line from a file are quite common. Below we provide a typical example of providing a specification.

Example 9. This example describes resolving reflective targets of a call to `Class.newInstance` in `javax.xml.transform.FactoryFinder` in the JDK in order to illustrate the power and limitation of using points-to information. Class `FactoryFinder` has a method `Class.newInstance` shown in Figure 20. The call to `Class.newInstance` occurs on line 9. However, the exact class instantiated at runtime depends on the `className` parameter, which is passed into this function. This function is invoked from a variety of places with the `className` parameter being read from initialization properties files, the console, etc. In only one case, when `Class.newInstance` is called from another function `find` located in another file, is the `className` parameter a string constant.

This example makes the power of using points-to information apparent — the `Class.newInstance` target corresponding to the string constant is often difficult to find by just looking at the code. The relevant string constant was passed down through several levels of method calls located in a different file; it took us more that five minutes of exploration with a powerful code browsing tool to find this case in the source. Resolving this `Class.newInstance` call also requires the user to provide input for four specification points: along with a constant class name, our analysis identifies two specification points, which correspond to file reads, one access of system properties, and another read from a hash table. □

In most cases, the majority of calls to `Class.forName` are fully resolved. However, a small number of unresolved calls are potentially responsible for a

```
 1. private static Object newInstance(String className,
 2.                 ClassLoader classLoader) throws ConfigurationError {
 3.     try {
 4.         Class spiClass;
 5.         if (classLoader == null) {
 6.             spiClass = Class.forName(className);
 7.         }
 8.         ...
 9.         return spiClass.newInstance();
10.     } catch (...)
11.         ...
12.     }
```

Figure 20: Reflection resolution using points-to results in `javax.xml.transform.FactoryFinder` in the JDK.

large number of specification points the user has to provide. For POINTS-TO, the average number of specification points per invocation site ranges from 3 for `freetts` to 9 for `gruntspud`. However, for `jedit`, the average number of specification points is 422. Specification points computed by the pointer analysis-based approach can be thought of as "hints" to the user as to where provide specification.

In most cases, the user is likely to provide specification at program input points where he knows what the input strings may be. This is because at a reflective call it may be difficult to tell what all the constant class names that flow to it may be, as illustrated by Example 9. Generally, however, the user has a choice. For problematic reflective calls like those in `jedit` that produce a high number of specification points, a better strategy for the user may be to provide reflective specifications at the `Class.forName` *calls themselves* instead of laboriously going through all the specification points.

6.5 Casts & Reflection Resolution (CASTS)

Type casts often provide a good first static approximation to what objects can be created at a given reflective creation site. There is a pretty significant increase in the number of `Class.forName` calls reported in Figure 19 in a few cases, including 93 newly discovered `Class.forName` calls in `gruntspud` that apprear due to a bigger call graph when reflective calls are resolved. In all cases, the majority of `Class.forName` calls have their targets at least partially

resolved. In fact, as many as 95% of calls are resolved in the case of `jedit`.

As our experience with the Java reflection APIs would suggest, most `Class.newInstance` calls are post-dominated by a cast operation, often located within only a few lines of code of the `Class.newInstance` call. However, in our experiments, we have identified a number of `Class.newInstance` call sites, which were not dominated by a cast of any sort and therefore the return result of `Class.newInstance` could not be constrained in any way. As it turns out, most of these unconstrained `Class.newInstance` call sites are located in the JDK and `sun.*` sources, Apache libraries, etc. Very few were found in application code.

The high number of unresolved calls in the JDK is due to the fact that reflection use in libraries tends to be highly generic and it is common to have "`Class.newInstance` wrappers" — methods that accept a class name as a string and return an object of that class, which is later cast to an appropriate type in the caller method. Since we rely on *intraprocedural* post-dominance, resolving these calls is beyond our scope. However, since such "wrapper" methods are typically called from multiple invocation sites and different sites can resolve to different types, it is unlikely that a precise approximation of the object type returned by `Class.newInstance` is possible in these cases at all.

6.5.1 Precision of Cast Information

Many reflective object creation sites are located in the JDK itself and are present in all applications we have analyzed. For example, method `lookup` in package `sun.nio.cs.AbstractCharsetProvider` reflectively creates a subclass of `Charset` and there are 53 different character sets defined in the system. In this case, the answer is precise because all of these charsets can conceivably be used depending on the application execution environment. In many cases, the cast approach is able to *uniquely* pinpoint the target of `Class.newInstance` calls based on cast information. For example, there is only one subclass of class `sun.awt.shell.ShellFolderManager` available to `gruntspud`, so, in order for the cast to succeed, it must be instantiated.

In general, however, the cast-based approach provides an imprecise upper bound on the call graph that needs to be analyzed. Because the results of `Class.newInstance` are occasionally cast to very wide types, such as `java.lang.Cloneable`, many potential subclasses can be instantiated at the `Class.newInstance` call site. The cast-based approach is likely to yield more precise results on applications that use Java generics, because those applica-

tions tend to use more narrow types when performing type casts.

6.6 A Sound Call Graph Approximation (SOUND)

Providing a specification for unresolved reflective calls allows us to achieve a sound approximation of the call graph. In order to estimate the amount of effort required to come up with a specification for unresolved reflective calls, we decided to start with POINTS-TO and add a reflection specification until the result became sound. Because providing a specification allows us to discover more of the call graph, two or three rounds of specification were required as new portions of the program became available. In practice, we would start without a specification and examine all unresolved calls and specification points corresponding to them. Then we would come up with a specification and feed it back to the call graph construction algorithm until the process converges.

Coming up with a specification is a difficult and error-prone task that requires looking at a large amount of source code. It took us about ten hours to incrementally devise an appropriate specification and ensure its completeness by rerunning the call graph construction algorithm. After providing a reflection specification stringing with POINTS-TO, we then estimate how much of the user-provided specification can be avoided if we were to rely on type casts instead.

6.6.1 Specification Statistics

The first part of Figure 21 summarizes the effort needed to provide specifications to make the call graph sound. The second column shows the number of specifications of the form

$$\textit{reflective call site} \;\Rightarrow\; \textit{type}$$

as exemplified by Figure 16. Columns 3–5 show the number of reflection calls sites covered by each specification, breaking them down into sites that located within library vs application code. As can be seen from the table, while the number of invocation sites for which specifications are necessary is always around 20, only a few are part of the application. Moreover, in the case of `jfreechart`, *all* of the calls requiring a specification are part of the library code.

	Starting with STRINGS					Starting with CASTS				
Benchmark	Specs	Sites	Libs	App	Types/site	Specs	Sites	Libs	App	Types/site
jgap	1,068	25	21	4	42.72	16	2	2	0	8.0
freetts	964	16	14	2	60.25	0	4	3	1	0.0
gruntspud	1,014	27	26	1	37.56	18	4	4	0	4.5
jedit	1,109	21	19	2	52.81	63	3	2	1	21.0
columba	1,006	22	21	1	45.73	16	2	2	0	8.0
jfreechart	1,342	21	21	0	63.90	18	4	4	0	4.5

Figure 21: User-provided specification statistics.

Since almost all specification points are located in the JDK and library code, specification can be shared among different applications. Indeed, there are only 40 *unique* invocation sites requiring a specification across all the benchmarks. Column 6 shows the average number of types specified per reflective call site. Numbers in this columns are high because most reflective calls within the JDK can refer to a multitude of implementation classes.

The second part of Figure 21 estimates the specification effort required if were were to start with a cast-based call graph construction approach. As can be seen from columns 8–10, the number of `Class.forName` calls that are not constrained by a cast operation is quite small. There are, in fact, only 14 unique invocation sites — or about a third of invocation sites required for POINTS-TO. This suggests that the the effort required to provide a specification to make CASTS sound is considerably smaller than our original effort that starts with POINTS-TO.

6.6.2 Specification Difficulties

In some cases, determining meaningful values to specify for `Class.forName` results is quite difficult, as shown by the example below. One such problematic example was the BeanShell interpreter in `jedit` first described in Section 3.7.

Example 10. In order to come up with a conservative superset of classes that may be invoked by the BeanShell interpreter for a *given installation* of `jedit`, we parse the scripts that are supplied with `jedit` to determine imported Java classes they have access to. (We should note that this specification is only sound for the default configuration of `jedit`; new classes may need to be added to the specification if new macros become available.)

It took us a little under an hour to develop appropriate Perl scripts to do the parsing of 125 macros supplied with `jedit`. The `Class.forName` call can instantiate a total of 65 different types, which is, of course, an improvement over an overly conservative approximation that assumes that *any* class in the system may be instantiated. □

We should emphasize that the conservativeness of the call graph depends on the conservativeness of the user-provided specification. If the specification missed potential relations, they will be also omitted from the call graph. Furthermore, a specification is typically only conservative for a given configuration of an application: if initialization files are different for a different program installation, the user-provided specification may no longer be conservative.

6.6.3 Remaining Unresolved Calls

Somewhat surprisingly, there are *still* some `Class.forName` calls that are not fully resolved given a user-provided specification, as can be seen from the last column in Figure 19. In fact, this is not a specification flaw: no valid specification is *possible* for those cases, as explained below.

Example 11. The audio API in the JDK includes method `javax.sound.sampled.AudioSystem.getDefaultServices`, which is not called in Java version 1.3 and above. A `Class.forName` call within that method resolves to constant `com.sun.media.sound.DefaultServices`, however, this class is absent in post-1.3 JDKs. However, since this method represents dead code, our answer is still sound. Similarly, other unresolved calls to `Class.forName` located within code that is not executed for the particular application configuration we are analyzing refer to classes specific to MacOS and unavailable on Linux, which is the platform we performed analysis on. In other cases, classes were unavailable for JDK version 1.4.2_08, which is the version we ran our analysis on. □

6.7 Effect of Reflection Resolution on Call Graph Size

Figure 22 compares the number of classes and methods across different analysis versions. Local analysis does not have any significant effect on the number of methods or classes in the call graph, even though most of the calls to `Class.forName` can be resolved with local analysis. This is due to the fact

that the vast majority of these calls are due to the use of the `T.class` idiom, which typically refer to classes that are already within the call graph. While these trivial calls are easy to resolve, it is the analysis of the other "hard" calls with a lot of potential targets that leads to a substantial increase in the call graph size.

Using Points-to increases the number of classes and methods in the call graph only moderately. The biggest increase in the number of methods occurs for `jedit` (293 methods). Using Casts leads to significantly bigger call graphs, especially for `gruntspud`, where the increase in the number of methods compared to None is almost two-fold.

The most noticeable increase in call graph size is observed in version Sound. Compared to None, the average increase in the number of classes is 3.2 times the original and the average increase for the number of methods is 3 times the original. The biggest increase in the number of methods occurs in `gruntspud`, with over 7,000 extra methods added to the graph.

Figure 22 also demonstrate that the lines of code metric is not always indicative of the size of the final call graph — programs are listed in the

Classes						
Benchmark	None	Local	Points-to	Casts	Sound	
`jgap`	264	264	268	276	1,569	5.94
`freetts`	309	309	309	351	1,415	4.58
`gruntspud`	1,258	1,258	1,275	2,442	2,784	2.21
`jedit`	1,660	1,661	1,726	2,152	2,754	1.66
`columba`	961	962	966	1,151	2,339	2.43
`jfreechart`	884	881	886	1,560	2,340	2.65

Methods						
Benchmark	None	Local	Points-to	Casts	Sound	
`jgap`	1,013	1,014	1,038	1,075	6,676	6.58
`freetts`	1,357	1,358	1,358	1,545	5,499	4.05
`gruntspud`	7,321	7,321	7,448	14,164	14,368	1.96
`jedit`	11,230	11,231	11,523	13,487	16,003	1.43
`columba`	5,636	5,642	5,652	6,199	12,001	2.13
`jfreechart`	5,374	5,374	5,392	8,375	12,111	2.25

Figure 22: Number of classes and methods in the call graph for different analysis versions.

	None		Local		Points-to		Casts		Sound	
Benchmark	**Solver**	**Total**	**Solver**	**Total**	**Solver**	**Total**	**Solver**	**Total**	**Solver**	**Total**
`jgap`	30	54	34	61	43	73	477	692		
`freetts`	16	32	19	35	23	39	180	250		
`gruntspud`	193	268	239	318	392	481	5,702	5,860		
`jedit`	836	1,236	983	1,394	1,925	2,401	0	7,300		
`columba`	106	158	125	179	173	237	1,161	1,456		
`jfreechart`	272	395	335	462	450	592	2,578	3,351		

Figure 23: Running times for different analysis versions, in seconds.

increasing order of line counts, yet, `jedit` and `gruntspud` are clearly the biggest benchmarks if we consider the method count. This can be attributed to the use of large libraries that ship with the application in binary form as well as considering a much larger portion of the JDK in version Sound compared to version None.

6.8 Running Times

Figure 23 presents the running times for the different versions our our analysis. For each analysis version, we specify the `bddbddb` solver time as well as the total wall clock time, which includes the time to load the relations and save the program representation and call graph information. The overhead involved in these steps can range from 38% to over 100%. When there is a client analysis that uses the results of reflection resolution, these additional saving operations can be avoided by running the client analysis together with reflection resolution while the relevant relations are still within `bddbddb`.

While the version with no reflection resolution runs relatively fast, other analysis versions take considerably more time to complete.

SECTION 7

Related Work

General treatments of reflection in Java are given in Forman and Forman [FF04] and Guéhéneuc et al. [GCSD02]. The rest of the related work falls into the following broad categories: projects that explicitly deal with reflection in Java and other languages; approaches to call graph construction in Java; and finally, static and dynamic analysis algorithms that address the issue of dynamic class loading.

7.1 Reflection and Metadata Research

The metadata and reflection community has a long line of research originating in languages such as Scheme [Thi96]. We only mention a few highly relevant projects here. The closest static analysis project to ours we are aware of is the work by Braux and Noyé on applying partial evaluation to reflection resolution for the purpose of optimization [BN99]. Their paper describes extensions to a standard partial evaluator to offer reflection support. The idea is to "compile away" reflective calls in Java programs, turning them into regular operations on objects and methods, given constraints on the concrete types of the object involved. The type constraints for performing specialization are provided by hand.

Our static analysis can be thought of as a tool for inferring such constraints, however, as our experimental results show, in many cases targets of reflective calls cannot be uniquely determined and so the benefits of specialization to optimize program execution may be limited. Braux and Noyé present a description of how their specialization approach may work on examples extracted from the JDK, but lacks a comprehensive experimental evaluation. In related work for languages other than Java, Ruf explores the use of partial evaluation as an optimization technique in the context of CLOS [Ruf93]. Masuhara et al. explore the use of partial evaluation as applied to an abstract object-oriented language [MY98].

The issue of *specifying* reflective targets is explicitly addressed in Jax [TLSS99]. Similarly, the Spark pointer analysis implemented within the Soot compiler uses specifications of many reflective targets in the JDK during call graph construction [LH03]. Just like our technique, Spark also used on-the-fly call graph construction. Potential reflective call targets are

automatically added to the set of root methods in the beginning of the analysis. Unlike Spark, which comes with models of many `native` methods, our approach is oblivious to `native` methods. While this is generally unlikely, not handling such methods can render our results not fully sound.

Jax is concerned with reducing the size of Java applications in order to reduce download time; it reads in the class files that constitute a Java application, and performs a whole-program analysis to determine the components of the application that must be retained in order to preserve program behavior. Clearly, information about the true call graph is necessary to ensure that no relevant parts of the application are pruned away. Jax's approach to reflection is to employ user-provided specifications of reflective calls. While our framework also supports user-provided annotations, as illustrated in Section 6.4, determining targets of reflective calls can often be error-prone, if delegated to the user. To assist the user with writing complete specification files, Jax relies on dynamic instrumentation to discover the missing targets of reflective calls. Our analysis based on points-to information can be thought of as a tool for determining where to insert reflection specifications.

A precise analysis of strings by Christensen et al. mentions reflections as one of the potential uses of their approach [CMS03]. They treat `Class.forName` calls as "hotspots" for their analysis, then trying to determine what the exact values passed as parameters may be. Their approach, however, relies on an external pointer analysis to determine the propagation of strings throughout the program. The paper applies their approach to programs that are all under 4,000 lines long and lacks a detailed experimental evaluation of the precision of their approach. Their technique, however, can potentially address reflective calls that have much more complex string expressions passed as reflective arguments. For example, knowing that the argument of `Class.forName` must end in string `"Configuration"` will allow the analysis to substantially limit the number of possibly instantiated classes.

7.2 Call Graph Construction

A lot of effort has been spent of analyzing function pointers in C [EGH94, MRR01, MRR04] as well as virtual method calls in C++ [AH96, BS96, CG94, PR96] and Java [GC01, GDDC97, RRHK00, SHR+00, TP00]. They are described in more detail below.

7.2.1 Function Pointers in C

Emami et al. describe how a context-sensitive pointer analysis for C integrated with call graph construction in the presence of function pointers [EGH94]. Their approach introduces the notion of call graph discovery when the call graph is unavailable in advance.

Milanova et al. evaluate the precision of call graph construction in the presence of function pointers using an inexpensive pointer analysis approach [Zha98] and conclude that it is sufficient for most cases [MRR01, MRR04].

7.2.2 Virtual Calls in C++

Bacon et al. compare the "unique name", RTA, and CHA virtual call resolution approaches [BS96, Bac98]. They conclude that RTA is both fast and effective and able to resolve 71% of virtual calls on average.

Aigner and Hölzle investigate the effect virtual call elimination using CHA has on the runtime of large C++ programs and report a median 18% performance improvement over the original programs [AH96]. The number of virtual function calls is reduced by a median factor of five.

7.2.3 Virtual Calls in Java

Grove et al. present a parameterized algorithmic framework for call graph construction [GC01, GDDC97]. They empirically assess a multitude of call graph construction algorithms by applying them to a suite of medium-sized programs written in Cecil and Java. Their experience with Java programs suggests that the effect of using context sensitivity for the task of call graph construction in Java yields only moderate improvements.

Tip and Palsberg propose a propagation-based algorithm for call graph construction and investigate the design space between existing algorithms for call graph construction such as 0-CFA and RTA, including RA, CHA, and four new ones [TP00]. Sundaresan et al. go beyond the tranditional RTA and CHA approaches in Java and and use type propagation for the purpose of obtaining a more precise call graph [SHR+00]. Their approach of using variable type analysis (VTA) is able to uniquely determine the targets of potentially polymorphic call sites in 32% to 94% of the cases.

Agrawal et al. propose a demand-driven algorithm for call graph construction [ALS02]. Their work is motivated by the need for just-in-time or dynamic

compilation as well as program analysis used as part of software development environments. They demonstrate that their demand-driven technique has the same accuracy as the corresponding exhaustive technique. The reduction in the graph construction time depends upon the ratio of the cardinality of the set of influencing nodes to the set of all nodes.

Rayside et al. explore the effect various call graph construction techniques have on automatic clustering approaches used to extract the high level structure of the program under study [RRHK00]. They also used a slightly different notion of the call graph that supports weighted edges.

7.3 Dynamic Analysis Approaches

Our work is motivated to a large extend by the need of error detection tool to have a static approximation of the true conservative call graph of the application. This largely precludes dynamic analysis that benefits optimizations such as method inlining and connectivity-based garbage collection.

A recent paper by Hirzel, Diwan, and Hind addresses the issues of dynamic class loading, native methods, and reflection in order to deal with the full complexity of Java in the implementation of a common pointer analysis [HDH04]. Their approach involves converting the pointer analysis [And94] into an online algorithm: they add constraints between analysis nodes as they are discovered at runtime. Newly generated constraints cause re-computation and the results are propagated to analysis clients such as a method inliner and a garbage collector at runtime. Their approach leverages the class hierarchy analysis (CHA) to update the call graph. Our technique uses a more precise pointer analysis-based approach to call graph construction.

Their paper also contains a comprehensive overview of analysis approaches that address dynamic class loading. Here we briefly mention some of the highlights. However, none of the projects mentioned below fully address the issue of reflection.

The ARE tool presented in Gscwind et al. allows tracing of method parameter and return values at runtime for program comprehension [GOP03]. They point out that ignoring reflection leads to program traces that are incomplete. ARE instruments the program and collects data that allows it to provide targets of reflective method calls. These reflective targets are subsequently displayed by way of sequence diagrams.

Pechtchanski and Sarkar [PS01] present a framework for interprocedural whole-program analysis. They discuss how the analysis is triggered (when

newly loaded methods are compiled), and how to keep track of what to de-optimize (when optimistic assumptions are invalidated). Qian and Hendren [QH04] adapt Tip and Palsbergs XTA [TP00] to deal with dynamic class loading. The main contribution of their work is a low-overhead call edge profiler, which yields a precise call graph upon which XTA is based.

SECTION 8

Conclusions

This paper presents the first static analysis for call graph construction in Java that addresses reflective calls. Our algorithm uses the results of a points-to analysis to determine potential reflective call targets. When the calls cannot be fully resolved, user-provided specification is requested. As an alternative to providing specification, type cast information can be used to provide a conservative approximation of reflective call targets.

We applied our static analysis techniques to the task of constructing call graphs for six large Java applications, some consisting of more than 190,000 lines of code. Our evaluation showed that as many as 95% of reflective `Class.forName` could at least partially be resolved to statically determined targets with the help of points-to results and cast information *without* providing any specification.

While most reflective calls are relatively easy to resolve statically, *precisely* interpreting some reflective calls requires a user-provided specification. Our pointer analysis-based approach also identified specification points — places in the program corresponding to file and system property read operations, etc., where user input is needed in order to obtain a full call graph. Our evaluation showed that the construction of a specification that makes the call graph conservative is a time-consuming and error-prone task. Fortunately, our cast-based approach can drastically reduce the specification burden placed on the user by providing a conservative, albeit potentially imprecise approximation of reflective targets.

Our experiments confirmed that ignoring reflection results in missing significant portions of the call graph, which is not something that effective static analysis tools can afford. While the local and points-to analysis techniques resulted in only a moderate increase in call graph size, using the cast-based approach resulted in call graphs with as many as 1.5 times more methods than the original call graph. Furthermore, providing a specification resulted in much larger conservative call graphs that were almost 7 times bigger than the original. For instance, in one our benchmark, an additional 7,047 methods were discovered in the conservative call graph version that were not present in the original.

SECTION 9

Acknowledgements

This work was supported by NSF Grant No. 0326227 and an Intel Graduate Fellowship. We would like to thank Chris Unkel as well as the anonymous reviewers for helpful suggestions about improving this work.

References

[AH96] Gerald Aigner and Urs Hölzle. Eliminating virtual function calls in C++ programs. In *Proceedings of the 10th European Conference on Object-Oriented Programming*, pages 142–166. Springer-Verlag, 1996.

[ALS02] Gagan Agrawal, Jinqian Li, and Qi Su. Evaluating a demand driven technique for call graph construction. In *Computational Complexity*, pages 29–45, 2002.

[And94] L. O. Andersen. *Program analysis and specialization for the C programming language*. PhD thesis, University of Copenhagen, 1994.

[ASU86] A.V. Aho, R. Sethi, and J.D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, 1986.

[Bac98] David Francis Bacon. *Fast and Effective Optimization of Statically Typed Object-Oriented Languages*. PhD thesis, University of California at Berkeley, 5, 1998.

[BD03] Jason Brittain and Ian F. Darwin. *Tomcat: The Definitive Guide*. O'Reilly and Associates, 2003.

[BN99] Mathias Braux and Jacques Noyé. Towards partially evaluating reflection in Java. In *Proceedings of the ACM Workshop on Partial Evaluation and Semantics-based Program Manipulation*, pages 2–11, 1999.

[BS96] David F. Bacon and Peter F. Sweeney. Fast static analysis of C++ virtual function calls. In *Proceedings of the 11th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 324–341, 1996.

[CG94] Brad Calder and Dirk Grunwald. Reducing indirect function call overhead in C++ programs. In *Conference Record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 397–408, Portland, Oregon, 1994.

[CMS03] Aske Simon Christensen, Anders Møller, and Michael I. Schwartzbach. Precise analysis of string expressions. In *Proc. 10th International Static Analysis Symposium, SAS '03*, volume 2694 of *LNCS*, pages 1–18. Springer-Verlag, June 2003. Available from `http://www.brics.dk/JSA/`.

[DGC95] Jeffrey Dean, David Grove, and Craig Chambers. Optimization of object-oriented programs using static class hierarchy analysis. *Lecture Notes in Computer Science*, 952:77–101, 1995.

[EGH94] Maryam Emami, Rakesh Ghiya, and Laurie J. Hendren. Context-sensitive interprocedural points-to analysis in the presence of function pointers. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 242–256, 1994.

[FF04] Ira R. Forman and Nate Forman. *Java Reflection in Action*. Manning Publications, 2004.

[GC01] David Grove and Craig Chambers. A framework for call graph construction algorithms. *ACM Trans. Program. Lang. Syst.*, 23(6):685–746, 2001.

[GCSD02] Yann-Gaël Guéhéneuc, Pierre Cointe, and Marc Ségura-Devillechaise. Java reflection exercises, correction, and FAQs. `http://www.yann-gael.gueheneuc.net/Work/Teaching/Documents/Practical-ReflectionCourse.doc.pdf`, 2002.

[GDDC97] David Grove, Greg DeFouw, Jeffrey Dean, and Craig Chambers. Call graph construction in object-oriented languages. In *Proceedings of the ACM Conference on Object-oriented Programming, Systems, Languages, and Applications*, pages 108–124, 1997.

[GOP03] Thomas Gschwind, Johann Oberleitner, and Martin Pinzger. Using run-time data for program comprehension. In *IWPC '03: Proceedings of the 11th IEEE International Workshop on Program Comprehension*, page 245, Washington, DC, USA, 2003. IEEE Computer Society.

[HDH04] Martin Hirzel, Amer Diwan, and Michael Hind. Pointer analysis in the presence of dynamic class loading. In *Proceedings of the*

European Conference on Object-Oriented Programming, Systems, Languages, and Applications, pages 96–122, 2004.

[KPK02] L. Koved, M. Pistoia, and A. Kershenbaum. Access rights analysis for Java. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 359 – 372, 2002.

[LH03] Ondřej Lhoták and Laurie Hendren. Scaling Java points-to analysis using Spark. In G. Hedin, editor, *Compiler Construction, 12th International Conference*, volume 2622 of *LNCS*, pages 153–169, Warsaw, Poland, April 2003. Springer.

[LL05a] V. Benjamin Livshits and Monica S. Lam. Finding security errors in Java programs with static analysis. Technical report, Stanford University, August 2005.

[LL05b] V. Benjamin Livshits and Monica S. Lam. Finding security errors in Java programs with static analysis. In *Proceedings of the 14th Usenix Security Symposium*, pages 271 – 286, August 2005.

[LWL+05] Monica S. Lam, John Whaley, V. Benjamin Livshits, Michael C. Martin, Dzintars Avots, Michael Carbin, and Christopher Unkel. Context-sensitive program analysis as database queries. In *Proceedings of the ACM Symposium on Principles of Database Systems*, pages 1 – 12, June 2005.

[MRR01] A. Milanova, A. Rountev, and B. Ryder. Precise call graph construction in the presence of function pointers. Technical report, Rutgers University, 2001.

[MRR02] Ana Milanova, Atanas Rountev, and Barbara G. Ryder. Parameterized object sensitivity for points-to and side-effect analyses for Java. In *ISSTA '02: Proceedings of the 2002 ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 1–11, 2002.

[MRR04] A. Milanova, A. Rountev, and B. G. Ryder. Precise and efficient call graph onstruction for programs with function pointers. *Journal of Automated Software Engineering*, 2004.

[MY98] Hidehiko Masuhara and Akinori Yonezawa. Design and partial evaluation of meta-objects for a concurrent reflective language. In *Proceedings of the European Conference on Object-Oriented Programming*, pages 418–439. Springer-Verlag, 1998.

[Nie] Patrick Niemeyer. BeanShell 2.0. http://www.beanshell.org/BeanShellSlides.pdf.

[PR96] Hemant D. Pande and Barbara G. Ryder. Data-flow-based virtual function resolution. In *SAS '96: Proceedings of the Third International Symposium on Static Analysis*, pages 238–254. Springer-Verlag, 1996.

[PS01] Igor Pechtchanski and Vivek Sarkar. Dynamic optimistic interprocedural analysis: a framework and an application. In *Proceedings of the 16th ACM SIGPLAN conference on Object oriented programming, systems, languages, and applications*, pages 195–210, 2001.

[QH04] Feng Qian and Laurie Hendren. Towards dynamic interprocedural analysis in JVMs. In *Usenix VM*, 2004.

[RRHK00] Derek Rayside, Steve Reuss, Erik Hedges, and Kostas Kontogiannis. The effect of call graph construction algorithms for object-oriented programs on automatic clustering. In *Proceedings of the 8th International Workshop on Program Comprehension*, page 191. IEEE Computer Society, 2000.

[RSS+04] Darrell Reimer, Edith Schonberg, Kavitha Srinivas, Harini Srinivasan, Bowen Alpern, Robert D. Johnson, Aaron Kershenbaum, and Larry Koved. SABER: Smart Analysis Based Error Reduction. In *Proceedings of International Symposium on Software Testing and Analysis*, pages 243 – 251, 2004.

[Ruf93] Erik Ruf. Partial evaluation in reflective system implementations. In *Workshop on Reflection and Metalevel Architecture*, October 1993.

[SHR+00] Vijay Sundaresan, Laurie Hendren, Chrislain Razafimahefa, Raja Vallée-Rai, Patrick Lam, Etienne Gagnon, and Charles Godin.

Practical virtual method call resolution for Java. *ACM SIGPLAN Notices*, 35(10):264–280, 2000.

[Thi96] Peter Thiemann. Towards partial evaluation of full Scheme. In *Reflection '96*, 1996.

[TLSS99] Frank Tip, Chris Laffra, Peter F. Sweeney, and David Streeter. Practical experience with an application extractor for Java. *ACM SIGPLAN Notices*, 34(10):292–305, 1999.

[TP00] Frank Tip and Jens Palsberg. Scalable propagation-based call graph construction algorithms. *ACM SIGPLAN Notices*, 35(10):281–293, 2000.

[WL04] John Whaley and Monica Lam. Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*, pages 131 – 144, 2004.

[WN04] Wesley Weimer and George Necula. Finding and preventing run-time error handling mistakes. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 419 – 431, 2004.

[Zha98] Sean Zhang. *Practical Pointer Aliasing Analyses for C.* PhD thesis, Rutgers University, August 1998.