Hacking SandMark

Christian Collberg

May 31, 2002

Contents

I Working with SandMark **9**

1 Introducing SandMark **11**
1.1 Introduction . . . 11
Trying out the Watermarker . . . 11
1.2 Installing SandMark . . . 13
Building SandMark . . . 14
1.3 Scripting SandMark . . . 15

2 The SandMark Code-base **17**
2.1 SandMark Packages . . . 17
sandmark . . . 17
sandmark.gui . . . 17
sandmark.html . . . 18
sandmark.obfuscate . . . 18
sandmark.obfuscate.setfieldspublic . . . 18
sandmark.obfuscate.splitclass . . . 19
sandmark.obfuscate.util . . . 19
sandmark.optimise . . . 19
sandmark.util . . . 19
sandmark.util.exec . . . 20
sandmark.util.graph . . . 21
sandmark.util.graph.codec . . . 21
sandmark.util.javagen . . . 21
sandmark.watermark . . . 22
sandmark.watermark.CT . . . 23
sandmark.watermark.CT.embed . . . 23
sandmark.watermark.CT.encode . . . 23
sandmark.watermark.CT.encode.ir . . . 23
sandmark.watermark.CT.encode.ir2ir . . . 24
sandmark.watermark.CT.encode.storage . . . 25
sandmark.watermark.CT.recognize . . . 25
sandmark.watermark.CT.trace . . . 25
sandmark.watermark.constantstring . . . 25
sandmark.statistics . . . 25
2.2 Dynamic Class Loading . . . 26

3 Extending SandMark **27**
3.1 Adding an Obfuscator . . . 27
3.2 Obfuscation Configuration . . . 28
Method Summary . . . 28
3.3 Adding a Watermarker . . . 28
3.4 Adding a Graph Codec . . . 31

3.5 Documentation . . . 32

II SandMark Algorithms 35

4 The Collberg-Thomborson Watermarking Algorithm 37
4.1 Introduction . . . 37
4.2 Annotation . . . 37
4.3 Tracing . . . 40
An Example . . . 41
Choosing `mark()`-Locations . . . 41
4.4 Embedding . . . 42
Building the Graph . . . 42
Splitting the Graph . . . 43
Generating Intermediate Code . . . 44
Generating Java Code . . . 46
Inserting the Java Code . . . 49
4.5 Recognition . . . 49
4.6 Passing Roots in Formal Parameters . . . 50

5 The ConstantString Static Watermarking Algorithm 53
5.1 Introduction . . . 53
5.2 Embedding . . . 53
5.3 Recognition . . . 53

6 The RenameMethodsFields Watermarking Algorithm 55
6.1 Introduction . . . 55
6.2 Embedding . . . 55
6.3 Recognition . . . 55

7 Project VIII: Class-Splitting Obfuscations 57
7.1 Mission Document . . . 57
7.2 Basic Idea . . . 57
7.3 How we plan to implement . . . 57
7.4 Project Tasks . . . 57
7.5 Reference . . . 57

8 The Bogus Switch Watermarking Algorithm 59
8.1 Introduction . . . 59
8.2 Embedding . . . 59
8.3 Recognition . . . 59

9 Boolean Variable Splitting 61
9.1 Introduction . . . 61
9.2 Boolean variable recognition . . . 61
Observations . . . 61
Code pattern matching method . . . 62
Variable range based method . . . 62
9.3 Splitting Techniques . . . 64
XOR Splitting . . . 64
Parity Splitting . . . 64
Equality Splitting . . . 64
9.4 Code Transformation Techniques . . . 65
Simple Variable Transformations . . . 65

Boolean Array Splitting Techniques . . . 65

10 Bogus Initializers Watermarking Algorithm 67
10.1 Introduction . . . 67
10.2 Embedding . . . 67
10.3 Recognition . . . 68

11 The AddMethField Watermarking Algorithm 69
11.1 Introduction . . . 69
11.2 Embedding . . . 69
11.3 Recognition . . . 70

12 The HatTrick Watermarking Algorithm 73
12.1 Introduction . . . 73
12.2 Embedding . . . 73
12.3 Recognition . . . 74

13 The AddBogusFields Obfuscation Algorithm 75
13.1 Introduction . . . 75
13.2 Apply . . . 75

14 The PromoteLocals Obfuscating Algorithm 77
14.1 Introduction . . . 77
14.2 Apply . . . 77

15 Bogus Predicates Obfuscating Algorithm 79
15.1 Introduction . . . 79
15.2 Apply . . . 79

16 The Method 2R Madness Obfuscation Algorithm 81
16.1 Introduction . . . 81
16.2 Pass Two - Method Splitting . . . 81
16.3 Pass Three - Primitive Promoting . . . 82
16.4 Pass Four - Parameter Reordering . . . 83
16.5 Pass Five - Signature Bludgeoning . . . 83
16.6 Pass Six - Method Merging . . . 84

17 Class-Splitting and False Refactoring Obfuscation Algorithms 87
17.1 Introduction . . . 87
17.2 Class Splitting Obfuscation . . . 87
Dependency Graph Generation . . . 87
Topological sort of the Dependency Graph . . . 87
Results . . . 89
17.3 False Re-factoring of Classes . . . 89
17.4 Reference . . . 90

18 The Degrading Obfuscation Algorithm 91
18.1 Introduction . . . 91
18.2 Promotion Obfuscation . . . 91
Special `int`'s . . . 91
Promoting Primitives . . . 92
Controling Promotion . . . 93
18.3 Thread Contention Obfuscation . . . 94
Apply . . . 94
18.4 Sample Results . . . 99

19 The Robust Object Watermarking Algorithm **101**
19.1 Introduction . . . 101
19.2 Vector Instruction Group Identification . . . 101
19.3 Building the CodeBook : . . . 103
19.4 Vector Extraction . . . 103
19.5 The overall watermarking algorithm . . . 103
19.6 Insertion . . . 104
19.7 Comparison with the previous implementation on x86 assembly code . . . 106
19.8 Analysis of attacks . . . 106
19.9 Brief analysis of our work and scope for future Work . . . 106
19.10Conclusion: . . . 109
19.11Bibliography . . . 109

20 Potkonjak's Watermarking Algorithm **111**
20.1 Introduction . . . 111
20.2 Embedding . . . 111
Convert watermark to binary . . . 111
Generate the original coloring . . . 111
Bridge construct and generate new coloring . . . 111
Determine if the method can be watermarked . . . 112
Generate the pool of graphs . . . 112
Pick the correct coloring . . . 112
Embed the chosen coloring . . . 112
20.3 Recognition . . . 112
Determine the current coloring . . . 112
Generate the original coloring for the method . . . 112
Compare the two colorings . . . 113
Convert binary to string . . . 113
20.4 Future Work . . . 113

III Appendices **115**

A Useful Tools **117**
A.1 Examining Java Classfiles . . . 117
javap . . . 117
Jasmin . . . 117
BCEL's Class Construction Kit . . . 118
BCEL's listclass . . . 118
BCEL's JustIce . . . 119
A.2 BCEL . . . 119
Getting started . . . 119
Editing a Method . . . 120
Editing Instructions . . . 120
Local variables . . . 121
Example 1: Creating a New Class . . . 122
Branches . . . 123
Exceptions . . . 123
Types . . . 124
Method Calls . . . 125
Example 2: Modifying an Existing Program . . . 126
Control Flow Graphs . . . 126
Miscellaneous . . . 126
A.3 BLOAT . . . 127

Introduction . . . 127
The Basics . . . 127
Editing Class Files . . . 128
Committing Edits . . . 131
Instruction Representation . . . 131
Advanced Features . . . 131

B Working with the SandMark Code-base **137**
B.1 Coding standard . . . 137
B.2 Using CVS . . . 138
Installing CVS . . . 138
Getting Started . . . 138
Checking out code . . . 138
Updating . . . 140
Deleting files . . . 140
Summary . . . 140

Part I

Working with SandMark

Chapter 1

Introducing SandMark

1.1 Introduction

SandMark is designed to be the Swiss-Army-Chainsaw of software protection research. In other words, we hope to build an infrastructure that makes it easy to implement algorithms for

1. code obfuscation,
2. softeware watermarking, and
3. tamper-proofing.

In fact, we hope to implement *every* software protection algorithm know to man, so that we can compare and evaluate them.

You normally interact with sandmark through its graphical interface. Start it by typing `smark` in the `smbin` directory:

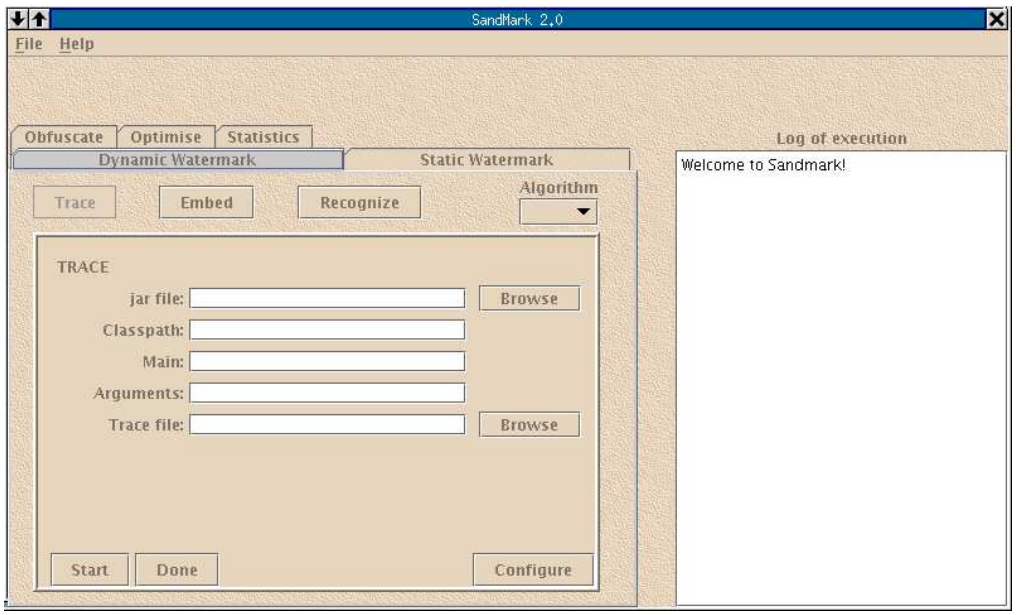

Trying out the Watermarker

To get started try watermarking the `TTT` (tic-tac-toe) application. It can be found in the `smapps2` directory. Start SandMark, then do the following:

1. In the *trace* pane enter

   Jar-file to watermark: `TTT.jar`

   Main class name: `TTTApplication`

2. Hit START.

3. Click on a few X's and O's. Remember the order in which you do the clicks! It should look something like this:

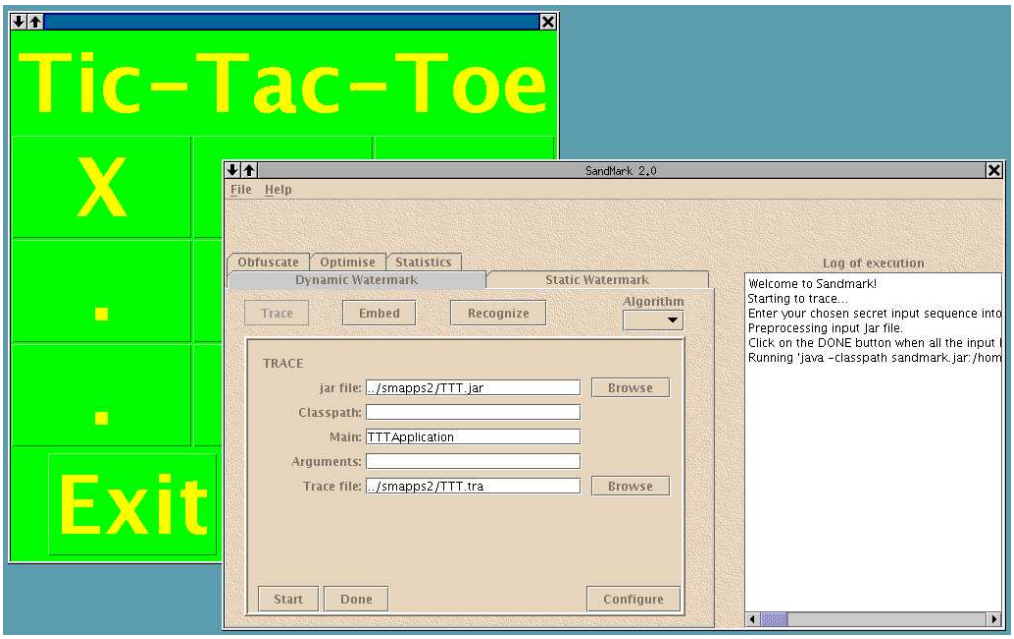

4. Hit DONE.

5. Go to the *embed* pane and enter the watermark value `123456`, like this:

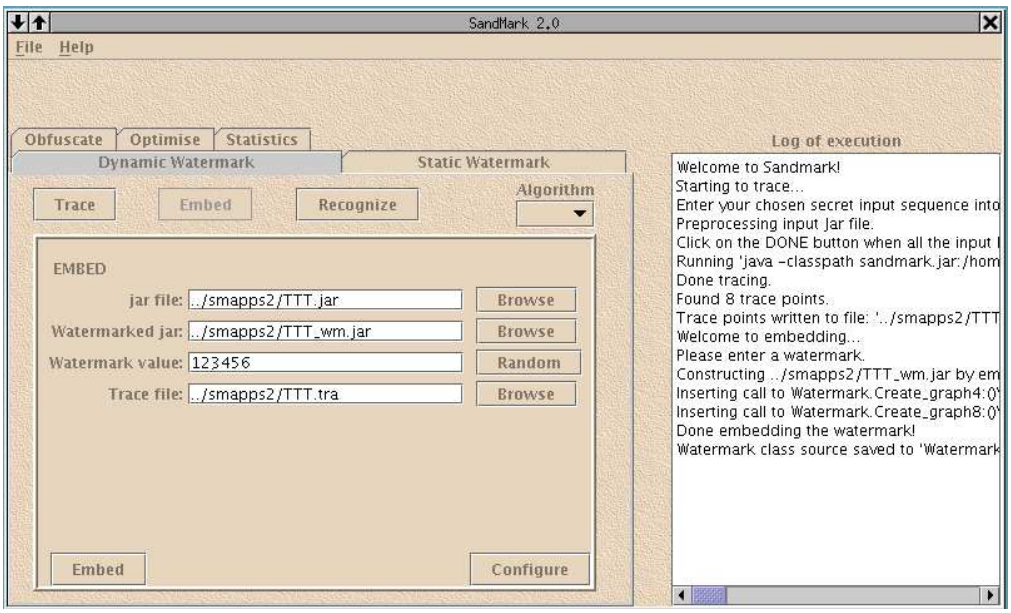

6. Hit EMBED.
7. Go to the *recognize* pane and hit the START button.
8. Click on the same X's and O's as you did in step 3), in the same order.
9. Hit DONE:

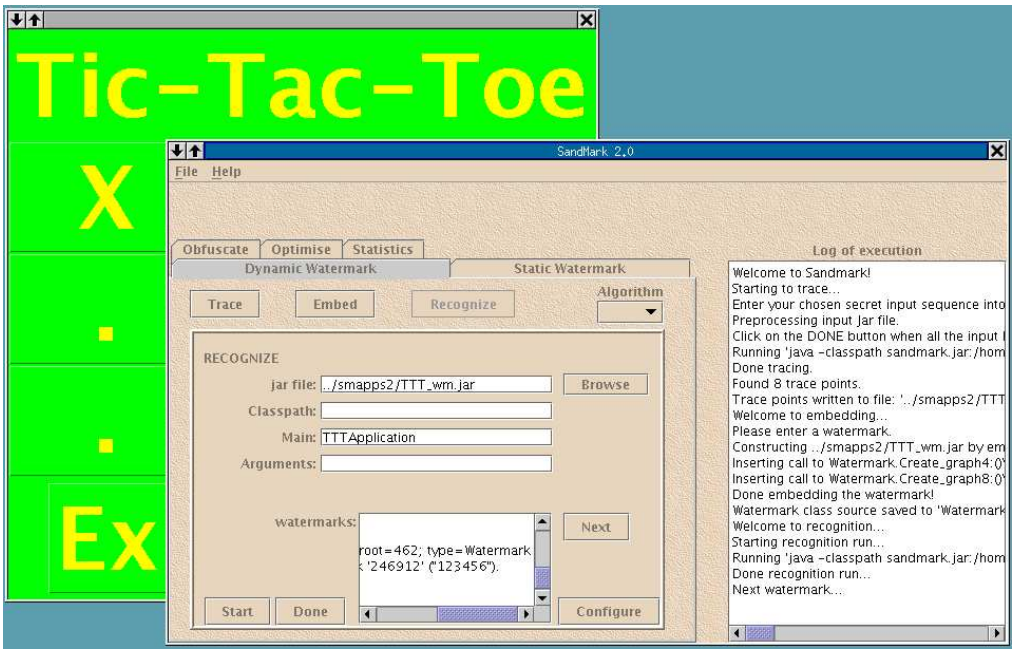

You should see the watermark `123456` extracted from the watermarked `TTT` application.

You can also try

```
> smark -f ../smapps2/TTT.script
```

This will run a script that traces, embeds, and recognizes a watermark in the TTT application. You still have to enter the X's and O's and hit the DONE buttons in the trace and recognize panes.

1.2 Installing SandMark

SandMark's source code is stored in a CVS repository at `cvs.cs.arizona.edu`. You can get the source anonymously (in which case you can't make any changes to it): To get the sources anonymously, do the following

```
> setenv CVS_RSH ssh
> cvs -d :pserver:anonymous@cvs.cs.arizona.edu:/cvs/wmark login
> cvs -d :pserver:anonymous@cvs.cs.arizona.edu:/cvs/wmark \
   checkout -P smark2 smextern smapps2 smbloat2 smbin
```

or, if you have got write access to SandMark (i.e an account on `cvs.cs.arizona.edu`) you should instead do

```
> setenv CVS_RSH ssh
> cvs -d :ext:MyLogin@cvs.cs.arizona.edu:/cvs/cvs/wmark \
       checkout -P smark2 smextern smapps2 smbloat smbin
```

where `MyLogin` is your account name on `cvs.cs.arizona.edu`.

You should now have four directories:

smark2: The SandMark sources.

smbin: Scripts.

smextern: External Java code i.e. jar and zip files needed to run SandMark.

smapps2: Some simple applications you can use to try out SandMark

smbloat2: Some simple test cases for Bloat. Read these to get a feel for how Bloat should be used.

Once you have checked out SandMark you can get the latest version by running the `cvs update` command:

```
> cd smark2
> cvs update -dP
```

To add a new file you do `cvs add file`, to remove it `cvs rm file` and to commit your changes to the repository at `cvs.cs.arizona.edu` you say `cvs commit`.

Building SandMark

Note that you will need Java 1.4 to run SandMark properly. Get the latest version from `http://java.sun.com/j2se/1.4`.

Do the following to build SandMark:

```
> cp smark2/Makedefs.std smark2/Makedefs
# Make the 'obvious' changes to Makedefs.
# In particular, you should set these variables:
#   JDK =
#   HOME =

> cp smbin/smark.std smbin/smark
# Make the 'obvious' changes to smark.
# Again, you should set these variables:
#   JDK =
#   HOME =

> chmod a+rx smbin/smark
> make

# Build applications to watermark
> cd smapps2
> make

# Start SandMark
> smbin/smark
```

SandMark gets compiled into a jar-file `sandmark.jar`. To execute it you also need some other packages (`bloat`, `BCEL`, etc.), which can be found in the `smextern` directory. The `smark`-script takes care of setting Java's classpath correctly so that these get picked up.

You can also build the manual:

```
> cd smark2/doc
> make
> make again
> make again
```

Finally, you can generate html or latex from the JavaDoc comments in the source:

```
> make jdoc
> make jtex
```

This generates a PostScript file `classes.ps` and a directory `jdoc` of html files.

1.3 Scripting SandMark

SandMark can be scripted. Either start SandMark from the command line with the `-f` option:

```
> smark -f file.script
```

or enter the script file in SandMark's `file` pull-down manu.

- You can set a property value using the command

  ```
  set PROPERTY VALUE
  ```

  The following property values are recognized:

  DWM_CT_AnnotatorClass:

  DWM_CT_Encode_NodeType:

  DWM_CT_Encode_ParentClass:

  DWM_CT_Encode_ClassName:

  DWM_CT_Encode_AvailableEdges:

  DWM_CT_Encode_StoreWhat:

  DWM_CT_Encode_StoreMethods:

  DWM_CT_Encode_StoreLocation :

  DWM_CT_Encode_ProtectionMethods:

  DWM_CT_Encode_IndividualFixups:

  DWM_CT_Encode_Encoding: Should be one of `perm` or `radix`. `"*"` picks a random encoding method.

  DWM_CT_Encode_Components:

  DWM_CT_Encode_Package:

  DWM_MaxTracePoints:

  DWM_CT_Encode_StoreLocation: One of `formal` or `global`.

  DWM_CT_DumpIR:

  DWM_ClassPath:

- To run tracing use the command

  ```
  trace input.jar trace.tra MAINCLASS ARGUMENTS
  ```

  The classpath is set through the command

  ```
  set ClassPath ...
  ```

  Tracepoints are written to the `trace.tra` file.

- Embed watermark in input.jar using the command

```
embed input.jar output.jar watermark trace.tra
```

Read the tracepoints from the file `trace.tra`.

- Obfuscate input.jar, creating output.jar:

```
obfuscate input.jar output.jar
```

- Run recognition.

```
recognize input.jar watermark_count MAINCLASS ARGUMENTS
```

The classpath is set through the command

```
set ClassPath ...
```

- If the first non-blank character on a line is # the rest of the line is ignored.
- Commands are case insensitive, arguments are case sensitive.

Chapter 2

The SandMark Code-base

2.1 SandMark Packages

SandMark is a large application. Currently it is made up of some 29000 lines of java code distributed over 200 classes. The code is organized in a fairly deep package hierarchy.

sandmark

Classes at the top level:

sandmark.Console: The main entry point to SandMark. It contains code that starts up the GUI and acts as an interface between the GUI and the rest of the application.

sandmark.CLI: The command line interpreter.

sandmark.Scripting: The script interpreter.

sandmark.gui

`sandmark.gui` contains the code that build up the graphical interface.

sandmark.gui.AboutDialog:

sandmark.gui.ButtonRenderer:

sandmark.gui.CodeDialog:

sandmark.gui.EmbedConfigDialog:

sandmark.gui.ExtensionFileFilter:

sandmark.gui.GUIListener:

sandmark.gui.HelpFrame:

sandmark.gui.IntegerInput:

sandmark.gui.LayoutConstraints:

sandmark.gui.LightRenderer:

sandmark.gui.MultiHeaderRenderer:

sandmark.gui.ObTableModel:

sandmark.gui.ObfuscateConfigDialog:

sandmark.gui.RelativeLayout:

sandmark.gui.SandMarkFrame:

sandmark.gui.SandMarkGUIConstants:

sandmark.gui.SkinPanel:

sandmark.gui.StatDialog:

sandmark.gui.StatTableModel:

sandmark.gui.TablePanel:

sandmark.gui.SMarkGUIConstants:

sandmark.gui.SandMarkOSConstants:

sandmark.html

`Html` help files are kept here.

sandmark.obfuscate

Each obfuscation lives in its own subdirectory in `sandmark.obfuscate`. `sandmark.obfuscate.setfieldspublic`, for example, is an obfuscator that sets all fields to `public`.

sandmark.obfuscate.AllClassesObfuscator:

sandmark.obfuscate.AppObfuscator: The abstract base-class from which any obfuscator that works on an entire application should inherit. It inherits from `sandmark.obfuscate.GeneralObfuscator`.

sandmark.obfuscate.ClassObfuscator: The abstract base-class from which any obfuscator that works on a single class should inherit. It inherits from `sandmark.obfuscate.GeneralObfuscator`.

sandmark.obfuscate.GeneralObfuscator: The top-level abstract class for an obfuscator.

sandmark.obfuscate.MethodObfuscator: The abstract base-class from which any obfuscator that works on a single method should inherit. It inherits from `sandmark.obfuscate.GeneralObfuscator`.

sandmark.obfuscate.NameOverloadObfuscator: An advanced name obfuscator based on Paul Tyma's patent.

sandmark.obfuscate.Obfuscate: This class handles the incoming obfuscation requests from the GUI. It is called by `sandmark.Console`.

sandmark.obfuscate.Obfuscator: The top-level obfuscator that decides on overall strategy – which algorithms should be applied to which pieces of code.

sandmark.obfuscate.setfieldspublic

An obfuscator that sets all fields and methods to public.

sandmark.obfuscate.setfieldspublic.SetFieldsPublic:

sandmark.obfuscate.splitclass

An obfuscator that splits a class in two pieces. Non-functional.

sandmark.obfuscate.splitclass.DummyClass:

sandmark.obfuscate.splitclass.Obf:

sandmark.obfuscate.splitclass.SplitClass:

sandmark.obfuscate.util

`sandmark.obfuscate.util` contains classes of interest to all obfuscators.

sandmark.optimise

Routines to optimize Java programs. Currently based solely on the BLOAT package.

sandmark.optimise.Optimise:

sandmark.optimise.Main:

sandmark.optimise.Optimiser:

sandmark.util

`sandmark.util` contains classes of interest to all SandMark tools.

sandmark.util.Compile:

sandmark.util.BCEL: Utility routines to work with the `BCEL` (formerly `JavaClass`) package.

sandmark.util.ByteCodeLocation: Represents a location within a Java application. Essentially a tuple

$$\langle \texttt{sandmark.util.MethodID}, \texttt{lineNumber}, \texttt{byteCodeIndex} \rangle.$$

sandmark.util.CallGraphNode: The node in a call graph.

sandmark.util.CircularBuffer: A bounded buffer of arbitrary objects.

sandmark.util.ClassFileCollection: Routines for loading the classes of the program to be watermarked, obfuscated, etc. A class can be parsed into a `BCEL` or a `BLOAT` object.

sandmark.util.ClassHierarchy: A package for building the complete class hierarchy from a jarfile. There are many routines for querying the hierarchy, such as which methods override a particular method, which classes extend a particular class, etc.

sandmark.util.DependencyGraph: Used by `sandmark.obfuscate.splitclass.SplitClass`.

sandmark.util.Editor: Utility class for use with `BLOAT`.

sandmark.util.EditedClass: A class for maintaining a cache of classes and methods edited with `BCEL`.

sandmark.util.FileClassLoader: Used to load classes on the fly.

sandmark.util.FindClassFiles: Used to load classes on the fly.

sandmark.util.GraphViewer: Graphical display of `sandmark.util.graph.Graph`.

sandmark.util.InstructionTree:

sandmark.util.LabeledGrid:

sandmark.util.Log: Routines for printing out messages (both to a file and to the screen) to the user.

sandmark.util.MethodID: Represents a method within a Java application. Essentially a tuple

$$\langle \texttt{name}, \texttt{signature}, \texttt{className} \rangle.$$

sandmark.util.MethodNode:

sandmark.util.Misc:

sandmark.util.Options: Routines for cracking command line parameters.

sandmark.util.PriorityQueue:

sandmark.util.SparseVector: A specialized version on `java.util.Vector`.

sandmark.util.StackFrame: Used to build call graphs.

sandmark.util.StatisticsRecord:

sandmark.util.Stats:

sandmark.util.StringInt: A routine to package up an int inside a string.

sandmark.util.TempDir: Routines to manage temporary directories.

sandmark.util.Time: Routines for timing Java programs.

sandmark.util.exec

`sandmark.util.exec` contains classes that allow SandMark to run another Java application under debugging. This uses Java's JDI package. JDI allows you to start up a Java program, set breakpoints, trace method calls, examine variables, etc. We use this in the implementation of the trace and recognition steps of the dynamic watermarking algorithms.

sandmark.util.exec.Breakpoint: A class for creating breakpoints.

sandmark.util.exec.DumpAll:

sandmark.util.exec.EventHandler:

sandmark.util.exec.EventThread:

sandmark.util.exec.Heap: A class for iterating through all the objects on the heap.

sandmark.util.exec.HeapData: Objects returned by `sandmark.util.exec.Heap`.

sandmark.util.exec.MethodCallData:

sandmark.util.exec.Output: A class that handles input and output from the application being executed.

sandmark.util.exec.Overseer: Extend this class to run a program under JDI.

sandmark.util.exec.TracingException:

sandmark.util.graph

Many watermarking algorithms make use of graphs. This package contains a multigraph package `sandmark.util.graph.Graph` as well as operations on such graphs, such as all-pairs-shortest-path, depth-first-search, etc.

sandmark.util.graph.GraphOp: Additional algorithms on graphs.

sandmark.util.graph.Dfs: Classify the edges of a graph as `tree`, `cross`, and `back`.

sandmark.util.graph.Edge: Edge objects of a `sandmark.util.graph.Graph`.

sandmark.util.graph.Graph: A class for constructing multi-graphs.

sandmark.util.graph.Matrix: A class where graphs are represented as an adjacency matrix.

sandmark.util.graph.Node: Node objects of a `sandmark.util.graph.Graph`.

sandmark.util.graph.Path: A class for representing a path in a `sandmark.util.graph.Graph`.

sandmark.util.graph.codec

This package implements various algorithms for embedding a number into the topology of a graph.

sandmark.util.graph.codec.DecodeFailure:

sandmark.util.graph.codec.GraphCodec: Base-class for graph codecs.

sandmark.util.graph.codec.PPCT: Non-functional.

sandmark.util.graph.codec.PermutationGraph: A permutation graph encoding.

sandmark.util.graph.codec.RadixGraph: A radix graph encoding.

sandmark.util.javagen

`sandmark.util.javagen` contains classes for building up Java methods and classes. It is essentially a layer on top of `BCEL`. `sandmark.util.javagen` allows you to build up a Java AST and then turn that into bytecode.

sandmark.util.javagen.AssignField:

sandmark.util.javagen.AssignIndex:

sandmark.util.javagen.AssignStatic:

sandmark.util.javagen.Block:

sandmark.util.javagen.Cast:

sandmark.util.javagen.Class:

sandmark.util.javagen.Comment:

sandmark.util.javagen.CondNotNullExpr:

sandmark.util.javagen.Discard:

sandmark.util.javagen.EmptyStatement:

sandmark.util.javagen.Expression:

sandmark.util.javagen.Field:

sandmark.util.javagen.FieldRef:

sandmark.util.javagen.Formal:

sandmark.util.javagen.IfNotNull:

sandmark.util.javagen.Java:

sandmark.util.javagen.List:

sandmark.util.javagen.LiteralInt:

sandmark.util.javagen.LiteralString:

sandmark.util.javagen.LoadIndex:

sandmark.util.javagen.Local:

sandmark.util.javagen.LocalRef:

sandmark.util.javagen.Method:

sandmark.util.javagen.MyClass:

sandmark.util.javagen.New:

sandmark.util.javagen.NewArray:

sandmark.util.javagen.Null:

sandmark.util.javagen.Return:

sandmark.util.javagen.Statement:

sandmark.util.javagen.StaticCall:

sandmark.util.javagen.StaticFunCall:

sandmark.util.javagen.StaticRef:

sandmark.util.javagen.Test:

sandmark.util.javagen.Try:

sandmark.util.javagen.VirtualCall:

sandmark.util.javagen.VirtualFunCall:

sandmark.watermark

This package contain all the watermarking algorithms. Each algorithm resides in its own directory.

sandmark.watermark.DynamicEmbed: Handles the embedding phase of a dynamic watermarker.

sandmark.watermark.DynamicRecognize: Handles the recognition phase of a dynamic watermarker.

sandmark.watermark.DynamicTrace: Handles the tracing phase of a dynamic watermarker.

sandmark.watermark.DynamicWatermarker: The base-class for all dynamic watermarkers.

sandmark.watermark.GeneralWatermarker: The base-class of all watermarkers

sandmark.watermark.StaticWatermarker: The base-class for all static watermarkers.

sandmark.watermark.StaticEmbed: Handles the embedding phase of a static watermarker.

sandmark.watermark.StaticRecognize: Handles the recognition phase of a static watermarker.

sandmark.watermark.WatermarkingException:

sandmark.watermark.Watermarking:

sandmark.watermark.CT

This package implements the Collberg-Thomborson dynamic watermarking algorithm.

sandmark.watermark.CT.CT: The main class for the Collberg-Thomborson algorithm. Extends `sandmark.watermark.DynamicWatermarker`.

sandmark.watermark.CT.embed

This package implements the embedding of a graph into a Java program. This is part of the Collberg-Thomborson watermarking algorithm.

sandmark.watermark.CT.embed.AddParameters:

sandmark.watermark.CT.embed.CallGraphPath:

sandmark.watermark.CT.embed.DeleteMarkCalls:

sandmark.watermark.CT.embed.EmbedData:

sandmark.watermark.CT.embed.Embedder:

sandmark.watermark.CT.embed.InsertStorageCreators:

sandmark.watermark.CT.embed.PrepareTrace:

sandmark.watermark.CT.embed.ReplaceMarkCalls:

sandmark.watermark.CT.encode

This package implements the encoding of a graph into a Java program. This is part of the Collberg-Thomborson watermarking algorithm.

sandmark.watermark.CT.encode.Encoder:

sandmark.watermark.CT.encode.Graph2IR:

sandmark.watermark.CT.encode.Split:

sandmark.watermark.CT.encode.ir

This package implements the encoding of a graph into a simple intermediate code. This is part of the Collberg-Thomborson watermarking algorithm.

sandmark.watermark.CT.encode.ir.AddEdge:

sandmark.watermark.CT.encode.ir.Build:

sandmark.watermark.CT.encode.ir.Construct:

sandmark.watermark.CT.encode.ir.Create:

sandmark.watermark.CT.encode.ir.CreateNode:

sandmark.watermark.CT.encode.ir.CreateStorage:

sandmark.watermark.CT.encode.ir.Debug:

sandmark.watermark.CT.encode.ir.Destroy:

sandmark.watermark.CT.encode.ir.Destruct:

sandmark.watermark.CT.encode.ir.Field:

sandmark.watermark.CT.encode.ir.Fixup:

sandmark.watermark.CT.encode.ir.FollowLink:

sandmark.watermark.CT.encode.ir.Formal:

sandmark.watermark.CT.encode.ir.IR:

sandmark.watermark.CT.encode.ir.Init:

sandmark.watermark.CT.encode.ir.List:

sandmark.watermark.CT.encode.ir.LoadNode:

sandmark.watermark.CT.encode.ir.Method:

sandmark.watermark.CT.encode.ir.NodeStorage:

sandmark.watermark.CT.encode.ir.PrintGraph:

sandmark.watermark.CT.encode.ir.ProtectRegion:

sandmark.watermark.CT.encode.ir.SaveNode:

sandmark.watermark.CT.encode.ir.StaticCall:

sandmark.watermark.CT.encode.ir2ir

This package implements transformations of the intermediate code used in the encoding of graphs in the Collberg-Thomborson watermarking algorithm.

sandmark.watermark.CT.encode.ir2ir.AddFields:

sandmark.watermark.CT.encode.ir2ir.AddFormals:

sandmark.watermark.CT.encode.ir2ir.Builder:

sandmark.watermark.CT.encode.ir2ir.CleanUp:

sandmark.watermark.CT.encode.ir2ir.Debug:

sandmark.watermark.CT.encode.ir2ir.Destructors:

sandmark.watermark.CT.encode.ir2ir.InlineFixups:

sandmark.watermark.CT.encode.ir2ir.Protect:

sandmark.watermark.CT.encode.ir2ir.SaveNodes:

sandmark.watermark.CT.encode.ir2ir.Transformer:

sandmark.watermark.CT.encode.storage

This package implements operations on the storage of graph nodes in Collberg-Thomborson watermarking algorithm.

sandmark.watermark.CT.encode.storage.NodeStorage:

sandmark.watermark.CT.encode.storage.Array:

sandmark.watermark.CT.encode.storage.Hash:

sandmark.watermark.CT.encode.storage.GlobalStorage:

sandmark.watermark.CT.encode.storage.Pointer:

sandmark.watermark.CT.encode.storage.StorageClass:

sandmark.watermark.CT.encode.storage.Vector:

sandmark.watermark.CT.recognize

This package implements the recognition part of Collberg-Thomborson's software watermarking algorithm. It should be generalized to handle recognition in *any* dynamic software watermarking algorithm.

sandmark.watermark.CT.recognize.Heap2Graph:

sandmark.watermark.CT.recognize.RecognizeData:

sandmark.watermark.CT.recognize.Recognizer:

sandmark.watermark.CT.trace

This package currently implements the trace part of Collberg-Thomborson's software watermarking algorithm. It should be generalized to handle tracing in *any* dynamic software watermarking algorithm.

sandmark.watermark.CT.trace.Annotator:

sandmark.watermark.CT.trace.CallForest:

sandmark.watermark.CT.trace.Preprocessor:

sandmark.watermark.CT.trace.TracePoint:

sandmark.watermark.CT.trace.Tracer:

sandmark.watermark.constantstring

This package implements a trivial static watermarking algorithm.

sandmark.watermark.constantstring.ConstantString:

sandmark.watermark.constantstring.Test:

sandmark.watermark.constantstring.Program:

sandmark.statistics

This package counts all types of elements of a Java program.

sandmark.statistics.Statistics:

2.2 Dynamic Class Loading

Dynamic class loading is used in several places in sandmark to allow for extensibility. The major problem in dynamic class loading in sandmark is finding classes to load. In particular, sandmark can be run from a directory or a jar file, and it may or may not know the location of that directory or jar file or even which of the two it is running from.

All of this detail is abstracted by sandmark.util.classloading.ClassFinder. ClassFinder allows a caller to specify one of a small number of classes (defined in sandmark.util.classloading.IClassFinder), and ClassFinder will return the names of all classes derived from that class. For example, the following code prints to standard output all classes derived from sandmark.Algorithm.

```
1  java.util.Collection c =
2     sandmark.util.classloading.ClassFinder.getClassesWithAncestor
3        (sandmark.util.classloading.IClassFinder.ALGORITHM);
4  java.util.Iterator it = c.iterator();
5  while(it.hasNext()) {
6     String className = (String)c.next();
7     System.out.println(className);
8  }
```

The class names returned by ClassFinder are those returned by an implementor of sandmark.util.classloading.IClassFinder. ClassFinder uses the first implementation in the following list whose default constructor does not throw an exception:

1. **FileClassFinder** This class reads class names and from which of the interfaces in sandmark.util.classloading.IClassFinder they derive from a file call **Algorithms.txt**. This file is opened by a call to **ClassLoader.getSystemClassLoader().getResourceAsStream("Algorithms.txt")**. If this call fails, FileClassFinder's default constructor throws an exception.

2. **JarClassFinder** This class searches through all entries in the jar file named by System.getProperty("SMARK_PATH"), loads all class files contained therein, and determines from which classes in IClassFinder each of these classes derives. If opening the jar file for reading throws an exception, JarClassFinder's default constructor throws an exception.

3. **DirClassFinder** This class is similar to JarClassFinder, but it searches through the file system hierarchy rooted at System.getProperty("SMARK_ROOT"). Its default constructor throws an exception if the contents of SMARK_ROOT cannot be read.

Chapter 3

Extending SandMark

3.1 Adding an Obfuscator

SandMark is designed to make it easy to add a new obfuscation algorithm. Assume that we want to add a new obfuscation `ReorderMethods`. The process would be the following:

1. Create a new directory `sandmark/obfuscate/reorderMethods`.
2. Create a new class `sandmark/obfuscate/reorderMethods/ReorderMethods.java`.
3. The `ReorderMethods` class should extend one of the base classes `AppObfuscator` (if the algorithm works on the entire user program), `MethodObfuscator` (if the algorithm works one method at a time), or `ClassObfuscator` (if the algorithm works on one individual class at a time). Let's assume that our algorithm reorders methods within one class. `ReorderMethods` should therefore extend `ClassObfuscator`, which looks like this:

```
1   package sandmark.obfuscate;
2   public abstract class ClassObfuscator extends GeneralObfuscator {
3   protected ClassObfuscator(String label) {
4       super(label);
5   }
6   abstract public void apply(
7       sandmark.util.ClassFileCollection cfc, String classname)
8       throws Exception;
9   }
10  public String toString() {
11      return "ClassObfuscator(" + getLabel() + ")";
12  }
13  }
```

4. The `ReorderMethods` class should look something like this:

```
1   public class ReorderMethods extends sandmark.obfuscate.ClassObfuscator {
2      public ReorderMethods(String label) {
3         super(label);
4      }
5      public void apply(
6          sandmark.util.ClassFileCollection cfc, String classname) throws Exception {
7         // Your code goes here!
8      }
9   }
```

5. Use `BCEL` or `BLOAT` to implement your obfuscation. The `cfc` parameter represents the set of classes to be obfuscated. Use routines in `sandmark.util.ClassFileCollection` to open a class to be edited by `BCEL` or `BLOAT`.

6. Type `make` at the top-level sandmark directory (`smark`). The new obfuscation should be loaded automagically at runtime.

3.2 Obfuscation Configuration

To customize Obfuscation Algorithm Application and Obfuscation Level Settings for a whole jar file ,or a class file and its methods or a method alone the configure button in the obfuscation panel is to be used .On clicking The Configure button in the obfuscate panel , a new window pops up which is known as Obfuscation Dialog ,which helps the users to select specific obfuscation algorthims and obfuscation level .

Method Summary

```
2   /*  Find if a particular algorithm is marked for obfuscation in the case
3    *  of a class or jar file . Default is that all algorithms are to be applied
4    *  for all classes and methods in the jar file .
5    */
6       public boolean IsMarked(String algo,String class_or_jarname);
7
8   /*  Find if a particular algorithm is marked for obfuscation in the case
9    *  of a method . Please specify the  signature enclosed in brackets.
10   */
11      public boolean IsMarked(String algo,String classname,String methodname,
12           String signature);
13
14  /*  Obtain the obfuscation level for a class or jar file. Possible String
15   *  values are high , medium, low,none. Default is high .
16   */
17      public String getLevel(String class_or_jarname) ;
18
19  /*  Obtain the obfuscation  level for a method .The  signature should
20   *   be enclosed in brackets .Possible values are high ,medium,low,none.
21   *   Default is high .
22   */
23      public String getLevel(String classname,String methodname,String signature)
```

3.3 Adding a Watermarker

Adding a new watermarking algorithm is similar to adding an obfuscator. Algorithms are loaded dynamically at run-time, so there is no need to explcitly link them into the system.

To create a new watermarking algorithm `wm` you

1. create a new directory `sandmark.watermark.wm`,

2. create a new class `sandmark.watermark.wm.WM` which extends `sandmark.watermark.StaticWatermarker` or `sandmark.watermark.DynamicWatermarker`. To build a new static watermarker you just have to implement two methods, one to embed the watermark into a jarfile and the other to extract it:

```
1   package sandmark.watermark;
2
3   public abstract class StaticWatermarker
```

```
4      extends sandmark.watermark.GeneralWatermarker {
5
6   public StaticWatermarker() {}
7
8   /* Embed a watermark value into the program. The props argument
9    * holds at least the following properties:
10   *  <UL>
11   *     <LI> WM_Encode_Watermark: The watermark value to be embedded.
12   *     <LI> WM_Embed_JarInput: The name of the file to be watermarked.
13   *     <LI> WM_Embed_JarOutput: The name of the jar file to be constructed.
14   *     <LI> SWM_Embed_Key: The secret key.
15   *  </UL>
16   */
17  public abstract void embed(
18     java.util.Properties props)
19        throws sandmark.watermark.WatermarkingException,
20               java.io.IOException;
21
22
23  /* Return an iterator which generates the watermarks
24   * found in the program. The props argument
25   * holds at least the following properties:
26   *  <UL>
27   *     <LI> WM_Recognize_JarInput: The name of the file to be watermarked.
28   *     <LI> SWM_Recognize_Key: The secret key.
29   *  </UL>
30   */
31  public abstract java.util.Iterator recognize(
32     java.util.Properties props)
33        throws sandmark.watermark.WatermarkingException,
34               java.io.IOException;
35
36  }
```

3. Use `BCEL` or `BLOAT` to implement your watermarker. Have a look at the trivial static watermarker `sandmark.watermark.constantstring.ConstantString` for an example.

4. Type `make` at the top-level sandmark directory (`smark`). The new watermarker should be loaded automagically at runtime.

Implementing a dynamic watermarker is more complex, since you have to provide methods for running the application during tracing and recognition:

```
1   package sandmark.watermark;
2
3   public abstract class DynamicWatermarker
4      extends sandmark.watermark.GeneralWatermarker {
5
6   /**
7    * Start a tracing run of the program. Return an iterator
8    * object that will generate the trace points encountered
9    * by the program.
10   */
11  public abstract java.util.Iterator startTracing (
12      java.util.Properties props) throws sandmark.util.exec.TracingException;
```

```
13
14  /*
15   * Wait for the program being traced to complete.
16   */
17  public abstract void waitForTracingToComplete()
18     throws sandmark.util.exec.TracingException;
19
20  /**
21   * This routine should be called when the tracing run has
22   * completed. tracePoints is a vector of generated
23   * trace points generated by the iterator returned by
24   * startTracing.
25   */
26  public abstract void endTracing(
27      java.util.Properties props,
28      java.util.Vector tracePoints) throws sandmark.util.exec.TracingException;
29
30  /**
31   * Force the end to a tracing run of the program.
32   */
33  public abstract void stopTracing(
34      java.util.Properties props) throws sandmark.util.exec.TracingException;
35
36
37  /* Embed a watermark value into the program. The props argument
38   * holds at least the following properties:
39   *  <UL>
40   *      <LI> WM_Encode_Watermark: The watermark value to be embedded.
41   *      <LI> DWM_Embed_TraceInput: The name of the file containing trace data.
42   *      <LI> WM_Embed_JarInput: The name of the file to be watermarked.
43   *      <LI> WM_Embed_JarOutput: The name of the jar file to be constructed.
44   *      <LI> DWM_CT_Encode_ClassName: The name of the Java file that builds the watermark.
45   *  </UL>
46   */
47  public abstract void embed(
48     java.util.Properties props);
49
50  /**
51   * Start a recognition run of the program.
52   */
53  public abstract void startRecognition (
54      java.util.Properties props) throws sandmark.util.exec.TracingException;
55
56  /**
57   * Return an iterator object that will generate
58   * the watermarks found in the program.
59   */
60  public abstract java.util.Iterator watermarks();
61
62  /**
63   * Force the end to a tracing run of the program.
64   */
65  public abstract void stopRecognition(
```

```
66       java.util.Properties props) throws sandmark.util.exec.TracingException;
67
68   }
```

3.4 Adding a Graph Codec

Several watermarking algorithms encode the watermark as a graph. SandMark contains several methods for making this encoding, stored in the `sandmark.util.graph.codec` package.

Adding a new graph coder/decoder *codec* algorithm is similar to adding an obfuscator or watermarker: just add a new class to the `codec` directory, make sure it extends the appropriate class, type `make`, and the new algorithm will have been added to the system.

Every graph codec should extend `sandmark.util.graph.codec.GraphCodec`:

```
1   package sandmark.util.graph.codec;
2   public abstract class GraphCodec {
3      public java.math.BigInteger value = null;
4      public sandmark.util.graph.Graph graph = null;
5
6   /**
7    * Codecs should implement this method to convert
8    * the 'value' into 'graph'.
9    */
10     abstract void encode();
11
12  /**
13   * Codecs should implement this method to convert
14   * the 'graph' into 'value'. Whenever the decoding
15   * failes (eg. because the graph has the wrong
16   * shape) the codec should simply throw an exception.
17   */
18     abstract void decode() throws sandmark.util.graph.codec.DecodeFailure;
19
20  /**
21   * Constructor to be used when encoding an integer into a graph.
22   *   @param value   The value to be encoded.
23   */
24     public GraphCodec (
25         java.math.BigInteger value) {
26         this.value = value;
27         encode();
28     }
29
30  /**
31   * Constructor to be used when decoding a graph to an integer.
32   *   @param graph   The graph to be decoded.
33   *   @param root    The root of the graph.
34   *   @param kidMap  An array of ints describing which field
35   *                  should represent the first child, the
36   *                  second child, etc.
37   */
38     public GraphCodec (
39         sandmark.util.graph.Graph graph,
40         int kidMap[]) throws sandmark.util.graph.codec.DecodeFailure {
```

```
41          this.graph = graph;
42          this.kidMap = kidMap;
43          decode();
44      }
45  }
```

Note that there are two constructors. One is used when encoding an integer (`java.meth.BigInteger`) into a graph (a `sandmark.util.graph.Graph`), and another when decoding a graph into an integer.

For example, the simplest graph codec, `RadixGraph` looks like this:

```
1   package sandmark.util.graph.codec;
2
3   public class RadixGraph extends sandmark.util.graph.codec.GraphCodec {
4
5       public final static String FULLNAME  = "Radix Graph";
6       public final static String SHORTNAME = "radix";
7
8   // Used when encoding.
9   public RadixGraph (java.math.BigInteger value)  {
10      super(value);
11      this.fullName = FULLNAME;
12      this.shortName = SHORTNAME;
13  }
14
15  void encode() { ... }
16
17  // Used when decoding.
18     public RadixGraph (
19          sandmark.util.graph.Graph graph,
20          int kidMap[]) throws sandmark.util.graph.codec.DecodeFailure {
21          super(graph, kidMap);
22          this.fullName = FULLNAME;
23          this.shortName = SHORTNAME;
24      }
25      void decode() throws sandmark.util.graph.codec.DecodeFailure { ... }
26  }
```

3.5 Documentation

To document a new obfuscator or watermarker, do the following:

1. create a new file *name*`.tex` in `smark/doc`:

   ```
   \algorithm{The ... Algorithm}{Authors}
    ....
   ```

   Where `Authors` is a comma-separated list of authors of this algorithm implementation.

2. add an `input`-statement to `smark/manual.tex`:

   ```
   \part{\SM\ Algorithms}
   \input{CollbergThomborson}
   ...
   \input{name.tex}
   ```

3. Update `smark/makefile`:

```
SECTIONS = \
   ...
   CollbergThomborson.tex\
   ...
   name.tex \
```

4. Type `make`.

Part II

SandMark Algorithms

Chapter 4

The Collberg-Thomborson Watermarking Algorithm

C. Collberg, J. Nagra, G. Townsend

4.1 Introduction

The Collberg-Thomborson watermarking algorithm (henceforth, CT) is a dynamic algorithm. The idea is that rather than embedding the watermark directly in the *code* of the application, code is embedded that *builds* the watermark at runtime. The algorithm assumes a secret key $\mathcal{K}$ which is necessary to extract the watermark. $\mathcal{K}$ is a sequence of inputs $I_0, I_1, \ldots$ to the application. As seen in Figure 4.1, the watermark (a graph structure) is built by the application only when the user runs it with the special input $I_0, I_1, \ldots$. Figure 4.3 shows a simple example of a what a program may look like after having been watermarked.

In the SandMark implementation of CT, watermark embedding and extraction runs in several steps (See Figure 4.2):

Annotation: Before the watermark can be embedded the user must add *annotation* (or *mark*) points into the application to be watermarked. These are calls of the form

```
1    sandmark.trace.Annotate.mark();
2    String S = ...;
3    sandmark.trace.Annotate.mark(S);
4    long L = ...;
5    sandmark.trace.Annotate.mark(L);
```

The `mark()` calls perform no action. They simply indicate to the watermarker locations in the code where (part of) a watermark-building code can be inserted. The argument to the `mark()` call can be any string or integer expression that (directly or indirectly) depends on user input to the application.

Tracing: When the application has been annotated the user should do a *tracing* run of the program. The application is run with the chosen secret input sequence, $\mathcal{K}$. During the run one or more annotation points are hit. Some of these points will be the locations where watermark-building code will later be inserted.

Embedding: During the embedding stage the user enters a watermark, a string or an integer. A string is converted to an integer. From this number a graph is generated, such that the topology of the graph embeds the number. The graph is split into a number of subgraphs, depending on the number of locations where watermarking code should be inserted. Each subgraph is converted to Java bytecode that builds the graph. The relevant `mark()`-calls are replaced with this graph-building code.

Recognition: During recognition the application is again run with the secret input sequence as input. The same `mark()`-locations will be hit as during the tracing run. Now, however, these locations will contain code for building the watermark graph. When the last part of the input has been entered, the heap is examined for graphs that could potentially be watermark graphs. The graphs are decoded and the resulting watermark number or string is reported to the user.

4.2 Annotation

The CT watermark consists of dynamic data-structures. This means that the code inserted in the application will look like this:

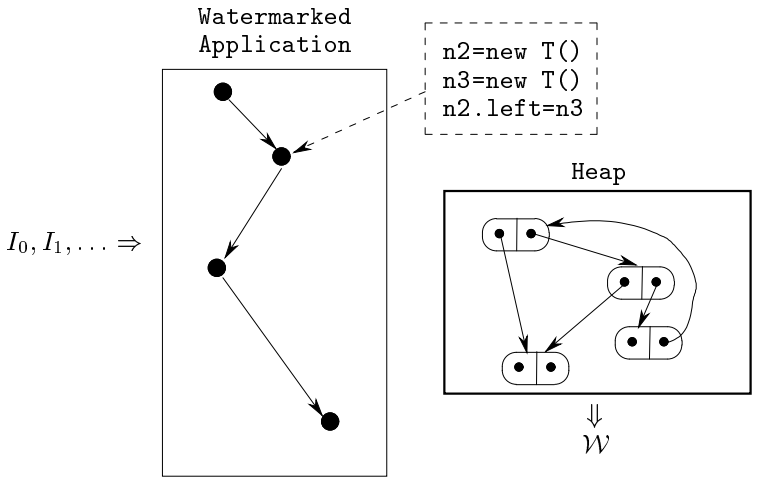

Figure 4.1: Overview of how the CT algorithm recognizes a watermark $\mathcal{W}$. At runtime the watermarked application will – given the special secret input key sequence $I_0, I_1, \ldots$ – traverse certain points in the program. At these points code has been inserted which builds a graph $G_{\mathcal{W}}$ on the heap. The topology of the graph embeds the watermark $\mathcal{W}$.

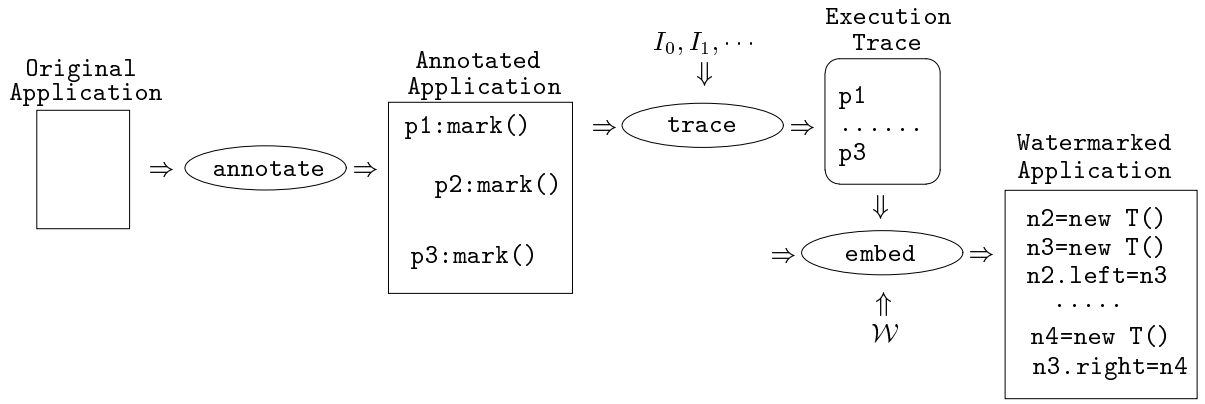

Figure 4.2: Overview of how the CT algorithm watermarks an application. First, the user adds *annotation points* (`mark()`-calls) to the application. These are locations where watermarking code may be inserted. Secondly, the application is run with a secret input sequence, $I_0, I_1, \ldots$ and the trace of `mark()`-calls hit during this run is recorded. Finally, code is embedded into the application (at certain `mark()`-call locations) that builds a graph G_W at runtime. The topology of $G_{\mathcal{W}}$ embeds the watermark $\mathcal{W}$.

```
1   public class Simple {
2      static void P(String i) {
3         System.out.println("Hello " + i);
4      }
5      public static void main(String args[]) {
6         P(args[0]);
7      }
8   }
9
10  public class Simple_W {
11     static void P(String i) {
12        if (i.equals("World")) Watermark.Create_G4();
13        System.out.println("Hello " + i);
14     }
15     public static void main(String args[]) {
16        Watermark.Create_G2();
17        P(args[0]);
18     }
19  }
20
21  public class Watermark extends java.lang.Object {
22    public Watermark edge1;
23    public Watermark edge2;
24    public static java.util.Hashtable sm$hash = new java.util.Hashtable();
25    public static Watermark[] sm$array = new Watermark[4];
26
27    public static void Create_G2 () {
28       Watermark n3 = new Watermark();
29       Watermark n2 = new Watermark();
30       Watermark.sm$array[1] = n2;
31       n2.edge1 = n3;
32       n2.edge2 = n3;
33    }
34    public static void Create_G4 () {
35       Watermark n1 = new Watermark();
36       Watermark n4 = new Watermark();
37       Watermark.sm$hash.put(new java.lang.Integer(4), n4);
38       n4.edge1 = n1;
39       Watermark n2 = Watermark.sm$array[1];
40       n1.edge1 = n2;
41       Watermark n3 = (n2 != null)?n2.edge1:new Watermark();
42       n3.edge1 = n1;
43    }
44  }
```

Figure 4.3: Simple watermarking example. The class `Simple` is modified into `Simple_W` by adding calls into the generated watermark class `Watermark`. The static methods `Watermark.Create_G2()` and `Watermark.Create_G4()` are only called when the application is run with the secret input argument `"World"`. When this happens, the watermark graph is built on the heap.

```
1       Node n1 = new Node();
2       Node n2 = new Node();
3       n1.edge = n2;
4       ...
```

Hence, we should prefer mark locations that

- allocate objects and manipulate pointers, and
- directly depend on user input.

We should avoid mark locations that

- are hot-spots, and
- are executed non-deterministically.

In other words, `mark()`-calls should be added to locations where the resulting watermark code will be fit in (is *stealthy*), won't affect performance, and will be executed consistently from run to run, depending only on user actions.

For example, the following code is undesirable since `Math.random()` may generate different values during different runs of the program:

```
1       if (Math.random() < 0.5) {
2          ...
3          sandmark.trace.Annotate.mark();
4       }
```

Similarly, if thread scheduling, network activity, processor load, etc. can affect the order in which some locations are executed, these locations are not valid annotations points and should be avoided.

4.3 Tracing

SandMark makes heavy use of Java's `JDI` (*Java Debugging Interface*) framework. During tracing and recognition SandMark starts up the user's application as a subprocess running under debugging. This allows SandMark to set breakpoints, examine variables, and step through the application – all the operations that can be done under an interactive debugger. During tracing we are interested in obtaining a trace of the `mark()`-calls that are hit while the user enters their secret input. We also want to know the argument to the `mark()`-call and the stack trace at the point of the call.

Unfortunately, `JDI` is not yet a perfect product and we have to jump through a couple of hoops to make it do what we want. First of all, examining the value of the argument to the `mark()`-call may or may not work. Examining static global variables seems to work, however, so we always start by storing the argument in a global, and then call the placeholder method `MARK()`. See Figure 4.4. During tracing we only have to put a breakpoint on the `MARK()` method.

The second problem is that we need a stack trace at the point of each `mark()`-call. This trace is used during embedding to compute an accurate call-graph of the program at each `mark()` location. The call graph allows us to compute ways to pass information between `mark()`-calls in method parameters. While `JDI` allows us to examine the stack frames at any point in the program, it is not possible to tell if two stack-frames are *the same*. That is, JDI stack-frames do not have unique identities. To solve this problem we add the following statement to the beginning of every method in the program:

```
1       long sm$stackID = sandmark.trace.Annotator.stackFrameNumber++;
```

This is done prior to tracing in `sandmark.trace.Preprocessor`.

When, during tracing, a `mark()`-call is hit we walk the stack, collecting the `sm$stackID`s in each frame.

At the end of tracing run we have gathered a list of `sandmark.trace.TracePoint`-objects. Each object represents a `mark()`-call that was hit during the trace and contains three pieces of information:

```
1   package sandmark.trace;
2   public class Annotator {
3      static String VALUE = "";
4      public static long stackFrameNumber=0;
5      public static void MARK(){}
6      public static void mark() {
7         long sm$stackID = sandmark.trace.Annotator.stackFrameNumber++;
8         VALUE = "----";
9         MARK();
10     }
11     public static void mark(String s) {
12        long sm$stackID = sandmark.trace.Annotator.stackFrameNumber++;
13        VALUE = "\"" + s + "\"";
14        MARK();
15     }
16     public static void mark(long v) {
17        long sm$stackID = sandmark.trace.Annotator.stackFrameNumber++;
18        VALUE = Long.toString(v);
19        MARK();
20     }
21  }
```

Figure 4.4: The class `sandmark.trace.Annotate`.

1. the location in the bytecode where the `mark()` was located (a `sandmark.util.ByteCodeLocation`);
2. the value that the user supplied as an argument to the `mark()`-call (a `String`);
3. a list of the stack-frames active when the `mark()`-call was hit (`sandmark.util.StackFrame[]`).

An Example

Consider the following example application:

```
1      public class SimpleA {
2         static void P(int i) {
3            sandmark.trace.Annotator.mark(6*i+9);
4         }
5         public static void main(String args[]) {
6            P(3);
7         }
8      }
```

After tracing we will have found only one trace point. It is described by a structure like this:

$$\langle \texttt{value} = 27, \texttt{location} = \langle \texttt{P}, \texttt{pc} = 8\rangle, \texttt{stack} = [\langle \texttt{P}, \texttt{pc} = 8, \texttt{frame} = 1\rangle, \langle \texttt{main}, \texttt{pc} = 8, \texttt{frame} = 0\rangle]\rangle$$

We have stored the argument to the `mark()` call (`value=27`), the bytecode location where that call was made (`pc=8`), and complete stack trace (with unique identifiers for each frame) at this location.

Choosing `mark()`-Locations

The tracing phase will have generated one or more `mark()`-locations. However, cannot be used to build the watermark graph and have to be removed. Also, we need k locations to insert code to build a k-component graph, and any extra locations should be deleted.

`sandmark.watermark.CT.embed.PrepareTrace` examines the trace to find a set of `mark()`-locations that can be used to build the watermark graph. An annotation point $\langle value, location \rangle$ is valid if

1. there is exactly one trace point at *location*, or
2. there are multiple trace points at *location*, but they all have unique *value*s.

For example, consider the following `mark()`-points:

$$\langle -, L_0 \rangle$$
$$\langle 1, L_1 \rangle$$
$$\langle 1, L_1 \rangle$$
$$\langle 10, L_2 \rangle$$
$$\langle 11, L_2 \rangle$$
$$\langle 12, L_2 \rangle$$

`value` $=-$ is used for `mark()`-calls that take no argument. $\langle -, L_0 \rangle$ is valid, because it is the only `mark()`-point at location L_0. $\langle 1, L_1 \rangle$ is not valid because there are two identical annotation values at this location. If we were to insert watermark-building code at this location we would not be able to tell the difference between the first and the second time we arrive. $\langle 10, L_2 \rangle, \langle 11, L_2 \rangle, \langle 12, L_2 \rangle$ are valid because the *value*s are unique. If there is one unique value at a location, this `mark()`-call is said to be `LOCATION`-based, otherwise it is `VALUE`-based.

4.4 Embedding

Once the application has been traced we can finally start embedding the watermark. The input to this phase is tracing information (as described in the previous section), a watermark $\mathcal{W}$ to be embedded, and a jar-file containing the classfiles in which to embed the mark. The embedding is divided into five phases:

1. First we generate a graph G whose topology embeds $\mathcal{W}$.
2. Next, we split G into k subgraphs $\langle G_1, \ldots, G_k \rangle$.
3. From each subgraph G_i we generate an *intermediate code* C_i that builds this graph and connects it to the subgraphs $\langle G_1, \ldots, G_{i-1} \rangle$.
4. We translate each intermediate code C_i into a Java method M_i that, when executed, will build G_i.
5. Finally, based on the tracing information, we replace some of the `mark()`-calls with calls to one of the M_i-methods. The remaining `mark()`-calls are removed.

The result is a new jar-file that when executed with the special input sequence will execute the methods $\langle M_1, \ldots, M_k \rangle$ (in this order), and, consequently, build the watermark graph G on the heap.

Building the Graph

Eventually we hope to have a whole library of algorithms for building watermark graphs. Currently, only two have been implemented. The algorithms are located in `sandmark.util.graph.codec`. See Section 3.4 for a description of how to add a new *Graph Codec* to SandMark.

Radix Encoding is the simplest algorithm. The codec is in `sandmark.util.graph.codec.RadixGraph`. Figure 4.5 illustrates the idea of a Radix-k encoding using a circular linked list. An extra pointer field encodes a base-k digit in the length of the path from the node back to itself. A null-pointer encodes a 0, a self-pointer a 1, a pointer to the next node encodes a 2, etc.

The *Permutation Encoding* codec is in `sandmark.util.graph.codec.PermuationGraph`. The idea is to represent the watermark $\mathcal{W}$ by a permutation of the numbers $\langle 0, \ldots, n-1 \rangle$. For example, the number

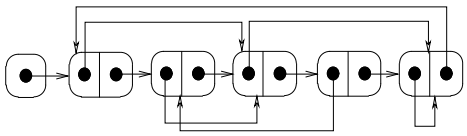

$$61 \times 73 = 3 \cdot 6^4 + 2 \cdot 6^3 + 3 \cdot 6^2 + 4 \cdot 6^1 + 1 \cdot 6^0$$

Figure 4.5: Radix-6 encoding. The right pointer field holds the `next` field, the left pointer encodes a base-k digit.

could be represented by the permutation

$$\langle 9, 6, 5, 2, 3, 4, 0, 1, 7, 8 \rangle$$

of the numbers

$$\langle 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 \rangle.$$

A permutation of length n is encoded into a graph structure similar to the one in Figure 4.5. The nodes of the graph form a circular linked list and a pointer from node i to node j represents the fact that i has been permuted with j.

It should be noted that the graphs we use in the CT algorithm are, in fact, hyper-graphs. They are implemented by the package `sandmark.util.graph`.

Splitting the Graph

When the watermark graph has been built it needs to be split into pieces. This is done in the package `sandmark.watermark.CT.encode.Split`. There should be one graph component per `mark()`-location that we intend to use. There are three things to consider when we split the graph:

1. The subgraphs should be of roughly equal size. (It is actually not quite clear that this is a reasonable requirement. For stealth reasons it might be better if the components are of random size. The current implementation, however, splits in equal-size pieces.)

2. The splitting of G should be done in such a way that each subgraph has a root, a special node from which all other nodes in the graph can be reached. This allows us to store only pointers to root nodes to prevent the garbage collector from collecting the subgraphs. (More about this later.)

3. We should attempt to split G in such a way that the number of edges between subgraphs is minimized. The reason for this restriction is that the more edges there are between subgraphs, the more Java code we will have to generate in order to connect the subgraphs into the complete graph G.

We use a graph-splitting algorithm by Kundu and Misra. It would, for example, split the graph on the left into the two components on the right:

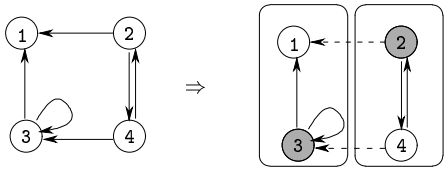

Root nodes have been shaded and inter-component edges have been dashed.

INSTRUCTION	DESCRIPTION
`AddEdge(`G_i`,`G_j`,`$n \xrightarrow{\texttt{edge}} m$`)`	Add an edge from node n in subgraph G_i to node m in G_j. Since the graphs are multi-graphs the out-edges are named.
`CreateNode(`G_i`,`n`)`	Create node n in subgraph G_i.
`CreateStorage(`G`,`S`)`	Create the global storage structure S.
`Debug(`*msg*`)`	Insert debugging code.
`FollowLink(`G_i`,`$n \xrightarrow{\texttt{edge}} m$`)`	Return m by following the edge `edge` from n.
`LoadNode(`G_i`,`n`,`L`)`	Load node n from global storage location L.
`PrintGraph()`	Insert code for printing the graph. Used for debugging.
`ProtectRegion(`*ops*`)`	The instructions *ops* may generate runtime errors, such as `null` dereference. Protect against such errors by, for example, putting *ops* inside a `try`-block.
`SaveNode(`G_i`,`n`,`L`)`	Save node n in global storage location L.

Table 4.1: Intermediate code instructions.

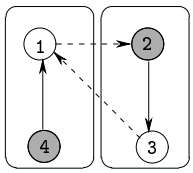

Figure 4.6: Two graph components G_2 and G_4. Components are named after their (shaded) root nodes.

Generating Intermediate Code

We could, of course, generate Java code directly from the graph components. However, it turns out to be advantageous to insert one intermediate step. From each graph component we generate a list of *intermediate code instructions*, much in the same way a compiler might generate an intermediate representation of a program, in anticipation of code generation and optimization. In a compiler, the intermediate code separates the front-end from the back-end, improving retargetability, and also providing a target-independent representation for optimizing transformations. Similarly, our intermediate representation provides

1. retargetability, in case one day we may want to generate C++ or C# code; and

2. transformability, i.e. the ability to optimize or otherwise transform the intermediate code prior to generating Java code.

In fact, we start by generating very simple intermediate code, and then run several transformations (in `sandmark.watermark.CT.encode.ir2ir`) over the code to optimize it, etc.

The intermediate code instructions are defined in the package `sandmark.watermark.CT.encode.ir`. The main operations are given in Table 4.1.

Consider the two graph components G_2 and G_4 in Figure 4.6. The following intermediate code is generated from G_2:

`create(`G_2`)`
  n_3 `= CreateNode(`G_2`)`
  n_2 `= CreateNode(`G_2`)`
  `SaveNode(`n_2`, `G_2`, '`n_2`:Array/global')`
  `AddEdge(`$n_2 \xrightarrow{\texttt{edge1}} n_3$`, `$G_2$`, `$G_2$`)`
  `AddEdge(`$n_2 \xrightarrow{\texttt{edge2}} n_3$`, `$G_2$`, `$G_2$`)`

Nodes are named n_1, n_2, etc. The `SaveNode` instruction is used to store the root node of a graph component in a global structure such as a hash table, vector, etc. We do this for two reasons:

1. Suppose we have two subgraphs G_1 and G_2, where G_1 is created first. After G_2 has been created, the two graphs need to be connected. At this point we need (at least) pointers to both their root nodes, so that we can access and link any two nodes in the graphs.

2. Every node in every subgraph must, at all times, be live or it may be deleted by the garbage collector.

As we will see later, we can often do away with these global pointers by passing root nodes as method arguments. This is much stealthier since most programs have few global variables but many method parameters.

Here is the intermediate code generated from the subgraph component G_4 in Figure 4.6:

```
create(G4)
    n1 = CreateNode(G4)
    n4 = CreateNode(G4)
    SaveNode(G4, n4, 'n4:Hash/global')
    AddEdge(G4, G4, n4 --edge1--> n1)
    n2 = LoadNode(G4, 'n2:Array/global')
    AddEdge(G4, G2, n1 --edge1--> n2)
    n3 := FollowLink(G2, n2 --edge1--> n3)
    AddEdge(G2, G4, n3 --edge1--> n1)
```

Note how node n_2 from graph G_2 as been loaded from global storage in order to connect n_1 to n_2. Note also how the `FollowLink` instruction is used to traverse G_2 from n_2 to get to node n_3, which can then be connected to n_1.

To generate intermediate code from a subgraph G_i we perform a depth-first search from the root of the graph. This is done in `sandmark.watermark.CT.encode.Graph2IR`. `CreateNode(`n`)`-instructions are generated from each node, in a reverse topological order. That is, leaves are generated first and the root node last. We can issue an `AddEdge(`G_i`, `G_i`, `$n \xrightarrow{\texttt{edge}} m$`)`-instruction as soon as `CreateNode(`m`)` and `CreateNode(`n`)` have both been generated.

The code for G_i must also contain instructions connecting G_i to all the previous subgraphs $G_0, \ldots, G_{i-1}$. If there is an inter-subgraph edge $m \xrightarrow{\texttt{edge}} n$ from G_k to G_i (i.e. m is a node in G_k and n is in G_i) then we must generate

1. one or more `FollowLink()`-instructions to reach node m by traversing G_k starting at its root node,

2. and a final `AddEdge()` instruction to link m to n.

This is done by finding the shortest path from k (the root of subgraph G_k) and m and issuing a `FollowLink()`-instruction for each.

`sandmark.watermark.CT.encode.Graph2IR` generates the basic list of intermediate instructions. These "programs" are then optimized and transformed in various ways by the transformers in `sandmark.watermark.CT.encode.ir2ir.*`. Some of the more important ones are

sandmark.watermark.CT.encode.ir2ir.AddFields:

sandmark.watermark.CT.encode.ir2ir.AddFormals: Add parameters to the `Create_`G_i methods to allow graph roots to be passed in formals rather than globals. (This will be described in more detail in Section 4.6).

sandmark.watermark.CT.encode.ir2ir.CleanUp: `LoadNode()`-instructions are added (in `sandmark.watermark.CT.encode.ir2ir.SaveNodes`. We do this in a greedy way, and the `CleanUp` transformer removes redundant loads.

sandmark.watermark.CT.encode.ir2ir.Debug: Add `Debug()`-instructions. These will print out trace messages as the watermark graphs are built at runtime.

sandmark.watermark.CT.encode.ir2ir.Destructors: Create bogus graph builders/destroyers that can be inserted in various places in the code. The destructors are created by modifying copies of the creator codes.

sandmark.watermark.CT.encode.ir2ir.InlineFixups: `sandmark.watermark.CT.encode.Graph2IR` generates special methods (called `Fixup_`G_i`_`G_j) which add the inter-graph links that connect too subgraphs G_i and G_j. Normally these are inlined into the code for G_j by this transformer.

sandmark.watermark.CT.encode.ir2ir.Protect: Add protection code when this is necessary to prevent `FollowLink()`-instructions from throwing unwanted exceptions at runtime.

sandmark.watermark.CT.encode.ir2ir.SaveNodes: Add code to load and store graph roots into global, static storage.

Generating Java Code

Generating Java code from the intermediate representation is relatively straight-forward. We use the `BCEL` library to generate a bytecode class `Watermark`. We can generate Java source also, but this is mostly used for debugging.

The intermediate code from the previous section is translated into the Java class in Figure 4.7. Method `Create_G2` builds subgraph G_2 and `Create_G4` subgraph G_4. Note, in particular, the statements

```
1       Watermark n2 = Watermark.sm$array[1];
2       Watermark n3 = (n2 != null)?n2.edge1:new Watermark();
3       n3.edge1 = n1;
```

which link nodes n_3 and n_1. To get access to G_2's node n_3 we follow the edge from G_2's root node n_2 to n_3. This will work provided G_2 has been created at this point. However, if we're not doing a recognition run (i.e. the input sequence in *not* $I_0, I_1, \ldots$) then G_2 may not have been created in which case `n2` may be `null`. We can protect against this in a variety of ways:

1. if `n2` is `null` we create a new node and assign it to `n2` (as above);

2. we can enclose the entire code segment in a `try-catch`-block:

```
1       Watermark n2 = Watermark.sm$array[1];
2       try {
3          Watermark n3 = n2.edge1;
4          n3.edge1 = n1;
5       } catch (Exception e){}
```

   ; or

3. we may simply not do the assignment if `n2` is `null`:

```
1       Watermark n2 = Watermark.sm$array[1];
2       if (n2 != null) {
3          Watermark n3 = n2.edge1;
4          n3.edge1 = n1;
5       }
```

It is useful to have a whole library of such protection mechanisms to prevent attacks by pattern matching.

```
1  public class Watermark extends java.lang.Object {
2    public Watermark edge1;
3    public Watermark edge2;
4    public static java.util.Hashtable sm$hash;
5    public static Watermark[] sm$array;
6
7    public static void Create_G2 () {
8       Watermark n3 = new Watermark();       // n3 = CreateNode(G2)
9       Watermark n2 = new Watermark();       // n2 = CreateNode(G2)
10      Watermark.sm$array[1] = n2;           // SaveNode(n2, G2, 'n2:Array/global')
11      n2.edge1 = n3;                        // AddEdge(G2, G2, n2-edge1->n3)
12      n2.edge2 = n3;                        // AddEdge(G2, G2, n2-edge2->n3)
13   }
14
15   public static void Create_G4 () {
16      Watermark n1 = new Watermark();       // n1 = CreateNode(G4)
17      Watermark n4 = new Watermark();       // n4 = CreateNode(G4)
18      Watermark.sm$hash.put(
19          new java.lang.Integer(4), n4);    // SaveNode(n4, G4, 'n4:Hash/global')
20      n4.edge1 = n1;                        // AddEdge(G4, G4, n4-edge1->n1)
21      Watermark n2 = Watermark.sm$array[1]; // n2 = LoadNode(G4, 'n2:Array/global')
22      n1.edge1 = n2;                        // AddEdge(G4, G2, n1-edge1->n2)
23      Watermark n3 =                        // n3 := FollowLink(G2, n2-edge1->n3)
24         (n2 != null)?n2.edge1:new Watermark();
25      n3.edge1 = n1;                        // AddEdge(G2, G4, n3-edge1->n1)
26   }
27 }
```

Figure 4.7: Java code generated from the graph components in Figure 4.6.

INSTRUCTION	JAVA
AddEdge(G_i,G_j,$n \xrightarrow{\text{edge}} m$)	`n.edge = m`
CreateNode(G_i,n)	`Watermark n = new Watermark()`
CreateStorage(G,S)	One of 1. `static java.util.Hashtable sm$hash = new java.util.Hashtable()`; 2. `static Watermark sm$array = new Watermark[`m`]`; 3. `static java.util.Vector sm$vec = new java.util.Vector(`m`); sm$vec.setSize(`m`)`; 4. `static Watermark sm$n1,sm$n2,...`; where m is the number of nodes in the graph and `sm$n1,sm$n2,...` are the root nodes of the subgraphs.
FollowLink(G_i,$n \xrightarrow{\text{edge}} m$)	`Watermark m = n.edge`
LoadNode(G_i,n,S)	One of 1. `Watermark` n `= (Watermark) sm$hash.get(new java.lang.Integer(`k`))`; 2. `Watermark` n `= Watermark.sm$arr[`$k-1$`]`; 3. `Watermark` n `= (Watermark) sm$vec.get(`$k-1$`)`; 4. `Watermark` n `= Watermark.sm$n`k depending on how n is stored. k is n's node number.
ProtectRegion(*ops*)	`try {` *ops* `} catch (Exception` e`) {}`
SaveNode(G_i,n,L)	One of 1. `sm$hash.put(new java.lang.Integer(`k`),` n`)`; 2. `Watermark.sm$arr[`$k-1$`] =` n; 3. `(Watermark) sm$vec.set(`$k-1$`,` n`)`; 4. `Watermark.sm$n`k `=` n depending on how n is stored. k is n's node number.

Table 4.2: Translation from intermediate code instructions to Java.

Inserting the Java Code

`sandmark.watermark.CT.embed.Embedder` is the main class for modifying the Java program to be watermarked.

The chosen `mark()`-locations are replaced by calls to `Watermark.Create_`G_i. (We're relying on the `BLOAT` optimizer to eventually inline these calls and remove the `Watermark` class.) Remaining `mark()`-calls are deleted.

As usual, this process is done in several passes over the code:

1. `sandmark.watermark.CT.embed.ReplaceMarkCalls` replaces the `mark()`-calls with calls to `Watermark.Create_`G_i. There are two cases, depending on whether the `mark()`-call is `LOCATION`-based or `VALUE`-based. A `LOCATION`-based `mark()`-call is simply replaced by a call

```
1       Watermark.Create_Gi();
```

   A `VALUE`-based `mark(expr)`-call is replaced by the call

```
1       if (expr==value)
2          Watermark.Create_Gi();
```

2. `sandmark.watermark.CT.embed.AddParameters` adds formal parameters in order to be able to pass graph root nodes in formals rather than in globals. See Section 4.6 for more details.

3. `sandmark.watermark.CT.embed.InsertStorageCreators` inserts code to create hashtables, arrays, vectors, etc. that are used to store subgraph root nodes.

4. `sandmark.watermark.CT.embed.DeleteMarkCalls` removes any traces of the `mark()`-calls from the application.

4.5 Recognition

`sandmark.watermark.CT.recognize.Recognizer` starts up the watermarked application as a subprocess under debugging, again using Java's JDI debugging framework. The user enters their secret input sequence $I_0, I_i, \ldots$ exactly as they did during the tracing phase. This causes the methods `Watermark.Create_`G_i to be executed and the watermark graph to be constructed on the heap. When the last input has been entered it is the recognizer's task to locate the graph on the heap, decode it, and present the watermark value to the user.

There may, of course, be an enormous number of objects on the heap and it would be impossible to examine them all. To cut down the search space we rely on the observation that the root node of the watermark graph will be one of the very last objects to be added to the heap. Hence, a good strategy would likely be to examine the heap objects in reverse allocation order. Unfortunately, `JDI` does not yet provide support for examining the heap in this way.

An elegant and efficient approach would be to modify the constructor for `java.lang.Object` to include a counter:

```
1  package java.lang;
2  public class Object {
3      public static long objCount = 0;
4      public long allocTime;
5      public Object() {allocTime = objCount++;}
6  }
```

Since every constructor must call `java.lang.Object.<init>` this means that we've assigned an allocation order to the objects on the heap at the cost of only an extra add and assign per allocation.

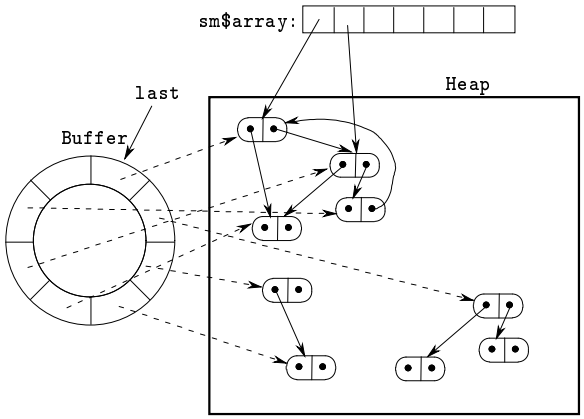

Figure 4.8: A view of memory during recognition. A circular linked buffer holds the last allocated objects. The recognizer examines the objects in reverse allocation order and extracts the subgraph reachable from each object. This is decoded into the watermark.

We've shied away from this approach, however, since it would require modifying the Java runtime system. Also, some Java compilers optimize away calls to `java.lang.Object.<init>` under the assumption that this constructor does nothing.

Instead, we rely on a more heavyweight but portable solution. Using `JDI` we add a breakpoint to every constructor in the program. Whenever an allocation occurs we add a pointer to the new object to a circular linked buffer, `sandmark.util.CircularBuffer`. This way, we always have the last 1000 (say) allocated objects available. The downside is a fairly substantial slowdown due to the overhead incurred by handling the breakpoints.

The recognition algorithm is as follows (see `sandmark.watermark.CT.recognize.Recognizer` and Figure 4.8):

```
1     static int kidMaps[][] = {{1,2},{2,1},{1,3},{3,1},{2,3},{3,1}};
2     for every object O on sandmark.util.CircularBuffer, starting with the last allocated do {
3         G := the graph consisting of the nodes reachable from O;
4         for every graph decoder D do {
5            for every kidmap K do {
6               W := decode G using D, assuming K;
7               print W;
8            }
9         }
10    }
```

The kidmaps are used to select 2 pointers out of each object as our outgoing pointers. The decoding is done by the codecs in `sandmark.util.graph.codec`. The class `sandmark.watermark.CT.recognize.Heap2Graph` is used to convert a heap structure to a `sandmark.util.graph.Graph` object. Note that the decoding can fail for a number of reasons. Some of the objects on the circular buffer may no longer be alive, the graph structure extracted from the heap may not be of the form expected by any of the codecs, etc. For this reason, we catch all possible exceptions (such as `null`-dereferences) and ignore any structures for which errors occur.

4.6 Passing Roots in Formal Parameters

In the descriptions above we have assumed that the roots of subgraphs are stored in static, global variables. In Figure 4.8, for example, the root of each subgraph is stored in a global array `sm$array`. This is obviously un-stealthy since programs typically only contain a few globals. Instead, we would like to pass roots in the

formal parameters of methods. This means we are going to have to find paths through the call-graphs from one `mark()`-call to the next.

More details about the call graph etc. here.

Chapter 5

The ConstantString Static Watermarking Algorithm

C. Collberg

5.1 Introduction

This is a trivial algorithm that embeds a string

sm$watermark=*WATERMARK*

in the constant pool of the application. *WATERMARK* is the string to be embedded. The code of this algorithm resides in `sandmark.watermark.constantstring`.

5.2 Embedding

To embed the watermark we pick one of the user's classes at random and adds the appropriate string:

```
1       String jarInput  = props.getProperty("WM_Embed_JarInput");
2       String jarOutput = props.getProperty("WM_Embed_JarOutput");
3       String watermark = props.getProperty("WM_Encode_Watermark");
4
5       sandmark.util.ClassFileCollection cfc =
6          new sandmark.util.ClassFileCollection(jarInput);
7
8       java.util.Iterator classes = cfc.classes();
9       String className = (String) classes.next();
10      de.fub.bytecode.classfile.JavaClass origClass = cfc.getClass(className);
11      de.fub.bytecode.generic.ClassGen cg = new de.fub.bytecode.generic.ClassGen(origClass);
12
13      de.fub.bytecode.generic.ConstantPoolGen cp = cg.getConstantPool();
14      int stringIndex = cp.addString("sm$watermark" + "=" + watermark);
15
16      de.fub.bytecode.classfile.JavaClass newClass = cg.getJavaClass();
17      cfc.addClass(newClass);
18      cfc.saveJar(jarOutput);
```

5.3 Recognition

During recognition we go through every class in the watermarked jar-file looking for a string in the constant pool that starts with `"sm$watermark"`:

```
1       String jarInput  = props.getProperty("WM_Recognize_JarInput");
2       sandmark.util.ClassFileCollection cfc =
3           new sandmark.util.ClassFileCollection(jarInput);
4       java.util.Iterator classes = cfc.classes();
5       while (classes.hasNext()) {
6          String className = (String) classes.next();
7          de.fub.bytecode.classfile.JavaClass clazz = cfc.getClass(className);
8          de.fub.bytecode.classfile.ConstantPool cp = clazz.getConstantPool();
9          for (int i=0; i<cp.getLength(); i++) {
```

```
10            de.fub.bytecode.classfile.Constant c = cp.getConstant(i);
11            if (c instanceof de.fub.bytecode.classfile.ConstantString) {
12                de.fub.bytecode.classfile.ConstantString s =
13                   (de.fub.bytecode.classfile.ConstantString) c;
14                String v = (String)s.getConstantValue(cp);
15                if (v.startsWith("sm$watermark")) {
16                   String w = v.substring("sm$watermark".length()+1);
17                   // w is the watermark
18                }
19            }
20         }
21      }
22      cfc.close();
```

Chapter 6

The RenameMethodsFields Watermarking Algorithm

Martin Stepp and Kelly Heffner

6.1 Introduction

This is an algorithm that embeds a watermark into a Java program by modifying the names of different methods within the program. The code of this algorithm resides in `sandmark.watermark.RenameMethodsFields`.

6.2 Embedding

To embed the watermark we first generate a list of all of the method equivalence classes. Two methods are in the same equivalence class if one method overrides the other. This list does not contain any special methods (overriden from Java, constructors, static initializers, etc.) From there, each letter of the watermark is appended on to a method name using a random number generator (with a particular seed which remains constant between embedding and recognition) to pick which method to use next. In the event that one method is used twice in order to embed the watermark, the letters of the watermark will be added in a LIFO (Last In First Out) order. A delimeter $ is used to mark the end of the watermark.

6.3 Recognition

During recognition we generate the same list of equivalence classes as in the embedding process, and use the same seed for the random number generator in order to rescan the methods, and reconstruct the watermark.

Chapter 7

Project VIII: Class-Splitting Obfuscations

Team : Ashok Purushotham and RathnaPrabhu Rajendran

7.1 Mission Document

For our final project we would like to implement a class-splitting obfuscator for Java .

7.2 Basic Idea

The basic idea behind the project is that there are two important object splitting techniques. The first splits each object into two, such that some fields belong to the first subpart, the remainder to the second subpart.

7.3 How we plan to implement

In the bytecode every "new"-operation turns into two "new"s and every pointer operation (getfield and setfield) turns into two pointer operations. The second splitting technique splits at the class level. A class C is broken into two classes C 1 and C 2 , such that C 2 inherits from C 1 . C 1 has fields and methods that only refer to themselves, whereas C 2 has fields and methods that can refer to themselves as well as fields and methods in C1 . Bytecode references to C will have to be replaced with references to C 2 .

The reason we want to implement this algorithm is that, to the best of our knowledge, it hasn't been implemented before. We would like to prove this by implementing and evaluating it. The implementation will be done within the SandMark framework. This will allow us to take advantage of the many excellent, well-documented libraries that are available within this tool.

7.4 Project Tasks

We believe this will be a difficult project with many potential pitfalls. First of all, the paper talks shallowly about our project and nothing else ,so we need to work out a reasonable algorithm. Second we can't simply split classes ,but splitting should be based on inheritance principles and understanding all of them and implementing them as a single task will be an uphill task .We therefore identify the following subtasks: 1. Understanding of concepts of splitting classes,going thourgh all relevant material in Sandmark and Collberg's paper. 2. Design 3. Discussing our design and implementation algorithms with Prof.Collberg and Coding 4. Testing and evaluation. 5. Documentation.

7.5 Reference

The reference for this project will be Collberg: Breaking Abstractions and Unstructuring Data Structures.

Chapter 8

The Bogus Switch Watermarking Algorithm

Andrew Huntwork and Xiangyu Zhang

8.1 Introduction

This is an algorithm that embeds a watermark in the 'case' statements of a switch block. The code of this algorithm resides in `sandmark.watermark.bogus_switch`.

8.2 Embedding

To embed the watermark we insert a lookup switch at the beginning of the instruction stream for a randomly chosen method:

8.3 Recognition

During recognition we find the switch we inserted, and decode its 'case' statements

Chapter 9

Boolean Variable Splitting

A. Huntwork and X. Zhang

9.1 Introduction

This algorithm detects boolean variables and arrays and modifies all uses and definitions of these variables. In particular, every variable or array element is split into 2 variables or array elements, and the state of the original variable is reflected in the combined state of these 2 variables or array elements. This algorithm consists of two components. One is boolean variable identification. We implement a dataflow analysis algorithm over the control flow graph which can identify the range of possible values of any variable. From the range the variable we can tell if it is a boolean variable (range between 0 and 1). The other is splitting the variables that are recognized. We implement 3 splitting algorithms, described below.

9.2 Boolean variable recognition

Observations

- Boolean variables in method fields have Z type
- Boolean variables in method parameters have Z type
- Boolean functions have Z return type
- Boolean local variables have implicit I type (which means iload is used to load the variable, istore is used to store)
- An assignment to a boolean variable in java source code generates the following java bytecode sequence:

```
ifeq l1
    iconst_0
    goto l2
    l1:
    iconst_1
    l2:
    istore boolean_var
```

- An boolean operation is implemented similarly. For example, b=a&&c generates the following code sequence:

```
iload a
    ifeq l1
    iload c
    ifeq l1
    iconst_1
    goto l2
    l1:
    iconst_0      l2:
    istore_1
```

- An boolean array variable has an explicit *newarray* with type *Boolean*

- The operations over boolean array elements are byte operations (baload, bastore)

From these observations, we can tell that it is easy to recognize a boolean field, method parameter, or array variable. The difficulty resides on the identification of boolean local variables.

Code pattern matching method

Based on our observations, we have developed a method based on code pattern matching which recognizes boolean variables by looking at the code sequences preceding stores to local variables. We compare these sequences with this pattern:

ifeq l1
iconst_0
goto l2
l1:
iconst_1
l2:
istore VAR

If all stores to a variable are preceded by such a pattern, VAR will be considered boolean variable. Both the Sun and IBM java compilers generate the same code pattern for boolean assignment. This pattern accounts for almost all assignments to boolean variables

But it does not work for the following cases:

b=a, whose code sequence is:

iload a
istore b b=f(), (f is a boolean function), whose code sequence is:
invokestatic ¡Method f¿
istore b

Even though this method discovers many boolean variables, it also misses many important cases. Therefore, we developed a better method.

Variable range based method

In this method, we do data flow analysis to identify the possbile value range of any variable, if it is [0,1] , we consider it as boolean variable.

BLOAT CFG code analysis

BLOAT does not have a dataflow analysis framework. However, it does have a control flow graph package which constructs the control flow graph and expression tree for a method. This CFG is used in its transformation, optimization and SSA packages.

The basic procedures of BLOAT CFG construction are:

1. Divide the code list into basic blocks (class Block extends Gnode). The division is simple. Every label starts a block, a block consists of the code sequence from one label to the next label. (This is in the buildBlocks method of class FlowGraph)

2. Starting from the first block (first label in the instruction list), BLOAT adds edges between blocks and build expression trees for block. The addition of edges and building of expression trees are in an depth first order. The depth first order means that if it finds a branch instruction, it calls itself with the target block as the parameter. Finally, BLOAT constructs a very complex graph. Blocks and the edges between them compose a conventional CFG. The CFG also consists of a forest of expression trees, the roots being statements, the edges representing the data dependence. A tree may include nodes from multiple basic blocks. (This is in the buildTreeForBlock method of class FlowGraph).

3. Add expression handler edges. Understanding this very complicated procedure is unnecessary for our purposes.

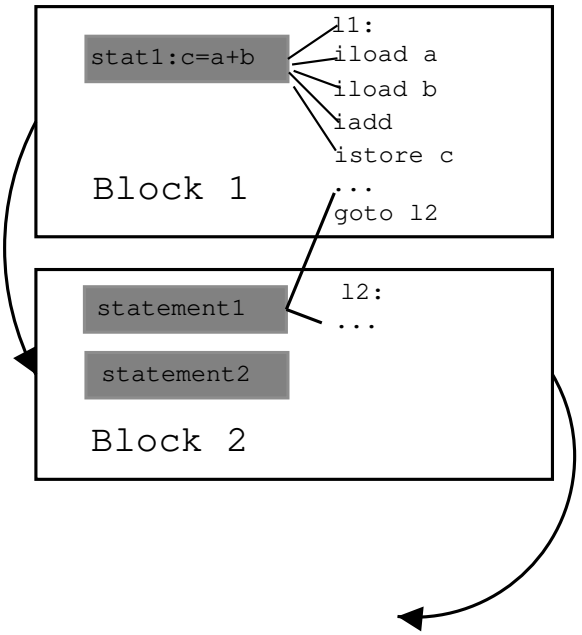

Expression tree construction is also a tedious procedure. BLOAT has a function to deal with the expression tree corresponding to **every instruction**. BLOAT also simulates the stack. For one instuction, it pops the operands from the stack, then pushes the expression of this instruction onto the stack and corresponding edges are added between this expression and the operand expressions. If this instruction does not represent an expression, nothing is pushed onto the stack. Instead, an statement is generated which is actually the root for the expression tree. (The code is in Tree class)

The CFG traversal is a DFS (anti-DFS) traversal over blocks. Inside a block, statements are traversed one by one.

Dataflow analysis

Given the complexity to build a pratical DF analyzer (the exception handler code in particular frightens the authors), we don't want to build a DF analyzer from scratch. Most of the components of a standard DF analyzer are already there. Blocks and edges between them compose a conventional CFG which can be used to do dataflow analysis. The functions which deal with expression construction for every instruction can be modified to do unit data flow procedure. The simulated stack can be used to flow the data. Our solution follows from these facts.

1. Divide the code list into blocks using the same method as CFG block division.
2. Add edges between blocks. This is also part of the CFG building function. At this point, we have a conventional CFG *dcfg*(without the confusing expression tree).
3. Generate a topological order for *dcfg*. We have to get the dominance to break a loop.
4. Traverse the blocks in a topological order. Store the range information in the stack. Store 0 for iconst_0, Store 1 for iconst_1, Store the current range of *var* for *iload var,* Store 0 for boolean function. etc. The meet operation of the data flow is simply the union of all the preceding stacks.
5. Finally, if a variable has [0,1] as its range, it is marked a boolean variable. We also mark boolean array, Boolean, Boolean array.

It turns out to be a nice solution to boolean local variable identification. It considers those integer variable with range [0,1] as boolean. But it doesn't matter.

9.3 Splitting Techniques

```
1
2   public class Test {
3       public static void main(String argv[]) {
4           int i;
5           boolean b,a[] = new boolean[16];
6           b = false;
7           for(i = 0 ; i < 16 ; i++) {
8               a[i] = b;
9               b = !b;
10          }
11          for(i = 0 ; i < 16 ; i++) {
12              System.out.println(a[i]);
13          }
14      }
15  }
16
```

The boolean variable splitting techniques presented below will be demonstrated on the sample code above.

XOR Splitting

The basic idea of the XOR splitting technique is as follows:

```
1           boolean a = true;
2           boolean b = false;
```

is converted to:

```
1           boolean a1 = true,a2 = false;
2           boolean b1 = false,b2 = false;
```

or

```
1           boolean a1 = false,a2 = true;
2           boolean b1 = true,b2 = true;
```

In general, a single boolean variable will be replaced with a pair of boolean variables such that the state of the original variable at any program point is equal to the XOR of the 2 new variables. Likewise, a boolean array 'a' will be replaced by an array 'b' with twice as many elements, in which the state of array element a[i] in the original array at any program point is equal to the XOR of b[2*i] and b[2*i + 1].

Parity Splitting

The basic idea of parity splitting is to replace a single boolean variable with a pair of integer variables that sum to an even number at every program point where the original variable is false, and sum to an odd number otherwise.

Equality Splitting

The basic idea of equality splitting is to replace a boolean variable with a pair of integer variables that have equal values at every program point where the original variable is true, and unequal values otherwise.

9.4 Code Transformation Techniques

Simple Variable Transformations

Our implementation only obfuscates boolean variables where all assignments to the variables are in one of the following forms:

```
1   iconst_0
2   istore_1
3
4   <or>
5
6   iconst_1
7   istore_1
```

or

```
1   <compute integer condition>
2   ifeq L1
3   iconst_1
4   goto L2
5   L1:
6   iconst_0
7   L2:
8   istore_1
```

These patterns are easily modified to work as split variables:

```
1   ifeq L1
2   iconst_0
3   iconst_1
4   goto L2
5   L1:
6   iconst_1
7   iconst_1
8   L2:
9   istore_1
10  istore_2
```

Boolean Array Splitting Techniques

Our implementation only obfuscates arrays that are allocated in a function and used in specific ways exclusively in that function. Allocation must look like this (all examples use the XOR obfuscation):

```
1   newarray boolean
2   astore_1
```

This pattern is converted to this pattern in the obfuscated code, which produces an array that is twice as long:

```
1   iconst_1
2   ishl
3   newarray boolean
4   astore_1
```

All stores must look approximately like this:

```
1   aload_1
2   iload_2
3   iload_2
4   ifeq L1
5   iconst_1
6   goto L2
7   L1:
8   iconst_0
9   L2:
10  bastore
```

This pattern is obfuscated as follows:

```
1   aload_1
2   iload_2
3   iconst_1
4   ishl
5   dup2
6   iload_3
7   ifeq L1
8   iconst_1
9   dup_x2
10  pop
11  iconst_0
12  goto L2
13  L1:
14  iconst_0
15  dup_x2
16  pop
17  iconst_0
18  L2:
19  bastore
20  bastore
```

All uses must look approximately like this:

```
1   aload_1
2   iload_2
3   baload
```

This pattern is obfuscated as follows:

```
1   aload_1
2   iload_2
3   iconst_1
4   ishl
5   dup
6   iconst_1
7   iadd
8   aload_1
9   swap
10  baload
11  dup_x2
12  pop
13  baload
14  ixor
```

Chapter 10

Bogus Initializers Watermarking Algorithm

Ashok Purushotham , RathnaPrabhu Rajendran

10.1 Introduction

This is an algorithm that embeds a number in the constant pool of the application. The code of this algorithm resides in `sandmark.watermark.constantstring`.

10.2 Embedding

To embed the watermark we pick one of the user's classes at random and add a few local variables to a method in the class. The initial values of the variables added, store the watermark.Suppose the number to be watermarked is 129091,we create three local variables sm$1, sm$2 and sm$3 and assign 12,90 and 91 to each of them respectively. If the number of digits in the number is an odd number, we append a zero to it, to make it even and split it among the variables appropriately. We also store the number of digits in the original watermark number, which is useful while recognizing the watermark.

```
1   int temp_len=watermark.length();
2   int stringIndex = cpg.addString("sm$len" + "=" + watermark.length();
3   if(temp_len%2 ==0)
4       bogus_ids_no=temp_len /2;
5   else
6     {
7       bogus_ids_no=temp_len /2 +1;
8       watermark=watermark+"0";
9       System.out.println(watermark);
10     }
11  de.fub.bytecode.classfile.Method[] methods=cg.getMethods()   ;
12  de.fub.bytecode.generic.MethodGen mg=
13     new de.fub.bytecode.generic.MethodGen(methods[methods.length-1],className,cpg);
14  int cp_index[] = new int[bogus_ids_no];
15  de.fub.bytecode.generic.InstructionList il = mg.getInstructionList();
16  for(int i=0;i<bogus_ids_no;i++){
17     String index=String.valueOf(i+1);
18     de.fub.bytecode.generic.LocalVariableGen lv=
19        mg.addLocalVariable("sm$"+index,de.fub.bytecode.generic.Type.INT,null,null);
20     cp_index[i]=lv.getIndex();
21     String str_value = watermark.substring((2*i),(2*i+1)+1);
22     int int_value=Integer.parseInt(str_value);
23     byte b = (byte)int_value;
24     il.insert( new de.fub.bytecode.generic.ISTORE(cp_index[i]));
25     il.insert( new de.fub.bytecode.generic.BIPUSH(b));
26     }
27
```

10.3 Recognition

During recognition, we go through every class in the watermarked jar-file looking for a string in the constant pool that starts with "sm$len" (This string holds the length the watermark added). Once we identify the string, we extract the values of the bogus variables added during embedding and combine them to reveal the watermark. We also take care of the extra '0' added to make the number of digits even, during the embedding process.

```
1   if (v.startsWith("sm$len")) {
2      String w = v.substring("sm$len".length()+1);
3      int no_of_vars=Integer.parseInt(w);
4      if(no_of_vars%2==0)
5          no_of_vars/=2;
6      else
7          no_of_vars=no_of_vars/2 +1;
8      String wm="";
9      for(int j=0;j<no_of_vars;j++){
10         de.fub.bytecode.generic.BIPUSH bip= (de.fub.bytecode.generic.BIPUSH)instr[2*j];
11         int sub_value= bip.getValue().intValue();
12         String str = String.valueOf(sub_value);
13         if( sub_value < 10)
14            str= "0"+str;
15         wm = str + wm;
16         }
17         if(wm.length()==Integer.parseInt(w))
18            result.add(wm);
19         else
20            result.add(wm.substring(0,wm.length()-1));
21
```

Chapter 11

The AddMethField Watermarking Algorithm

Miriam Miklofsky and Ginger Myles

11.1 Introduction

AddMethodField is a watermarking algorithm that embeds a watermark by splitting it into two parts. The first part is attached to a field name and the second is attached to a method name. Both the field and the method are new items inserted into the code. The code of this algorithm resides in `sandmark.watermark.addMethField`.

11.2 Embedding

To embed the watermark we split the watermark into two parts, wmPart1 and wmPart2. wmPart1 is combined with "sm$" to create the name of a new field that is added to the class. In order to embed wmPart2 we first randomly choose a method that already exists in the class. This method is chosen by seeding a random number generator with the key. Since the key is provided as a String we convert the String to a long. If a key is not specified then the seed has a default value. The randomly chosen method is used for two different purposes. The first purpose is for the naming of the new method. The name of the random method is used as the first part of the watermarked method's name, i.e. newMethodName = randomMethodName + "$" + wmPart2. The second purpose is that the chosen method will make a call to the watermarked method, so as to disguise the watermarked method. Since we have inserted a new method into the class, this method should do something. The watermarked method takes two parameters, whose values were also inserted into the chosen method, adds them together and places the result into the watermarked field. This assignment will aid in the recovery of the watermarked field.

```
1
2       //break the watermark into two parts
3       wmPart1 = first half of watermark
4       wmPart2 = second half of watermark
5
6
7       //get a pseudorandom number to indicate which method will make a call
8       //to the watermarked method.
9       //we are going to seed the random number generator with the key
10      if(key == empty string){
11         seed = 42;
12      }else{
13         seed = key converted to a long
14      }
15      methodIndex = randomGenerator(seed) % number of methods in class
16
17      //get the name of method where the watermarked method will be
18      //inserted. We are using this name to tack on the second part
19      //of the watermark.
20      String modMethodName = methods[methodIndex].getName();
21      String newMethodName = modMethodName + "$" + wmPart2;
22
23      create a new field
```

```
24
25      create a new method
26      //This new method takes the values of the two parameters, adds them
27      //together and places the result in the watermarked field.
28      //the instruction list for the method: load the value of the
29      //parameters, add, put the value in the watermarked field, return.
30      il.append(load1)
31      il.append(load2)
32      il.append(putfield(watermarkedField)
33      il.append(return)
34
35      add the new method so it becomes part of the class
36
37
38      //modify the randomly chosen method so that it makes a call to the new
39      //method. The instruction list that is to be inserted will be the
40      //following: store 5, store 10, invoke watermarked method.
41      il2.append(store5)
42      il2.append(store10)
43      il2.append(invoke watermarked method)
44      insert il2 into instruction list of chosen method
45
```

11.3 Recognition

During recognition the key is again used to seed a random number generator. This is used to determine which method contains the call to our watermarked method. The watermarked method can be distinguished from any other method calls because the begining of its name will be the same as the name of the method from which it is being called. Once the watermarked method is identified the instruction list of the watermarked method is searched to find a putfield instruction. The field that is being assigned to is the watermarked field. Once the two pieces of the watermark are identified they are tacked together to form the watermark that is returned.

```
1
2       //get a pseudorandom number to indicate which method will make a call
3       //to the watermarked method.
4       //we are going to seed the random number generator with the key
5       if(key == empty string){
6          seed = 42;
7       }else{
8          seed = key converted to a long
9       }
10      methodIndex = randomNumberGenerator(seed) % number of methods in class
11
12      String methodName = methods[methodIndex].getName()
13      //extract the second half of the watermark from the added method
14      //the watermarked method will have the same beginning as methodName
15      newMethodIndex = 0;
16      for(each method in the class) {
17         if(k != methodIndex) {
18            methodName2 = methods[k].getName();
19            if(methodName2.startsWith(methodName + "$" )) {
20               wmprt2 = methodName2.substring((methodName.length()+1), methodName2.length())
```

```
21              newMethodIndex = k;
22            }
23          }
24       }
25
26       //search the watermarked method's instruction list for putfield
27       if(instruction instanceof PUTFIELD){
28          fieldName = get field name
29          wmprt1 = fieldName - ``sm$''
30       }
31
32       //put pieces of watermark together
33       watermark = wmprt1 + wmprt2;
34
```

Chapter 12

The HatTrick Watermarking Algorithm

Danny Mandel and Anna Segurson

12.1 Introduction

This algorithm embeds a static watermark in a java *.jar file through the sandmark interface. It first converts the watermark (which is a String) to a number. This number is then encoded into the java byte code by adding bogus local variables to one method of one class in the *.jar file. The class in which the method is in is marked by a extra field variable that contains the name of the method.

Each number of the watermark is mapped to a specific type:

Number	Type
0	`java.util.GregorianCalendar`
1	`java.lang.Thread`
2	`java.util.Vector`
3	`java.util.Stack`
4	`java.util.Date`
5	`java.io.InputStream`
6	`java.io.ObjectInputStream`
7	`java.lang.Math`
8	`java.io.OutputStream`
9	`java.lang.String`

A local variable of the proper type is created for each digit of the watermark and the digit's location in the watermark is stored in the name of the bogus local variable. The rightmost digit is position 0.

The code of this algorithm resides in `sandmark.watermark.hattrick`.

12.2 Embedding

The watermark is embedded in the first non-abstract method found in the first class file of the jar file. The name of the method is recorded in the name of a field with the value of a hockey player's name. These are the local variables added to a class when the watermark "Howdy!" was embedded:

```
1       Math yzerman$0 = new Math();
2       String yzerman$1 = new String();
3       ObjectInputStream yzerman$2 = new ObjectInputStream();
4       Stack yzerman$3 = new Stack();
5       InputStream yzerman$4 = new InputStream();
6       Stack yzerman$5 = new Stack();
7       OutputStream yzerman$6 = new OutputStream();
8       InputStream yzerman$7 = new InputStream();
9       InputStream yzerman$8 = new InputStream();
10      OutputStream yzerman$9 = new OutputStream();
11      Thread yzerman$10 = new Thread();
12      Thread yzerman$11 = new Thread();
13      Thread yzerman$12 = new Thread();
14      ObjectInputStream yzerman$13 = new ObjectInputStream();
15      Stack yzerman$14 = new Stack();
```

This is the field that is added: `public static final String hat<init>Trick = ``Vrbata'';` where `<init>` is the name of the method the watermark is in.

12.3 Recognition

The watermark is recognized when the proper field is present which stores the name of the method. Every local variable that contains the `SECRET_NAME` in that method is part of the watermark.

In the example listed above `yzerman` is the secret name. To recover the number stored by the `SECRET_NAME` variables, the algorithm initializes a number, `wm`, to zero then loops through all of the locals in the method and adds the mapped value of the variable times 10 to the power of the number after the $ to `wm`. After the number is recovered it is converted back to the watermark string.

MethodRearranger is a static obfuscation algorithm that moves all of the static methods that it can into their own class. Only static methods that do not rely on private information within its residing class can be moved.

The MethodRearranger works as follows:

1. A list of movable methods is collected, and those methods are moved into MasterStaticClass.
2. Constant pool references within each moved method are updated.
3. Constant pool references for the method calls for those methods are updated to reference the new class.

Chapter 13

The AddBogusFields Obfuscation Algorithm

Miriam Miklofsky and Ginger Myles

13.1 Introduction

AddBogusFields is an Obfuscation algorithm that will work on a class file. It inserts assignments to a bogus field in specific places throughout the code. These locations are identified by choosing a "sister" field. The code of this algorithm resides in `sandmark.obfuscate.addbogusfields`.

13.2 Apply

To obfuscate the code a new field in added to the class and then assignments are made to this field. The way the obfuscation works is that a "sister" field is identified which is used to insert an assignment to the new field. The sister field is chosen randomly. Everytime there is an assignment to the sister field an assignment to the new field is inserted into the code.

```
1   newFieldName = ``sm$'' + i
2   sisterFieldIndex = randomNumberGen() % fields.length
3   insert newFieldName in the class
4   for(each method in class){
5      search instruction list for assignment to sisterField
6      insert assignment to newField
7   }
```

Chapter 14

The PromoteLocals Obfuscating Algorithm

Danny Mandel and Anna Segurson

14.1 Introduction

This is an algorithm that changes all of the local `int` variables in a method to local `java.lang.Integer` variables.

For this to work every reference to the old local `int` has to be replaced by the appropriate reference to the new local `java.lang.Integer`. We identified three java byte code instructions that needed to be modified to make this change. The following table lists the necessary modifications:

Old Byte Code	Replacement Byte Code

`istore <index>`	`new Ljava/lang/Integer;`
	`dup_x1`
	`dup_x1`
	`pop`
	`invokespecial java/lang/Integer/<init>(I)V`
	`astore <index>`

`iload <index>`	`aload <index>` `invokevirtual java/lang/Integer/intValue(V)I`

`iinc <index><num >`	`new Ljava/lang/Integer;`
	`dup`
	`ldc <num>`
	`aload <index>`
	`invokevirtual java/lang/Integer/intValue(V)I`
	`iadd`
	`invokespecial java/lang/Integer/<init>(I)V`
	`astore <index>`

The code of this algorithm resides in `sandmark.watermark.promotelocals`.

14.2 Apply

To apply this obfuscation we circulate through the byte code of a given method and if the instruction is one of the ones in the table above we replace that instruction with the corresponding replacement byte code.

For example, this simple class:

```
1   public class VerySimp {
2      public static void main(String[] args) {
3         int a,b,c;
4         a = 1;
5         b = 22;
6         c = a + b;
7         for (int i = 0; i < 3; i++) {
8            a+=2;
9         }
10     }
11  }
```

is obfuscated to the following class:

```
1  public class VerySimp {
2     public static void main(String[] as)
3     {
4        Integer integer2 = new Integer( 1 );
5        Integer integer3 = new Integer( 22 );
6        Integer integer4 = new Integer(integer2.intValue() + integer3.intValue());
7        Integer integer1 = null;
8        for( integer1 = new Integer( 0 ); integer1.intValue() < 3;
9             integer1 = new Integer( 1 + integer1.intValue() ) )
10          integer2 = new Integer( 2 + integer2.intValue() );
11    }
12 }
13
```

Chapter 15

Bogus Predicates Obfuscating Algorithm

Ashok Purushotham , RathnaPrabhu Rajendran

15.1 Introduction

This is an algorithm that implements simple boolean identities and add them to the user's code . The code of this algorithm resides in `sandmark.obfuscate.BogusPredicates`.

15.2 Apply

Our aim is to embed opaquely true constructs which must be stealthy . So we have selected some constructs based on algebraic properties and known facts in mathematics .For example ,we know for all x,y in I, (7(y X y) - 1) is not equal to (x X x) . A list of all the available constructs is maintained .At run time ,whenever we encounter a conditional expression ,we randomly select one among these to append to the current expression .Since the expression that we add is opaquely true,but the reverse engineer has to try out many inputs to find that this added expression is indeed opaquely true,if he isn't aware of mathematical properties . For example ,the original method and its transformation on applying bogusPredicates Algorithm:

```
3   main(){
4    int a=10;
5    int b=20;
6    if(a<30)
7      b=a+99;
8
9   }
10
11  main(){
12   int a=10;
13   int b=20;
14   int c;
15   if(a<30 && c(c+1)%2 ==0)
16     b=a+99;
17  }
18
```

The original byte code in a conditional expression was

```
1     iload_1
2     bipush 7
3     if_icmpne 29
```

The new added byte code as a result of our algorithm is

```
2     iload_1
3     bipush 7
4     if_icmpne 29
5     iload_3
6     dup
7     iconst_1
8     iadd
```

```
 9      imul
10      iconst_2
11      irem
12      iconst_0
13      if_icmpne 29
```

Chapter 16

The Method 2R Madness Obfuscation Algorithm

Martin Stepp and Kelly Heffner

16.1 Introduction

The Method 2R Madness Obfuscation (M2M) algorithm is based on the idea that useful information about program organization lies in the parameter list and separation of methods. In order to destroy that information, the M2M obfuscator performs a six pass obfuscation:

1. In pass one, all fields and methods are made public.
2. In pass two, all dynamic method bodies are moved into a static helper method. The dynamic shell is simply a call on this static helper.
3. In pass three, all primitive types residing in a static method signature are promoted to their respective object wrapper types.
4. In pass four, the parameter lists of all static methods are rearranged.
5. In pass five, the parameter lists of all static methods are condensed into one large `Object []` array.
6. In pass six, static methods in the same class with the same signature are merged into one method.

16.2 Pass Two - Method Splitting

As mentioned above, the second phase of the M2M obfuscation splits each dynamic method into two parts. The first part is the same dynamic method, except that the body of the method is changed to a simple call on a static helper method. The first parameter to this static method is the actual `this` reference which was implicitly argument zero in the dynamic method. Take for example the following class to computer Ackermann's function:

```
public class Ackermann
{
    public static void main(String[] args)
    {
      int num = Integer.parseInt(args[0]);
      System.out.println("Ack(3," + num + "): " + new Ackermann().Ack(3, num));
      }

    public int Ack(int m, int n)
    {
      return (m == 0) ? (n + 1) : ((n == 0) ? Ack(m-1, 1) :
                  Ack(m-1, Ack(m, n - 1)));
    }
}
```

After being split into a static section by the StaticSplit obfuscator phase, and then run through the SourceAgain decompiler, the code looks as follows:

```
//
// SourceAgain (TM) v1.10i (C) 2001 Ahpah Software Inc
//

import java.io.PrintStream;

public class Ackermann {

    public static void main(String[] as)
    {
        int i = Integer.parseInt( as[0] );

        System.out.println( "Ack(3," + i + "): " + (new Ackermann().Ack( 3, i )) );
    }

    public static int Ackintint0(Ackermann this, int arg0, int arg1)
    {
        return (arg0 == 0) ? arg1 + 1 : (arg1 == 0) ? Ack( arg0 - 1, 1 ) : Ack( arg0 - 1, Ack( arg0, ar
    }

    public int Ack(int arg0, int arg1)
    {
        return Ackintint0( this, arg0, arg1 );
    }
}
```

16.3 Pass Three - Primitive Promoting

The third phase of the M2M obfuscation consists of promoting the argument list and return type of the static methods so that all primitive typed arguments or returns are now reference types. A restriction of reference only parameters and returns allows for easier parameter reordering and array parameterize, and this primitive promoting obfuscation makes it possible for more static methods to fit the "reference only" criteria for reordering. We will again use the Ackermann example from above, and tack on the primitive promoting pass.

```
//
// SourceAgain (TM) v1.10i (C) 2001 Ahpah Software Inc
//

import java.io.PrintStream;

public class Ackermann {

    public static void main(String[] arg0)
    {
        int i = Integer.parseInt( arg0[0] );

        System.out.println( "Ack(3," + i + "): " + (new Ackermann().Ack( 3, i )) );
    }

    public static Integer get0(Ackermann this, Integer arg0, Integer arg1)
    {
        return new Integer( (arg0.intValue() == 0) ? arg1.intValue() + 1 : (arg1.intValue() == 0) ? Ack
```

```
    }

    public int Ack(int arg0, int arg1)
    {
        return get0( (Ackermann) new Object[3][0], (Integer) new Object[3][1], (Integer) new Object[3][
    }
}
```

Notice that the decompiler has a problem interpreting that the arguments to method calls were promoted using a temporary array of Objects.

16.4 Pass Four - Parameter Reordering

The parameter reorderer adds an extra dose of confusion to the algorithm, so that there is no direct mapping between the order of the nonstatic stub methods, and their static helpers. This also shuffles the `this` reference out of slot zero, making our static helpers less obvious.

```
//
// SourceAgain (TM) v1.10i (C) 2001 Ahpah Software Inc
//

import java.io.PrintStream;

public class Ackermann {

    public static void main(String[] arg0)
    {
        int i = Integer.parseInt( arg0[0] );

        System.out.println( "Ack(3," + i + "): " + (new Ackermann().Ack( 3, i )) );
    }

    public static Integer get0(Integer this, Ackermann arg0, Integer arg1)
    {
        return new Integer( (arg1.intValue() == 0) ? intValue() + 1 : (intValue() == 0) ? arg0.Ack( arg
    }

    public int Ack(int arg0, int arg1)
    {
        return get0( (Integer) new Object[3][0], (Ackermann) new Object[3][1], (Integer) new Object[3][
    }
}
```

16.5 Pass Five - Signature Bludgeoning

The bludgeon takes the final step in confusing the signatures of methods, by making all of the static method signatures the *same*, "(L[java/lang/Object;)Ljava/lang/Object;". This prepares the classes for the final pass, which is to combine any methods with the same signature (which is now all static methods).

```
//
// SourceAgain (TM) v1.10i (C) 2001 Ahpah Software Inc
//
```

```
import java.io.PrintStream;

public class Ackermann {

    public static void main(String[] arg0)
    {
        int i = Integer.parseInt( arg0[0] );

        System.out.println( "Ack(3," + i + "): " + (new Ackermann().Ack( 3, i )) );
    }

    public static Object get0(Object[] this)
    {
        return new Integer( (((Integer) this[2]).intValue() == 0) ? ((Integer) this[0]).intValue() + 1
    }

    public int Ack(int arg0, int arg1)
    {
        return ((Integer) get0( new Object[] { (Integer) new Object[3][0], (Ackermann) new Object[3][1]
    }
}
```

16.6 Pass Six - Method Merging

The method merger combines static methods with the same signature into one method, that is multiplexed on an extra `int` parameter. Consider the code for the original Ackermann, with an extra static method added in:

```
public class Ackermann
{
    public static void main(String[] args)
    {
      int num = Integer.parseInt(args[0]);
      System.out.println("Ack(3," + num + "): " + new Ackermann().Ack(3, num));
      }

public static void hi(long p)
{
System.out.println(p);
}

    public int Ack(int m, int n)
    {
      return (m == 0) ? (n + 1) : ((n == 0) ? Ack(m-1, 1) :
                  Ack(m-1, Ack(m, n - 1)));
    }
}
```

After a pass through all five other phases, the two static methods (which now share an Object[] to Object signature) are merged together as follows:

```
//
// SourceAgain (TM) v1.10i (C) 2001 Ahpah Software Inc
```

```
//

import java.io.PrintStream;

public class Ackermann {

    public static void main(String[] arg0)
    {
        int i = Integer.parseInt( arg0[0] );

        System.out.println( "Ack(3," + i + "): " + (new Ackermann().Ack( 3, i )) );
    }

    public int Ack(int arg0, int arg1)
    {
        return ((Integer) keldaANDmartio0( new Object[] { (Integer) new Object[3][0], (Integer) new Obj
    }

    public static Object keldaANDmartio0(Object[] keldaANDmartio0, int switchVar)
    {
        if( 1 == switchVar )
            return new Integer( (((Integer) keldaANDmartio0[1]).intValue() == 0) ? ((Integer) keldaANDm
        if( 0 != switchVar )
            ;
        System.out.println( ((Long) keldaANDmartio0[0]).longValue() );
        return null;
    }
}
```

Chapter 17
Class-Splitting and False Refactoring Obfuscation Algorithms

Ashok and Rathna Prabhu

17.1 Introduction

According to the Chidamber Metric [1], the complexity of a class grows with its depth in the inheritance hierarchy and the number of its direct descendents . This is an implementation of the class-splitting/False Refactoring obfuscation proposed in [2]. In the class-splitting obfuscation, we analyze the dependences between the fields and methods of a given class. Using this dependence information, the class is split into different classes forming an inheritance hierarchy, thus obfuscating the given class. In the second algorithm, we do refactoring of two unrelated classes i.e we identify some common fields in the two classes and move them into a bogus super class. Now these two classes appear to inherit from a common class. But in reality, the two classes are unrelated. The algorithms are described in the following sections with examples.

17.2 Class Splitting Obfuscation

A class C is broken into two classes C1 and C2 , such that C2 inherits from C1 . C1 has fields and methods that only refer to themselves, whereas C2 has fields and methods that can refer to themselves as well as fields and methods in C1 . Bytecode references to C will have to be replaced with references to C2. Consider the following simple code sample C.java, to which we'll apply the obfuscation technique:

```
1    class C {
2     float f;
3     int v;
4     public C()  {  }
5     public void P()
6     {
7      f=1.2;
8     }
9     public void Q()
10    {
11       v = 1;
12       P();
13    }
14   }
```

Figure 17.1 shows the different phases involved in the algorithm.

Dependency Graph Generation

We create a dependency graph G for class C .The nodes of G are the members of C. There is an edge from A to B in G, if the declaration of A must be in scope of B.If there is a path from C to Y in G , then Y must be declared in the child class during splitting. If there is a path from X to Y in G , then either X and Y are both declared in the same class or X is declared in the parent class C1. Figure 17.2 shows the dependency graph for the given code sample.

Topological sort of the Dependency Graph

We perform a topological sort of the dependency graph to know what methods and variables are to be put in which classes . For the dependency graph created, topological sort gives the following results:

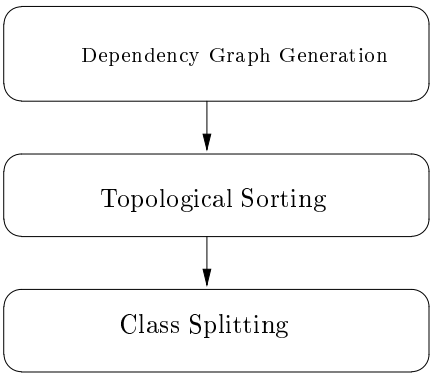

Figure 17.1: Phases in the Class-Splitting algorithm

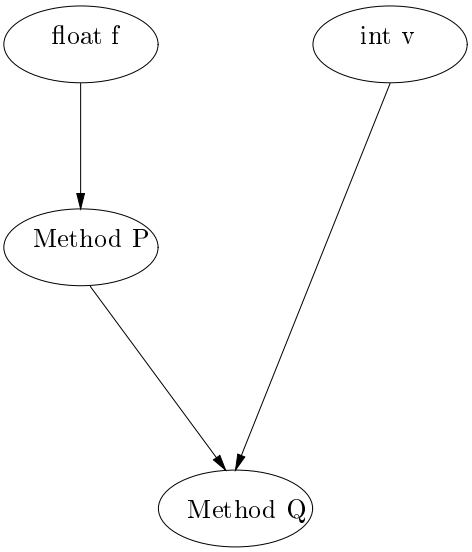

Figure 17.2: Dependency Graph

```
1   Level 0 : Field f
2   Level 1 : Method P
3   Level 2 : Method Q
```

From the above results we can decide on the placement of the methods and fields in the inheritance tree. For instance, the method P must be in the scope of the class in which Q is defined. So Q can be put in the same class as P or in some class inheriting from the class defining P. This ensures that the obfuscation is semantics preserving and gives the same results as before.

Results

The result of applying the algorithm to C.java is as follows :

```
1   class C0 {
2           float f;
3           int v;
4           public C0() { }
5    }
6   class C1 extends C0 {
7           public C1() { }
8           public void P() {
9            f=1.2
10           }
11   }
12  class C extends C1 {
13          public C() { }
14          public void Q() {
15            v=1;
16            P();
17          }
18   }
```

Note that the sub class has been named as C. This makes sure that the references to the original class C, in other java programs are still valid.

17.3 False Re-factoring of Classes

False refactoring is performed on two classes C1 and C2 that have no common behavior. If both classes have instance variables of the same type, these can be moved into the new parent class C3 . C3 's methods can be buggy versions of some of the methods from C1 and C2 . Consider the following sample code

```
1   C1.java
2   class C1 {
3           String customerName="";
4           public void process() {
5            //Process customer
6            customerName=...
7           }
8    }
9
10  C2.java
11  class C2 {
12          String furnitureName="";
13          public int polish(){
14           furnitureName=...
```

```
15            //Polish the furniture
16           }
17    }
```

After false refactoring the classes:

```
1   class SMBaseC1 {
2           String smstring=""
3           public void process() {
4             //Buggy code
5           }
6   }
7
8   class C1 extends SMBaseC1 {
9           public void process() {
10            smstring=...
11          }
12  }
13
14  class C1 extends SMBaseC1 {
15          public int  polish() {
16           smstring = ...
17
18          }
19  }
```

If one of the classes inherit from some other class other than java.lang.Object, then SMBaseC1 has to be an interface and the two classes have to implement SMBaseC1 instead of extending it.

17.4 Reference

1. Chidamber, S. and C. Kemerer, Towards a Metrics Suite for Object Oriented Design, Proceedings of OOPSLA, July 1991.

2. Breaking Abstractions and Unstructuring Data Structures - Christian Collberg, Clark Thomborson, Douglas Low

Chapter 18

The Degrading Obfuscation Algorithm

D. Mandel, A. Segurson

18.1 Introduction

The Degrading algorithm for obfuscation is designed to slow down java applications. It consists of two main parts, the promotion obfuscator, and the thread contention obfuscator. There is a motivation behind the desire to make java applications run slower: Software company X wants to release a new program. But, company X wants the users to be able to try the new software before they buy it, as a motivation to own it. If the customer already has a working version of the software, why would they still buy it? Because company X only released a slightly crippled version of the software and the customer will enjoy the optimal speed of the full version. Thus, company X will degrade their product and release it as a free demo, and sell the fastest working version.

What makes java programs run slower? Memory managment, and thread contention are two things in java code that we focus on that can slow things down. Our goal is to increse the amount of these two things in the java application with some control.

The following is the discussion of the two algorithms used in the Degrading Obfuscation, followed by some analysis of programs in which the degradation was applied.

18.2 Promotion Obfuscation

The main goal of the promotion obfuscator is to increase the amount of memory management in the java application. Almost everything in java is already an object, so the target was clear: primitive local variables. Primitive local variables are used quite frequently in programming, which in java includes the types: " `boolean`, `byte`, `char`, `int`, `double`, `float`, `long`, and `short`." Luckily, java already has provided us with objects that can hold these values.

The details on how to promotote the primitive local variables to an object local variables vary slightly from primitive to primitive. The basic idea follows these steps:

- cast all primitive variable parameters passed into the method to local objects
- when a primitive variable is loaded, replace the load with a call to the local object's value function
- when a primitive variable is stored, replace the store with a constructor for a new object and store the new object

Special `int`'s

The primitive `int` is probably the most common used primitive for a local variable. Also, on the java byte code level the primitives `boolean`, `byte`, `char`, and `short` are treated exactly the same as `int` variables. See the figure for an example. So, there is no support for `boolean`, `byte`, `char`, and `short`'s on the byte code level, they are all `int`'s.

As a result of java's treatment of the small primitives (`boolean`, `byte`, `char`, and `short`), the small primitives can all be stored in the object, `Integer`. There are objects for all of the small primitives, but detailed analysis of the program would be required to use the same object as small primitive, and there would be no beneficial results. For example, one could do data analysis on the local variable 2, and if it is an integer which is always 0 or 1 then it could be stored in a `Boolean` object, but saving it in the `Boolean` object offers no advantage over saving it in a `Integer` object. For the rest of the section `int`'s will include all small primitives.

```
1      public static char myMethod(short b, int c, boolean bb) {
2         if (bb) {
3            int x = b + c;
4            return 'a';
5         }
6         return 'b';
7      }
```

```
1   .method public static myMethod(SIZ)C
2   .limit stack 2
3   .limit locals 4
4
5   Label1:
6   .line 7
7           iload_2
8           ifeq Label0
9   .line 8
10          iload_0
11          iload_1
12          iadd
13          istore_3
14  .line 9
15          bipush 97
16          ireturn
17  Label0:
18  .line 11
19          bipush 98
20  Label2:
21          ireturn
22
23  .end method
```

Figure 18.1: The listing of the java, and the java byte code for the method, `myMethod`. `myMethod`'s signature in the byte code accepts a `short`, `int` and a `boolean` as paramaters, but when the method uses these paramaters it always uses an `iload`. Also, when the method returns the `char` it uses an `ireturn`.

Promoting Primitives

There are only four main types of primitives: `int`, `float`, `long` and `double`. To implement the promotion obfuscator four method obfuscator algorithms were added to Sandmark:

- DPromote - `double` promote
- FPromote - `float` promote
- IPromote - `int` promote
- LPromote - `long` promote

All four method obfuscator algorithms begin by loading all of the primitive parameters into an object. Figure 18.2 is the new byte code created for each primitive paramter type at the beginning of the method.

Now, there are only local variables that are objects. But, we still have to change the rest of the method to correct the stores and loads to reference the objects, instead of the primitives. Figures 18.3 - 18.4 itemize all the substitutions that need to be made.

Long	**Float**
new java/lang/Long dup lload 4 invokespecial java/lang/Long/<init>(J)V astore 4	new java/lang/Float dup fload_2 invokespecial java/lang/Float/<init>(F)V astore_ 2
Int	**Double**
new java/lang/Integer dup iload_3 invokespecial java/lang/Integer/<init>(I)V astore_3	new java/lang/Double dup dload_0 invokespecial java/lang/Double/<init>(D)V astore_0

Figure 18.2: The code inserted at the beginning of the method for each of the four primitive types. In this example the signature of the static method was `myMethod(DFIJ)V`.

Old Byte Code	Replacement Byte Code
istore <index>	new Ljava/lang/Integer; dup_x1 dup_x1 pop invokespecial java/lang/Integer/<init>(I)V astore <index>
iload <index>	aload <index> invokevirtual java/lang/Integer/intValue(V)I
iinc <index><num >	new Ljava/lang/Integer; dup ldc <num> aload <index> invokevirtual java/lang/Integer/intValue(V)I iadd invokespecial java/lang/Integer/<init>(I)V astore <index>

Figure 18.3: Substitution for integers.

Controling Promotion

The number of local variables that are promoted can be controled to a limited extent. The following parameter can be set in the Degrade configuration:

- `OBF_DE_PERCENT` - The approximate percent of the primitive local variables that you want to promote.

The algormithm will then apply the promotion obfuscation to methods in the jar file until it has reached or exceeded the desired percent. It uses the greedy algorithm to achieve this.

Old Byte Code	Replacement Byte Code

`lstore <index>`	`new Ljava/lang/Long;`
	`dup_x2`
	`dup_x2`
	`pop`
	`invokespecial java/lang/Long/<init>(J)V`
	`astore <index>`

`lload <index>`	`aload <index>` `invokevirtual java/lang/Long/longValue(V)J`

Old Byte Code	Replacement Byte Code

`fstore <index>`	`new Ljava/lang/Float;`
	`dup_x1`
	`dup_x1`
	`pop`
	`invokespecial java/lang/Float/<init>(F)V`
	`astore <index>`

`fload <index>`	`aload <index>` `invokevirtual java/lang/Float/floatValue(V)F`

Old Byte Code	Replacement Byte Code

`dstore <index>`	`new Ljava/lang/Double;`
	`dup_x2`
	`dup_x2`
	`pop`
	`invokespecial java/lang/Double/<init>(D)V`
	`astore <index>`

`dload <index>`	`aload <index>` `invokevirtual java/lang/Double/doubleValue(V)D`

Figure 18.4: Substitution for longs, floats and doubles.

18.3 Thread Contention Obfuscation

This algorithm attempts to automagically slow down some method in a program by inserting threads that perform busy-waiting while the original code in the method runs. While the idea of slowing down a program may seem odd, commercial software developers may find this idea interesting in order to give away a slower version of their product while charging for the complete version. This algorithm may also be used as a benchmarking scheme for java's multithreaded features. In addition, the obfuscation value of this algorithm (when it is completely implemented with sufficient stealthiness) should not be underestimated. Multithreaded code has been known to confuse even advanced programmers, and obfuscated multithreaded code should further help in this quest for confusion.

Apply

For the algorithm to work, the user must have configured the following parameters beforehand:

- `OBF_TC_ClassName` - The name of the class where the method resides
- `OBF_TC_MethodName` - The name of the method to obfuscate
- `OBF_TC_NumThreads` - The number of threads to insert into that method

If a method with the given class and name isn't found, then the obfuscation does nothing. Once these values have been set, the user may run the algorithm. A summary of the steps taken during the obfuscation are as follows:

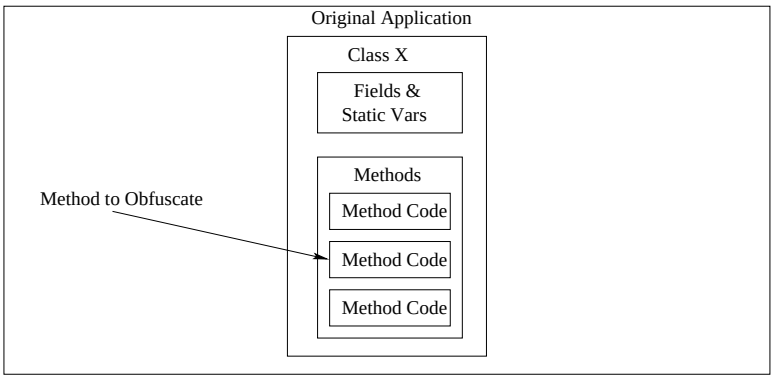

Figure 18.5: The original application to obfuscate. A method is selected for obfscation

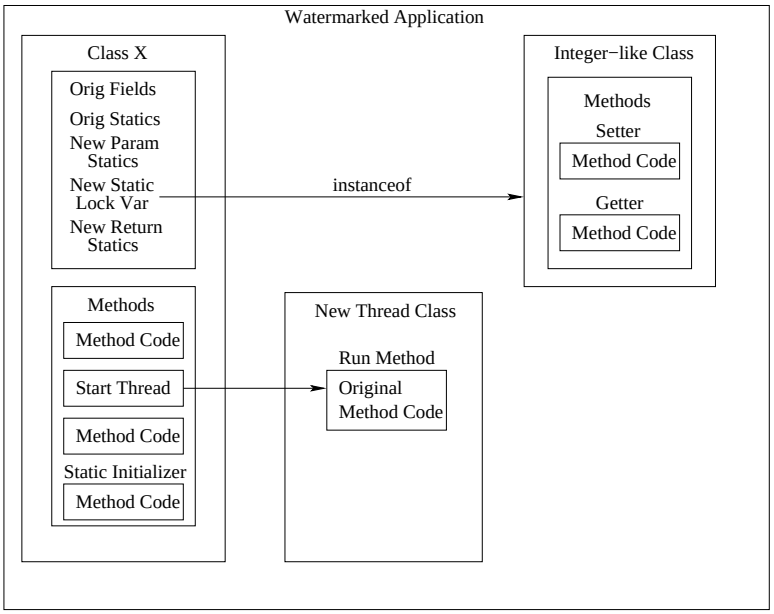

Figure 18.6: The application after obfuscation. The original method's code has been moved to a thread which the original method now starts. An Integer-like class has been added for locking purposes.

- A new `java.lang.Integer`-like class is added to the jar file
- A static initializer is either modified (if it already exists) or created (if it doesn't already exist) and a static lock variable is added to the original class
- A new inner-thread class that contains the original method code is added to the original class
- Some number of other inner-thread classes are added to the original class
- The original method is modified so it starts the new inner thread class and waits for that thread to complete

New java.lang.Integer Class

In order to support various locking methods, the first step in the obfuscation is to add a class similar to `java.lang.Integer`. The main difference between the new class and java's version is that the new class can change the value stored in the Integer object. This difference is significant because it allows a different way of coordinating threads. Java's built-in synchronization mechanism relies on monitors which function through object locks. For our purposes, we want to use this in addition to value-based synchronization mechanisms. The standard `java.lang.Integer` class cannot support value-based locking due to the value being immutable once the object has been created. If one wishes to change the value of an Integer, one must create a new Integer, which throws away the reference to the old Integer, which renders the java's built-in object locking mechanism useless.

Static Initializer and Lock Variable Modified or Added

For locking to take place, there needs to be some sort of locking variable. While typical locking in java uses the `synchronized` keyword in method headers, this synchronizes on the `this` object associated with that method, which defeats our purpose of having more than one thread attempting to run while our original method is running. To get around this, we introduce a new static class variable of our Integer-like class that the methods can synchronize on. A static initializer is added to the class in order to construct our new Integer variable.

New Inner-Thread Class

A new inner-thread class is created in order to run the old method code. The new class contains one method, `run()`, which is required in order to implement the `java.lang.Runnable` interface. The method body of the `run()` method contains the instructions that were formerly in the original method. However, there are two slight details that make this process difficult, and both of them stem from the signature of the `run()` method required by `java.lang.Runnable`. Since the signature of this method is set, it is impossible to pass parameters or return a value in the conventional manner. The workaround for this is to add static class variables to the original class so that both the original class and the new class can access the values. It should be noted that all generated code (with the exception of the first thread class which always has the name `<className>`$Inner) has random identifier strings associated with them. As well, all generated code is declared static so as to be accessible from both a static and dynamic context. For parameter passing, if the code to be edited originally looked like this:

```
1  public class ParamPassingExample {
2     public static void main(String[] args) {
3        for (int i = 0; i < args.length; i++) {
4           System.out.println(``Arg `` + i + `` = `` + args[i]);
5        }
6     }
7  }
```

Then we would add the new class and a static variable so the resulting code and modified `main(String[])` method would look like this:

```
1   public class ParamPassingExample {
2      // New static variable added to mimic parameter-passing
3      private static String[] d3345;
4
5      public static void main(String[] args) {
6         Thread newThread;
7         ParamPassingExample$Inner inner;
8
9         // Assign the parameters to the new static class variable
10        d3345 = args;
11
12        inner = new ParamPassingExample$Inner();
13        newThread = new Thread(inner);
14        newThread.start();
15     }
16
17     private static class ParamPassingExample$Inner implements java.lang.Runnable {
18        public void run() {
19           // Assign the static class variable to a local for manipulation like
20           // how a parameter would normally be manipulated
21           String[] local = d3345;
22
23           for (int i = 0; i < local.length; i++) {
24              System.out.println(``Arg `` + i + `` = `` + args[i]);
25           }
26        }
27     }
```

Return types are treated in a similar manner, in that a new static variable is added to the class. However, for return types to work correctly, the modified original method must spin on the lock until the thread has completed. To illustrate this, consider the following example:

```
1   public class ReturnTypeExample {
2
3      public static void main(String[] args) {
4         System.out.println(``Square root of 25 is: `` + squareRoot(25));
5      }
6
7      public static double squareRoot(int value) {
8         return Math.sqrt((double) value);
9      }
10  }
```

After adding the lock variable and modifying the `squareRoot(int)` method to construct the thread and spin until the thread has completed, we have the following code:

```
1   public class ReturnTypeExample {
2      // New static variable for param-passing
3      private static int d546;
4      // New static variable for return-type
5      private static double g6789;
6      // New static lock variable
7      private static ReturnTypeExample$n7632 k3452;
8
```

```
 9      // static initializer added to construct new lock variable
10      static {
11         k3452 = new ReturnTypeExample$n7632(0);
12      }
13
14      public static void main(String[] args) {
15         System.out.println(``Square root of 25 is: `` + squareRoot(25));
16      }
17
18      public static double squareRoot(int value) {
19         ReturnTypeExample$Inner inner;
20         Thread thread;
21
22         // Store the parameter in a static class variable
23         d546 = value;
24         // Then start the thread
25         inner = new ReturnTypeExample$Inner();
26         thread = new Thread(inner);
27         thread.start();
28         // Spin on the lock until the thread completes its work
29         while (k3452.j213() != 1) {
30         }
31         // Then return the newly added static variable
32         return g6789;
33      }
34
35      private class ReturnTypeExample$Inner implements java.lang.Runnable {
36         public void run() {
37            // Re-assign the value of the lock variable to 0 to delay execution
38            // of other threads
39            k3452.n4795(0);
40            // Assign the parameters to a local variable
41            int local = d546;
42            // Then assign the original return value to our static class member
43            g6789 = Math.sqrt(d546);
44            // Then assign the lock value to 1 so other threads can complete
45            k3452.n4795(1);
46         }
47      }
48
49      // Integer-like class
50      private class ReturnTypeExample$n7632 {
51         private int myValue;
52
53         public ReturnTypeExample$n7632(int value) {
54            super();
55            myValue = value;
56         }
57
58         // Getter method
59         public int j213() {
60            return myValue;
61         }
```

```
62
63          // Setter method
64          public void n4795(int value) {
65             myValue = value;
66          }
67       }
68    }
```

Other New Thread Classes

Any other new thread classes created during the process will have a `run()` method with an empty spin loop waiting until the original thread has completed. In future revisions of the algorithm, these additional threads may have domain-specific code inside of the spin loops, but for now they are relatively unstealthy. As an example, here is a possible inner-thread class:

```
1   private class SomeClass$s78 implements java.lang.Runnable {
2      public void run() {
3         // Just spin on some lock variable
4         while (b321.m987() != 1) {
5         }
6      }
7   }
```

Of course, the original method would also be modified to start this thread as well.

18.4 Sample Results

These experiments were performed on an Intel Pentium III 1.2 ghz machine running Red Hat Linux 7.2. The application that was obfuscated was the famous Sieve of Eratosthenes, an algorithm used to compute all prime numbers up to some **n**. In the test case, the **n** used was 99999999.

Number of threads used	*Percent Promotion*	*Seconds Needed*
1 (original)	0.0	12 seconds
1	1.0	52 seconds
2	0.0	25 seconds
2	1.0	73 seconds
3	0.0	38 seconds
3	1.0	122 seconds
4	0.0	64 seconds
4	1.0	411 seconds

Table 18.1: Sample results obtained using Degradation on the Sieve of Eratosthenes

Chapter 19

The Robust Object Watermarking Algorithm

Balamurugan Chirtsabesan, Tapas Sahoo

19.1 Introduction

The Robust Object Watermarking presents a new approach to watermarking. Instead of applying the watermarking scheme to the raw code directly, a new vector representation of the code is created. The basic idea revolving around this scheme is that instead of considering the overall structure of the code and its control flow, the code is viewed as a statistical object. The frequencies of groups of instructions in the entire code are taken into consideration in creating a new vector representation of the data. Spreading the watermark throughout the target code ensures a large measure of security against intentional and unintentional attacks.

Identifying the Vector instruction Groups: A set of instruction groups are identified that will form the watermark vector to be embedded. Ideally these instructions groups are instructions which have frequent occurence in the code and are chosen from some profiling information of the code. These instruction groups are small, each containing number of instructions in the range 1 to 4. This maximizes the probability of occurence of the instruction groups in the code while preserving the stealth.

Building the CodeBook: The CodeBook contains a mapping between semantically equivalent sets of instructions and is constructed in accordance with the vector instruction groups mentioned above. The basic motivation behind the construction of such a CodeBook is to facilitate the insertion of the watermark vector instructions through code subsitution. The CodeBook has been designed such that it can be upgraded as and when the vector instruction group is upgraded and new possible code equivalence is identified.

Vector Extraction: In this step, the vector length is chosen, the length of which is constrained by the number of the vector instruction groups. The frequency of occurence of the vector instruction groups in the code are identified and the intial vector is formed out of it.

Insertion: In this step, the watermark vector selected is to be embedded in the code. The frequencies of instructions in the vector instruction group are varied in accordance with the watermark vector by different implementation techniques such as code subsitution, new code insertion, cloning of method, etc. Care must be taken to ensure that all the operations done do not violate the semantics of the code.

Recognition: During recognition, the Vector Extraction procedure is invoked to extract the frequencies of the instructions in the vector group. The watermark is identified across a certain level of tolerance. A threshold is set prior to the extraction. Ideally, the watermark should be passed and the recognizer verifies the code and answers if the watermark vector is still present in the code within the threshold boundary.

19.2 Vector Instruction Group Identification

The instruction groups are identified keeping in mind the probability of occurence of the instructions in the code which is to be watermarked. Each vector group contains atmost 4 instructions. This is because a greater number of instructions in a vector group has less likelihood of occurence in the code. We have at present identified 8 vector instruction groups. The implementation has been done such that new groups can be appended to the existing vector groups as and when identified. The number of vector groups denotes the length of the watermark vector than can be embedded in the code.

```
1   package sandmark.watermark.objectwm;
2
3   public class ObjectWatermark extends sandmark.watermark.StaticWatermarker {
4
5      public ObjectWatermark();
6
7      public String getShortName();
8      public String getLongName();
9
10     public static sandmark.util.ConfigProperties getProperties();
11     public static void setProperties(sandmark.util.ConfigProperties);
12
13     public sandmark.util.ConfigProperties getConfigProperties();
14     public void setConfigProperties(sandmark.util.ConfigProperties);
15
16     public java.lang.String getAlgHTML();
17
18     public java.lang.String getAlgURL();
19
20     public void embed(java.util.Properties);
21        - Procedure to embed the watermark. Invokes the APIs in the Insertion class.
22
23     public java.util.Iterator recognize(java.util.Properties);
24        - Procedure for watermark recognizing.
25
26     class Recognizer implements java.util.Iterator {
27        public Recognizer(String jarInput);
28        public void generate();
29        public boolean hasNext();
30        public java.lang.Object next();
31     }
32  }
```

Figure 19.1: The general APIs used for implementing Algorithm for Robust Object Watermarking. label

```
1   package sandmark.watermark.objectwm;
2
3   public class CodeBook {
4
5      CodeBook();
6         - defines the CodeBook ie. the CodeGroup instruction groups, the vector
7           group instruction groups and the dependencies between them.
8
9      int getInstructionFromCodeBook(String[], int, int, int);
10        - invoked from the Insertion class to access the CodeBook.
11
12        private String[][] getParams(String[], int);
13
14        private String[] putParams(String[], int, String[][]);
15
16  }
```

Figure 19.2: The above figure shows the different APIs used for CodeBook access.

19.3 Building the CodeBook :

The CodeBook construction goes in parallel with the identification of the Vector Instruction Groups. Set of instructions are identified which are semantically equivalent and are placed in the same set. Moreover the equivalence should be such that it facilitates the increase of instruction vector frequencies. The following points are taken into consideration while building the CodeBook.

1. Each vector instruction group is dependent on atleast one CodeBook instruction group for its insertion.
2. Each instruction set in the CodeBook services exactly one vector instruction group. This simplifies the insertion procedure and isolates the frequency updates of different vector groups.
3. Each group in the CodeBook contains a set of instruction which has high frequency of occurence in the code. This gives more oppurtunities for code subsitution rather than new code insertion for vector frequency increment.
4. The code substitution increases the size of the code,but minimally.
5. The CodeBook has been designed such that it can be upgraded as and when new groups are identified with minimal effect on the APIs accessing the CodeBook.

A dependency array stores the dependence between each Vector Group Instruction and the set of CodeBook Instructions.

19.4 Vector Extraction

The vector extraction procedure is invoked at the initial stage prior to watermark insertion and during watermark recognition. The entire code is scanned for the occurence of the Vector Instruction Groups and the vector components are formed out of the frequencies of the corresponding groups.

19.5 The overall watermarking algorithm

Vector Extraction: Define n as a security parameter. Define a set S of n ordered groups of machine language instructions. For each group i in S, compute the frequency c_i of the group in the code, and form the vector

```
1   package sandmark.watermark.objectwm;
2
3   public class vectorExtraction {
4
5      public static java.util.Vector extractVector(String);
6         - implements the vector extraction procedure
7   }
```

Figure 19.3: VectorExtraction APIs

$c = (c_1, \ldots, c_n)$. This c is the extracted vector.

Algorithm
Initialization: Set a detection threshold σ.
Watermark Insertion:

- Apply the vector extraction step to obtain a vector c (of length n).
- Choose an n coordinates vector $w = (w_1, \ldots, w_n)$ whose coefficients are randomly distributed following a normal law with standard deviation α.
- Modify the code in such a way that the new extracted vector $\overline{c}$ is $c + w$.

Watermark Testing:

- Apply the extraction step to obtain a vector d.
- Compute a simliarity measure Q between $d - \overline{c}$ and w.
- If Q is higher than σ then the algorithm outputs *marked*, else it outputs *unmarked*.

19.6 Insertion

The insertion procedure is the most crucial step in the entire watermarking algorithm. For each Vector Instruction Group of the watermark vector, the algorithm intends to increase the frequency of that group by the vector value. It resorts to three techiniques to increase the vector group frequency.

- ***Code substitution:*** In this techinique we looks for group of instructions in the code that can be substituted by a equivalent set of instructions as enumerated in the CodeBook to increase the vector frequency. As an example, say vector group A is dependent on the CodeGroup[i,j]. We look for the occurence of CodeGroup[i,0] and substitute it by the equivalent set of instructions in the CodeGroup[i,j]. This increases the frequency of the vector by one. Some issues which we deal with while doing the substitution are enumerated below:

  1. We handle jump to a target out from this set of instructions appropriately by saving target instruction handles and later replacing it in the substituting instruction set.
  2. At present we do not handle jumps into the instruction set. So we discard the occurence of instruction sets which have branch targets as candidates for substitutions.
  3. The substitution increases the code length minimally as per the design of the CodeBook. We observed that finding an equivalence in which the vector frequency is increased as well the code size is reduced very difficult(unless the vector instruction group is restricted to almost one).

  Depending on the CodeBook implementation and the actual occurences of the instruction groups in the code there is an extent to which subsitution can be done to increment the vector frequency. Once this is exhausted, we resort to the next two techinques for vector increment.

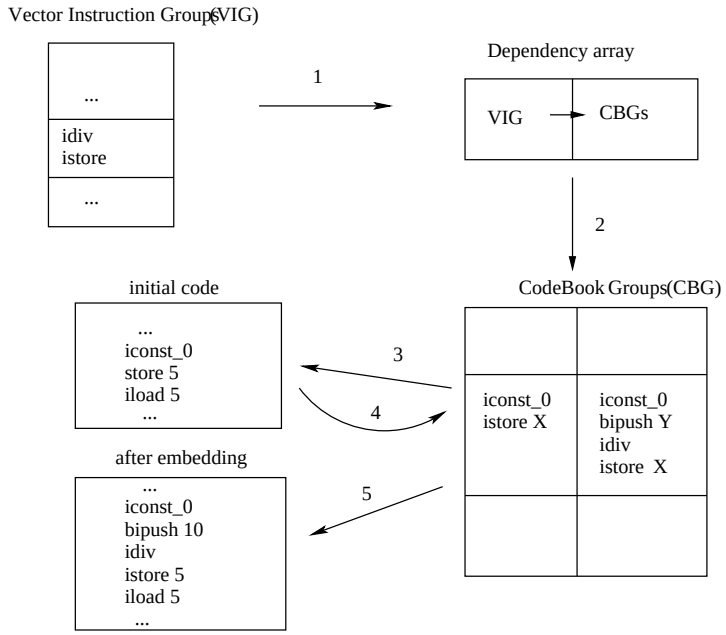

Figure 19.4:

- ***New code insertion:*** This is the most difficult part which we encountered in the entire implementation. First of all, instructions cannot be inserted at any random point in the code. After embedding, the modified code should pass the verifier. The following things were handled during the insertion:

  1. The substituting instructions corresponding to the vector are taken from the CodeBook group it is dependent on. To ensure that these set of instructions do not affect the stack size, we appended the instructions with required number of *pop* instructions.
  2. The insertion point is verified so that it does not breakup any other exisiting vector instruction group in the code.
  3. The variables used in the instructions to be inserted need to be declared and initialized in the method in which they are inserted prior to their use.
  4. To preserve the semantics of the code we inserted opaque predicates to jump over the newly inserted instructions. All the declarations for the predicate variables are done in the start of the method, though in the ideal case it could be anywhere before the point where they are used. Ideally we can have a set of such predicates in pool from which we can pick one in random and plugin into the method after inserting the vector instruction group codes.

- ***Method Cloning:*** In this techinique we identify the existing methods where we have occurence of the particular vector group instructions and create a clone of that method. This effectively increments the vector frequency. Some issues handled during this cloning are:

  1. Unlike other techniques, we have to keep track of the frequency updates of all the vectors since the replication might increase the vector frequencies of other vector groups.
  2. We maintain a threshold beyond which if the frequencies are modified, then we back track and do not replicate the method.

3. Ideally, we would make a bogus call to this clone method. But we do not implement such calls. In such a case a problem arises due to parameter passing. We need to initialize all the parameters in the clone method prior to all their use.

19.7 Comparison with the previous implementation on x86 assembly code

A previous implementation of this object watermarking was done in x86 assembly code. Each of these languages had its own constraints that restricted the implementation of the watermark embedding. In x86 assembly code we need to be careful in register selections during new instruction embedding as well as code substitution so that they dont interfere with each other. In the implementation we did on java bytecodes, we figured out that one needs to take into consideration the stack size while embedding new code. But no such problems arised during code substitutions which are semantically equivalent. The x86 code provides additional flexibility of swapping instruction sets or reordering instructions within an instruction group as long as the dependencies between the instructions are not violated, while such transformations have limited scope in the case of java bytecodes primarily due to the constraints of maintaining the proper stack height.

19.8 Analysis of attacks

The watermark faces the most serious threat from decompiling and recompiling the code. The compiler optimizations may change the sequence of instructions and even some instructions itself. The instructions in the vector group might well be the target of such modifications which results in the modification of the embedded frequency vector. But at the same time, since the watermark is spread across the entire code, the attacker has least idea about the exact location of the watermark. So to effect any substantial chenge in the vector frequency he needs to decompile and recompile the entire code rather than parts of the code which is tedious. Code compression poses the least threat since code has to be decompressed again to render it executable, so the watermark remains intact.

19.9 Brief analysis of our work and scope for future Work

We followed an extremely conservative approach during watermark embedding. One approach would have been to make an arbitrary change in the code and verify whether the new vector is closer to the target vector. If it digresses in the wrong direction then backtrack and try out another random modification. In our approach we always made sure prior to making any modification to the code that the resulting vector is always closer to the target vector. Such an approach was simple as well as preserved correctness of the code in all instances. The ratio of the code size before and after embedding increases minimally if the bulk of the code modification is done through substitution rather than new code insertion which apparently increases the code size drastically by a large factor. So its extremely important to identify the correct CodeBook and the Vector Instrution groups. The strengh of the watermark lies in the secret CodeBook which is maintained since there are always a wide range of possibilities of code equivalence and also the distribution of watermark over the entire code rather than being concentrated at a single point in the code. This increases the stealth of the watermark. Also we observed that codes with lower level of tolerance seem to withstand the tampering of this type of watermark more than code which provides more scope for modifications.

As a future work, we would like to identify some sort of profiling techniques to identify a proper set of vector and CodeBook instructions. We would also like to identify a few more approaches to embed the vector. The overall implementation needs to be modified to some extent so that we can have a framework in which we can plugin new code modification techniques.

```
1   package sandmark.watermark.objectwm;
2
3   public class Insertion {
4
5      public void modifyCode(java.util.Vector );
6         - Entry procedure for this 'Insertion' class
7
8
9      de.fub.bytecode.generic.InstructionHandle getCodeSubstPoint(String[], int, String[]);
10        - Searches and marks the instruction group occurence point in the code and returns the
11          instructions to be substituted.
12
13     void substituteCode(de.fub.bytecode.generic.InstructionHandle, String[], int);
14        - Subsitutes the new instructions  obtained from the codeBook into the code at the
15          marked instruction handle point.
16
17     void newsubstituteInstruction_Embed(String[], int);
18        - Implements the insertiong of previously non-existing code to increase the
19          vector frequency. This sets a random class-method-instruction insertion point for
20          insertion.
21
22
23        void substituteNewCode( de.fub.bytecode.generic.InstructionHandle, String[], int);
24           - Invoked by newsubstituteInstruction_Embed(); performs the actual low level insertion
25             operation.
26
27           String transformCode(String);
28              - invoked by substituteNewCode();  it populates the new instruction set to be
29                embedded with valid data and returns the final instruction to be inserted.
30
31              int getLocalVarIndex_CreateIndex();
32                 - invoked by transformCode(); returns a vaild local variable index. if none
33                   exist then it creates a new one and returns its index.
34
35
36           de.fub.bytecode.generic.Instruction extractInstrType(String);
37              - This API forms the instruction object to be inserted after parsing the input
38                instruction stream and returns the object.
39
```

Figure 19.5: The following APIs are implemented for the watermark Insertion procedure

```
39            void insertOpaque( de.fub.bytecode.generic.InstructionList,
40                               de.fub.bytecode.generic.InstructionHandle,
41                               de.fub.bytecode.generic.InstructionHandle,
42                               de.fub.bytecode.generic.InstructionHandle, int);
43               - Generates the opaque predicate and inserts it in appropriate position.
44
45
46            int getOpaqueDeclPoint(int);
47               - checks for and returns a $safe$ point for new instruction insertion.
48
49         boolean checkSplitVectorGrp(int, de.fub.bytecode.generic.InstructionHandle[]);
50            - Ensures that the insertion point for the new vector group does not split
51              another existing vector instruction group in the code and consequently
52              reduce its vector freqeuency.
53
54
55      int newcopyMethod_Embed(String[], int, int, int);
56         - Implements the method cloning procedure.
57
58
59      int remVecfreqUpdatesInThreshold(int, int);
60         - Ensures that the method cloning does not update the vecotr frequencies beyond
61           the {\em threshold}.
62
63
64      /* the various support APIs for the above APIs */
65
66      int getRandomValue(int, int);
67
68      void getInsertLocation(Integer[]);
69
70      int varTypeIsInt(int);
71
72   }
73
```

Figure 19.6: The following APIs are implemented for the watermark Insertion procedure

19.10 Conclusion:

The vector extraction paradigm produces a robust watermarking technique. It mainly takes advantage of the distribution of watermark over the code to maintain its stealth. But at the same time the effectiveness in the java bytecodes lies entirely on the exact implementation technique.

19.11 Bibliography

1. Julien P. Stern, Gael Hachez, Francois Koeune, and Jean-Jacques
Quisquate r. Robust Object Watermarking: Application to Code. In A. Pfitzmann, editor , Information Hiding '99, volume 1768 of Lectures Notes in Computer Science (LNCS), pages 368–378, Dresden, Germany, 2000. Springer-Verlag. http://ci teseer.nj.nec.com/stern00robust.html

2. I. Cox, J. Kilian, T. Leighton and T. Shamoon, "A secure, robust watermark for multimedia", in Proc. Workshop on Information Hiding, Univ. of Cambridg e, U.K., May 30 - June 1, 1996, pp. 175-190 http://citeseer.nj.nec.com/a rticle/cox96secure.html

Chapter 20

Potkonjak's Watermarking Algorithm

G. Myles

20.1 Introduction

Potkonjak's watermarking algorithm is based upon the ideas of graph coloring and register allocation. Java does not use registers to store values put instead uses local variables. So this algorithm will assign local variables based upon graph coloring instead of registers. The idea is to embed a watermark in a Java program that uniquely maps to one of the possible allocations of local variables. The reason that a watermark can be uniquely mapped to a local variable allocation is because a technique called bridge construction is used to generate multiple allocations from one specific allocation. By using this technique no additional code needs to be added so examining the code will not lead to the discovery of the watermark.

Currently this technique only works for a single method in a class, but in the future it will be extended so that an entire set of classes can be watermarked.

20.2 Embedding

The embedding of the watermark involves several different phases.

Convert watermark to binary

Converting the watermark, which is a string, to a binary is a fairly simple and straight forward task. This task is necessary because the binary is used to help choose the appropriate local variable allocation.

Generate the original coloring

The first real step of the algorithm is to determine the coloring of the method before any modifications have been made. This first step is necessary because a base is needed in order to base the binary relationship. It is not possible to used the local variable assignment that has already been allocated because BLOAT performs many optimizations in the process of completing the register allocation, which must be used in later phases of the algorithm. The majority of this step is handled by classes within BLOAT. To complete this task a control flow graph is created, SSA analysis is performed to generate the def-use information, liveness analysis is performed to generate the interference graph, and then register allocation takes place. Once these steps have been completed the instructions for the method are scaned to determine what loacal variable allocation was assigned.

Bridge construct and generate new coloring

Bridge construction is a technique used on graphs so that many colorings can be generated from the same graph. This is desirable since graph coloring is an NP problem and thus generating many colorings from the unmodified graph would be very intensive. With this algorithm it is necessary to generate as many colorings as possible so as to increase the chances of being able to embed the watermark in the method. Since the algorithm does not involve adding any code to the methods finding a suitable local variable allocation is crucial.

The bridge construction technique works by taking two unconnected nodes and connecting them and connecting the neighbors of one to the other. Each pair of nodes that can be connected generates 2^2 different colorings. If a set of three unconnected nodes can be found connecting them would generate 2^3 different colorings. It is also possible to combine these two sets to generate even more colorings. The bridge

construction is performed on the interference graph which was generated from the liveness analysis. Once the bridge construction has been completed the new interference graph is passed to the register allocator. Again after the register allocation has been completed the instructions are scaned to determine the allocation.

Determine if the method can be watermarked

Determining if the method can actually be watermarked is currently folded in with generating the colorings. Once the algorithm can be performed on multiple methods it will be necessary to check if the binary will fit perfectly in the program.

Generate the pool of graphs

In order to generate all of the different colorings it is necessary to know which edges were added to the interference graph. Currently the algorithm does not generate all of the colorings at once but instead generates certain sections of the coloring and determines if any of those possibilites will work. It they don't work then the method cannot be watermarked. The sections that are generated are the pairs of nodes that were connected. If two nodes were connected then those four solutions are generated and compared with the appropriate location in the binary to see if one will work. The nodes that are not connected will maintain their same coloring and therefore result in a 0 for the binary.

Pick the correct coloring

The coloring that is chosen is based upon the binary that needs to be embedded and the relationship between the original coloring and the generated colorings. By comparing the original to the generated coloring a binary relationship can be determined which corresponds to the binary watermark.

Embed the chosen coloring

Once a coloring has been selected it is passed to the register allocator along with the corresponding node label. The node labels are necessary because BLOAT does not appear to access the nodes of the graph in any particular order. This technique of passing the colors to the modified register allocator proved to be the easiest and most reliable because it takes care of both the defs and uses (stores and loads).

20.3 Recognition

Recognizing the watermark also involves several phases.

Determine the current coloring

Because the embedded watermark is based upon the comparison of two different local variable allocations, the first thing that must be done is to determine what allocation is currently assigned to the method. This is done simply by scanning the instructions in the method and looking for places where a store is perfomed. One a store has been found we can obtain the index of the local variable. The indexes are stored in a Vector, `currentColoring`, for later reference.

Generate the original coloring for the method

In order to generate the binary that dictated which coloring was chosen we have to determine the original local variable allocation for the method. This is done following the same steps that were used in order to generate the original. First we have to generate a control flow graph for the method. We then have to perform the SSA analysis in order to generate the def-use information that is used in the liveness analysis. After the SSA analysis the liveness analysis is performed so that an interference graph is generated. The interference graph is then used by the register allocator to assign local variable. Once all of this has been

completed it is necessary to scan the instructions again to determing the local variable allocation. Again we store this imformation in a Vector, `originalColoring`.

Compare the two colorings

Just as was done by in the embedding it is possible to compare the current coloring and the original coloring to get a binary. This binary will then become the watermark.

Convert binary to string

Again this is a simple task, but a necessary one. By converting the generated binary into a string we get the watermark back.

20.4 Future Work

Currently this watermarking technique has only been implemented so that a single method can be watermarked.

Part III

Appendices

Appendix A

Useful Tools

A.1 Examining Java Classfiles

There are a number of tools that are helpful for viewing Java classfiles.

javap

`javap` lists the contents of a Java classfile. It's particularly bad at displaying corrupted classfiles.

Normally, we call `javap` like this:

```
javap -c -s -verbose -l TTTApplication
```

These are the available options:

```
Usage: javap <options> <classes>...

where options include:
   -b                        Backward compatibility with javap in JDK 1.1
   -c                        Disassemble the code
   -classpath <pathlist>     Specify where to find user class files
   -extdirs <dirs>           Override location of installed extensions
   -help                     Print this usage message
   -J<flag>                  Pass <flag> directly to the runtime system
   -l                        Print line number and local variable tables
   -public                   Show only public classes and members
   -protected                Show protected/public classes and members
   -package                  Show package/protected/public classes
                             and members (default)
   -private                  Show all classes and members
   -s                        Print internal type signatures
   -bootclasspath <pathlist> Override location of class files loaded
                             by the bootstrap class loader
   -verbose                  Print stack size, number of locals and args for methods
                             If verifying, print reasons for failure
```

Jasmin

`Jasmin` is a Java bytecode assembler. It reads a text file containing Java bytecode instructions and generates a classfile. It can be found here: `http://mrl.nyu.edu/~meyer/jasmin/about.html`.

Here's is a simple class `hello.j` in the `Jasmin` assembler syntax:

```
1       .class public HelloWorld
2       .super java/lang/Object
3
4       ;
5       ; standard initializer (calls java.lang.Object's initializer)
6       ;
7       .method public <init>()V
8          aload_0
9          invokenonvirtual java/lang/Object/<init>()V
10         return
11      .end method
12
13      ;
14      ; main() - prints out Hello World
15      ;
16      .method public static main([Ljava/lang/String;)V
17         .limit stack 2   ; up to two items can be pushed
18
19         ; push System.out onto the stack
20         getstatic java/lang/System/out Ljava/io/PrintStream;
21
22         ; push a string onto the stack
23         ldc "Hello World!"
24
25         ; call the PrintStream.println() method.
26         invokevirtual java/io/PrintStream/println(Ljava/lang/String;)V
27
28         ; done
29         return
30      .end method
```

Call `Jasmin` like this:

```
java -classpath smextern/jasmin.jar jasmin.Main hello.j
```

which will produce a class file `HelloWorld.class`.

BCEL's Class Construction Kit

`cck` is an interactive viewer and editor of Java classfiles. It is built on `BCEL`. Call it like this:

```
java -classpath .:smextern/BCEL.jar -jar smextern/cck.jar
```

and open a classfile from the `file`-menu.

BCEL's listclass

`listclass` comes with the BCEL package. It's a replacement for `javap`, and very useful in cases when `javap` crashes. Call it like this:

```
java -classpath smextern/BCEL.jar listclass -code TTTApplication
```

The following options are available:

```
java listclass [-constants] [-code] [-brief] [-dependencies] \
     [-nocontents] [-recurse] class... [-exclude <list>]
```

`-code:` List byte code of methods.

`-brief:` List byte codes briefly

`-constants:` Print constants table (constant pool)

`-recurse:` Usually intended to be used along with

`-dependencies:` When this flag is set, listclass will also print i

BCEL's JustIce

`JustIce` verifies classfiles. Call it like this:

```
java -classpath .:smextern/BCEL.jar:smextern/JustIce.jar \
  de.fub.bytecode.verifier.GraphicalVerifier
```

or like this:

```
java -classpath .:smextern/BCEL.jar:smextern/JustIce.jar \
   de.fub.bytecode.verifier.Verifier TTTApplication.class
```

A.2 BCEL

SandMark reads Java classfiles, modifies them, and writes out the modified classfiles. There are a number of free packages available to parse, modify, and unparse classfiles. We are currently using two such packages, namely `BCEL` and `BLOAT`. Generally, `BLOAT` is comprehensive and hard to use, `BCEL` is simpler to use but not as complete.

`BCEL` (formerly `JavaClass`) allows you to load a class, iterate through the methods and fields, change methods, add new methods and fields, etc. `BCEL` also makes it easy to create new classes from scratch.

`BCEL` is here: `http://jakarta.apache.org/bcel/index.html`. There is a manual, and an on-line API description at `http://bcel.sourceforge.net/docs/index.html`. Unfortunately, neither are particularly informative.

SandMark contains a lot of `BCEL` code that might be worth studying:

- `sandmark.util.javagen` is a package for building up bytecode classes from scratch. Essentially, it allows you to build a Java AST (abstract syntax tree) and compile it down to bytecode. `javagen` is built on top of `BCEL`.
- `sandmark.util.BCEL` contains `BCEL` utility routines. The intention is that this class should grow in the future.
- The package `sandmark.watermark.CT.embed` contains several classes that manipulate Java classfiles using `BCEL`.

Getting started

To open a class file for editing you do the following:

```
1      sandmark.util.ClassFileCollection cfc = ...;
2
3      cfc = new sandmark.util.ClassFileCollection(jarFile);
4      de.fub.bytecode.classfile.JavaClass jc = cfc.getClass(className);
5      de.fub.bytecode.generic.ClassGen cg = new de.fub.bytecode.generic.ClassGen(jc);
6      de.fub.bytecode.generic.ConstantPoolGen cp =
7         new de.fub.bytecode.generic.ConstantPoolGen(jc.getConstantPool());
```

`jarFile` is the jar-file that contains the class file and `className` is the name of the class (possibly with a package prefix). `sandmark.util.ClassFileCollection` is the repository used within SandMark for loading and saving class files.

A `de.fub.bytecode.classfile.JavaClass`-object represents the parsed class file. You can query this object to get any information you need about the class, for example its methods, fields, code, super-class, etc.

A `de.fub.bytecode.generic.ClassGen`-object represents a class that can be edited. You can add methods and fields, you can modify existing methods, etc.

A `de.fub.bytecode.generic.ConstantPoolGen`-object represents a constant pool to which new constants can be added.

When you are finished editing the class you must close it and return it to the repository maintained by `sandmark.util.ClassFileCollection`:

```
8       de.fub.bytecode.classfile.JavaClass jc1 = cg.getJavaClass();
9       jc1.setConstantPool(cp.getFinalConstantPool());
10      cfc.addClass(jc1);
```

Editing a Method

`de.fub.bytecode.classfile.Method` represents a method that has been read from a class file, while `de.fub.bytecode.generic.MethodGen` is the `BCEL` class that represents a method being edited:

```
1       de.fub.bytecode.generic.ConstantPoolGen cp = ...;
2       de.fub.bytecode.generic.ClassGen cg = ...;
3
4       de.fub.bytecode.classfile.Method method =
5          cg.containsMethod(methodName, methodSignature);
6       de.fub.bytecode.generic.MethodGen mg =
7          new de.fub.bytecode.generic.MethodGen(method, className, cp);
```

Note that to open a method with `containsMethod` you need to know the class in which it is declared, as well as its name as well as its signature. The reason is that a Java class can have several methods with the same name (they're overloaded).

Alternatively, SandMark also provides a class `sandmark.util.EditedClass.open` than can be used to maintain a cache of classes and methods that have been opened by `BCEL`:

```
1        sandmark.util.ClassFileCollection cfc = ...;
2        sandmark.util.EditedClass ec = sandmark.util.EditedClass.open(cfc, className);
3       java.util.Iterator methods = ec.methods();
4       while (methods.hasNext()){
5           sandmark.util.MethodID m = (sandmark.util.MethodID) methods.next();
6           de.fub.bytecode.generic.MethodGen mg = ec.openMethod(m);
7           ....
8       }
```

A `sandmark.util.MethodID`-object represents a method within a Java application. It is essentially a tuple

$$\langle \texttt{name}, \texttt{signature}, \texttt{className} \rangle.$$

Editing Instructions

Once you have opened a method for editing you can get its list of instructions:

```
1       de.fub.bytecode.generic.MethodGen mg = ...;
2
3       de.fub.bytecode.generic.InstructionList il = mg.getInstructionList();
```

```
4      de.fub.bytecode.generic.InstructionHandle[] ihs = il.getInstructionHandles();
5      for(int i=0; i < ihs.length; i++) {
6         de.fub.bytecode.generic.InstructionHandle ih = ihs[i];
7         de.fub.bytecode.generic.Instruction instr = ih.getInstruction();
8            ....
9      }
```

`de.fub.bytecode.generic.InstructionLists` can be manipulated by appending or inserting new instruction, deleting instructions, etc.

All bytecode instructions in `BCEL` are represented by their own class. `de.fub.bytecode.generic.IADD` is the class that represents integer addition, for example. This allows us to use `instanceof` to classify instructions:

```
1      de.fub.bytecode.generic.Instruction instr = ...;
2      de.fub.bytecode.generic.InstructionHandle ih = ...;
3
4      if (instr instanceof de.fub.bytecode.generic.INVOKESTATIC) {
5         de.fub.bytecode.generic.INVOKESTATIC call =
6           (de.fub.bytecode.generic.INVOKESTATIC) instr;
7         String className  = call.getClassName(ec.cp);
8         String methodName = call.getMethodName(ec.cp);
9         String methodSig  = call.getSignature(ec.cp);
10
11        ih.setInstruction(new de.fub.bytecode.generic.NOP());
12     }
```

`setInstruction` is used to replace an original instruction with a new one.

Local variables

Local variables are quite complicated in Java bytecode. A Java method can have a maximum of 256 local variables. Note, however, that

- In a virtual (non-static) method local variable zero ("slot #0") is `this`.
- Formal arguments take up the first slots. In other words, in a virtual method with 2 formal parameters, slot #0 is `this`, slot #1 is the first formal, slot #2 is the second formal, and any local variables are slots #3,#4,...,#255.
- Local variables can be reused within a method. For example, consider the following method:

```
void P() {
   int x= 5;
   System.out.println(x);
   float y=5.6;
   System.out.println(y);
}
```

A clever compiler will allocate both `x` and `y` to the same slot, since they have non-overlapping lifetimes.

For this reason, `BCEL` requires that we indicate where a new local variable is accessible:

```
1      de.fub.bytecode.generic.Type T = ...;
2
3      de.fub.bytecode.generic.LocalVariableGen lg =
4         mg.addLocalVariable(localName, T, null, null);
5      int localIndex = lg.getIndex();
```

```
6
7       // push something here
8       de.fub.bytecode.generic.Instruction store =
9          new de.fub.bytecode.generic.ASTORE(localIndex);
10
11      de.fub.bytecode.generic.InstructionHandle start = ...;
12      lg.setStart(start);
```

`de.fub.bytecode.generic.LocalVariableGen` creates a new local variable. The two `null` arguments are the locations within the method where we want the variable to be visible. We start these locations out as `null` (unknown) and fill them in using the `setStart` method. It takes an `de.fub.bytecode.generic.InstructionHandle` as input, namely the first instruction where the variable should be visible.

Example 1: Creating a New Class

Here's a longer example that shows how you create a class from scratch.

We start by creating a new `de.fub.bytecode.generic.ClassGen` object representing the class. We provide the name of the class, the name of the parent class, the interfaces it implements, and its access flags:

```
1       int access_flags = de.fub.bytecode.Constants.ACC_PUBLIC;
2       String class_name = "MyClass";
3       String file_name = "MyClass.java";
4       String super_class_name = "java.lang.Object";
5       String[] interfaces = {};
6
7       de.fub.bytecode.generic.ClassGen cg =
8          new de.fub.bytecode.generic.ClassGen(
9             class_name, super_class_name, file_name,
10            access_flags, interfaces);
11      de.fub.bytecode.generic.ConstantPoolGen cp = cg.getConstantPool();
```

We can then add a field `field1` of type `int` to the new class:

```
12      de.fub.bytecode.generic.Type field_type =
13         de.fub.bytecode.generic.Type.INT;
14      String field_name = "field1";
15
16      de.fub.bytecode.generic.FieldGen fg =
17         new de.fub.bytecode.generic.FieldGen(
18            access_flags, field_type, field_name, cp);
19      cg.addField(fg.getField());
```

Next, we create a method `method1`:

```
20      int method_access_flags =
21         de.fub.bytecode.Constants.ACC_PUBLIC |
22         de.fub.bytecode.Constants.ACC_STATIC;
23
24      de.fub.bytecode.generic.Type return_type =
25         de.fub.bytecode.generic.Type.VOID;
26
27      de.fub.bytecode.generic.Type[] arg_types =
28         de.fub.bytecode.generic.Type.NO_ARGS;
29      String[] arg_names = {};
30
31      String method_name = "method1";
```

```
32      String class_name = "MyClass";
33
34      de.fub.bytecode.generic.InstructionList il =
35         new de.fub.bytecode.generic.InstructionList();
36      il.append(de.fub.bytecode.generic.InstructionConstants.RETURN);
37
38      de.fub.bytecode.generic.MethodGen mg =
39         new de.fub.bytecode.generic.MethodGen(
40            method_access_flags, return_type,  arg_types,
41            arg_names, method_name, class_name, il, cp);
42
43      mg.setMaxStack();
44      cg.addMethod(mg.getMethod());
```

The method has no arguments, returns `void`, and contains only one instruction, a `return`.

Note the call to `mg.setMaxStack()`. With every method in a class file is stored the maximum stack-size is needed to execute the method. For example, a method whose body consists of the instructions ⟨`iconst_1`, `iconst_2`, `iadd`, `pop`⟩ (i.e. compute $1 + 2$ and throw away the result) will have max-stack set to two. We can either compute this ourselves or let `BCEL`'s `mg.setMaxStack()` do it for us.

Branches

As is always the case with branches, we need to be able to handle forward jumps. The typical way of accomplishing this is to create a branch with unspecified target:

```
1       de.fub.bytecode.generic.InstructionList il = ...;
2
3       de.fub.bytecode.generic.IFNULL branch =
4          new de.fub.bytecode.generic.IFNULL(null);
5       il.append(branch);
6       ...
```

When we have gotten to the location we want to jump to we add a (bogus) `NOP` instruction which will serve as our branch target. The `append` instruction returns a handle to the `NOP` and we use this handle to set the target of the branch instruction:

```
7       de.fub.bytecode.generic.InstructionHandle h =
8          il.append(new de.fub.bytecode.generic.NOP());
9       branch.setTarget(h);
```

Exceptions

Here is the generic code for building a `try-catch`-block:

```
1       de.fub.bytecode.generic.InstructionList il = ...;
2       de.fub.bytecode.generic.MethodGen mg = ...;
3       String exception = "java.lang.Exception";
4
5       de.fub.bytecode.generic.InstructionHandle start_pc =
6          il.append(new de.fub.bytecode.generic.NOP());
7
8       // The code that builds the try-body goes here.
9
10      // Add code to jump out of the try-block when
11      // no exception was thrown.
12      de.fub.bytecode.generic.GOTO branch =
```

```
13          new de.fub.bytecode.generic.GOTO(null);
14      de.fub.bytecode.generic.InstructionHandle end_pc =
15           il.append(branch);
16
17      // Pop the exception off the stack when entering the catch block.
18      de.fub.bytecode.generic.InstructionHandle handler_pc =
19          il.append(new de.fub.bytecode.generic.POP());
20
21      // The code that builds the catch-body goes here.
22
23      // Add a NOP after the exception handler. This is
24      // where we will jump after we're through with the
25      // try-block.
26      de.fub.bytecode.generic.InstructionHandle next_pc =
27          il.append(new de.fub.bytecode.generic.NOP());
28      branch.setTarget(next_pc);
29
30      de.fub.bytecode.generic.ObjectType catch_type =
31          new de.fub.bytecode.generic.ObjectType(exception);
32
33      de.fub.bytecode.generic.CodeExceptionGen eg =
34          mg.addExceptionHandler(start_pc, end_pc, handler_pc, catch_type);
```

`start_pc` and `end_pc` hold the beginning and the end of the `try`-body, respectively. `handler_pc` holds the address of the beginning of the `catch`-block.

Types

Types (such as method signatures) are defined by the class `de.fub.bytecode.generic.Type`:

```
1       de.fub.bytecode.generic.Type return_type =
2          de.fub.bytecode.generic.Type.VOID;
3       de.fub.bytecode.generic.Type[] arg_types   =
4          new de.fub.bytecode.generic.Type[] {
5             new de.fub.bytecode.generic.ArrayType(
6                de.fub.bytecode.generic.Type.STRING, 1)
7          };
```

A source-level Java type is encoded into a classfile type descriptor using the following definitions:

BaseType	Type	Interpretation
B	byte	signed 8-bit integer
C	char	Unicode character
D	double	64-bit floating point value
F	float	32-bit floating point value
I	int	32-bit integer
J	long	64-bit integer
L<***classname***>**;**	reference	an instance of class <*classname*>
S	short	signed 16-bit integer
Z	boolean	`true` pr `false`
[	reference	one array dimension
V		void
(***BaseType****)***BaseType**		method descriptor

`BCEL` has quite a number of methods that convert back and forth between Java source format (such as `"java.lang.Object[]"`), bytecode format (such as `"[Ljava/lang/Object;"`), and `BCEL`'s internal format

(de.fub.bytecode.generic.Type). Unfortunately, these are by and large poorly documented. We will describe some of the more common methods here.

de.fub.bytecode.generic.Type.getType converts from Java bytecode format to Java source format. In other words, the following code

```
1       String S = "[Ljava/lang/Object;";
2       de.fub.bytecode.generic.Type T =
3          de.fub.bytecode.generic.Type.getType(S);
4       System.out.println(T);
```

prints out ⌜java.lang.Object[]⌝.

The method de.fub.bytecode.classfile.Utility.getSignature converts a type in Java bytecode format to Java source format. The code

```
1       String type = "java.lang.Object[]";
2       String S = de.fub.bytecode.classfile.Utility.getSignature(type);
3       System.out.println(S);
```

prints ⌜[Ljava/lang/Object;⌝.

The methods de.fub.bytecode.generic.Type.getArgumentTypes and de.fub.bytecode.generic.Type.getReturnType take a type in Java bytecode format as argument and extract the array of argument types and the return type, respectively:

```
1       String S = "(Ljava/lang/String;I)V;";
2       de.fub.bytecode.generic.Type[] arg_types =
3          de.fub.bytecode.generic.Type.getArgumentTypes(S);
4       de.fub.bytecode.generic.Type return_type =
5          de.fub.bytecode.generic.Type.getReturnType(S);
6       String M = de.fub.bytecode.generic.Type.getMethodSignature(return_type, arg_types);
7       System.out.println(M);
```

de.fub.bytecode.generic.Type.getMethodSignature converts the return and argument types back to a Java bytecode type string.

The method de.fub.bytecode.generic.Type.getSignature converts a BCEL type to the equivalent Java bytecode signature. This code

```
1       de.fub.bytecode.generic.Type T =
2          de.fub.bytecode.generic.Type.STRINGBUFFER;
3       String M = T.getSignature();
4       System.out.println(M);
```

will print out ⌜Ljava/lang/StringBuffer;⌝.

As a convenience, SandMark provides the method sandmark.util.javagen.Java.typeToByteCode, that converts a type in Java source format to a BCEL de.fub.bytecode.generic.Type:

```
1      String S = de.fub.bytecode.classfile.Utility.getSignature(type);
2      return de.fub.bytecode.generic.Type.getType(S);
```

Method Calls

To make a static method call we push the arguments on the stack, create a constant pool reference to the method, and then generate an INVOKESTATIC opcode that performs the actual call:

```
1      String className  = ...;
2      String methodName = ...;
3      String signature  = ...;
4
```

```
5       de.fub.bytecode.generic.InstructionList il = ...;
6       de.fub.bytecode.generic.ConstantPoolGen cp = ...;
7
8       // Generate code that pushes the actual arguments of the call.
9
10      int index = cp.addMethodref(className,methodName,signature);
11      de.fub.bytecode.generic.INVOKESTATIC call =
12          new de.fub.bytecode.generic.INVOKESTATIC(index);
13      il.append(call);
```

Making a virtual call is similar, except that we need an object to make the call through. This object reference is pushed on the stack prior to the arguments:

```
1       String className  = ...;
2       String methodName = ...;
3       String signature  = ...;
4
5       de.fub.bytecode.generic.InstructionList il = ...;
6       de.fub.bytecode.generic.ConstantPoolGen cp = ...;
7
8       // Generate code that pushes the object on the stack.
9
10      // Generate code that pushes the actual arguments of the call.
11
12      int index = cp.addMethodref(className,methodName,signature);
13      de.fub.bytecode.generic.INVOKEVIRTUAL s =
14          new de.fub.bytecode.generic.INVOKEVIRTUAL(index);
15      il.append(call);
```

Example 2: Modifying an Existing Program

Consider the following program:

```
1       public class C {
2          public static void main(String args[]) {
3             int a = 5;
4             int b = 6;
5             int c = a + b;
6             System.out.println(c);
7          }
8       }
```

Control Flow Graphs

The `JustIce` verifier that is built on top of `BCEL` has functions that construct control flow graphs for methods. Unfortunately, they don't construct real basic blocks: each instruction is in a block by itself.

Miscellaneous

`sandmark.util.javagen.Java.accessFlagsToByteCode` takes a list of string modifiers as argument and returns the corresponding access flag:

```
1       String[] attributes = {"public","static"};
2       int access_flags =
3          sandmark.util.javagen.Java.accessFlagsToByteCode(attributes);
```

A.3 BLOAT

Introduction

What is BLOAT?

BLOAT is a Java optimizer written entirely in Java. By optimizing Java byte code, code improvements can occur regardless of the compiler that compiled the byte code or the virtual machine on which the byte code is run. BLOAT performs many traditional program optimizations such as constant/copy propagation, constant folding and algebraic simplification, dead code elimination, and peephole optimizations. Additionally, it performs partial redundancy elimination of arithmetic and field access paths.

Beyond this BLOAT can also be used for class creation as well as method creation. In combination with byte code level editing BLOAT is can also be used as a framework for Java Class File obfuscation.

What to expect?

This how-to is in no way related to the creators of BLOAT. It was written by Richard Smith, a computer science undergraduate at the University of Arizona, to aid in research efforts by Christian Collberg. This guide does not demonstrate the optimization capabilities of BLOAT, but does demonstrate its ability to be used as a framework for byte code obfuscation. From this knowledge you can easily infer how to do optimizations. For further questions about this how-to email: `rsmith@cs.arionza.edu`. You may also look at The BLOAT Book which can be found at `http://www.cs.purdue.edu/s3/projects/bloat`. At the time of writing this how-to this book referred to BLOAT-0.8a which lacks many of the features of BLOAT-1.0.

A number of examples of how BLOAT is used can be found in the cvs directory `smbloat2`. Section 1.2 describes how to check out this directory.

The Basics

Architecture

BLOAT is a very complicated beast. This will be a brief overview, but will be enough to begin with. The first you must understand is the concept of a context. A context supplies a means of loading and editing classes. All contexts implement the `EditorContext`. This interface supplies the generic means of committing edits, editing classes, and a utility to load classes. There are five major contexts:

1. `BloatContext`,
2. `InlineContext`,
3. `EditorContext`,
4. `PersistentBloatContext`, and
5. `CachingBloatContext`.

The `InlineContext` and the `EditorContext` are both interfaces. The `InlineContext` and the `EditorClass` are both implemented by the abstract class `BloatContext`. The `PersistentBloatContext` and the `CachingBloatContext` are both subclasses of the `BloatContext`. The two contexts that you should worry about are the `PersistentBloatContext` and the `CachingBloatContext`. These contexts are used for the same purpose, which is to act as a repository for all of the edits made to java class files. The difference between the two contexts is that the `PersistentBloatContext` keeps every piece of data relevant to your work, the `CachingBloatContext` on the other hand manages your edits so that when a editor is no longer dirty it drops the instance of the editor.

Editors consist of the `ClassEditor`, `MethodEditor`, and `FieldEditor`. The names of these class should give evidence enough of their subjects of manipulation. It is through these editors that you will make actual changes to byte code. It should be noted that the `ClassEditor` constructor takes a `BloatContext` and all other editors accepts a `ClassEditor`. In some ways this is the use of a decorator pattern or wrapper methodology.

```
1   EDU.purdue.cs.bloat.file.ClassFileLoader classFileLoader = new
2      EDU.purdue.cs.bloat.file.ClassFileLoader();
3
4   EDU.purdue.cs.bloat.context.PersistentBloatContext bloatContext =
5      new EDU.purdue.cs.bloat.context.PersistentBloatContext(
6          classFile.loader());
7
8   String className = "ClassFileToTest";
9   EDU.purdue.cs.bloat.reflect.ClassInfo info = null;
10
11  try {
12     info = bloatContext.loadClass(className);
13  } catch (ClassNotFoundException ex) {
14     System.err.println("Couldn't find class: " +ex.getMessage());
15     System.exit(1);
16  }
17
18  EDU.purdue.cs.bloat.file.ClassFile classFile =
19        (EDU.purdue.cs.bloat.file.ClassFile)info;
```

Figure A.1: Loading a class file.

Editing Class Files

Where to Start?

BLOAT seems to have endless constructors that depend on other objects which in turn depend on other objects. In truth unless you know where to begin you will wrap yourself around the API for a few hours trying to find the object you need create in order to create the object you want to create.

The first thing that you need in BLOAT is a `ClassFileLoader`. Unlike most classes in BLOAT the constructor for this class takes no arguments. The `ClassFileLoader` provides you with tools to manipulate the system CLASSPATH, open zip/jar archives, and create new classes. Overall these tools will not be useful to you. They can be used as opposed to the process I am about to demonstrate to you. However, if you chose to do so you will not be using contexts as the central point of your BLOAT manipulations which was fore mentioned as the core of the BLOAT architecture.

After creating a `ClassFileLoader` the next step is to create either a `PersistentBloatContext` or a `CachingBloatContext`. The choice between these two contexts depends on the resources available to you. In most case you should chose the `PersistentBloatContext` as memory is cheap.

Once you have a BLOAT context you can load the class files you wish to edit. This is done by using the `loadClass` method of the BLOAT context. You could also use the same method in the `ClassFileLoader`, however again this does not take advantage of the context architecture.

After loading the class file the `loadClass` method returns a `ClassInfo` object. This is really an interface implemented by the `ClassFile` object so you will want to cast the `ClassInfo` object to a `ClassFile` to take advantage of some more advanced features found in the `ClassFile` object. The `ClassFile` object, as one might think, contains all of the relevant information about the class you load. This includes information pertaining to fields and methods as well as more byte code specific information. Figure A.1 will illustrate this process so far.

The `ClassEditor`

Once you created a `ClassFile` object you can begin byte-code manipulation. The next step is to create a `ClassEditor`. There are two ways to go about this, create a `ClassEditor` from scratch or create a `ClassEditor` from an existing `ClassFile`.

To create a `ClassEditor` from an existing `ClassFile` simply use the method `editClass` supplied by the `BloatContext`. This method returns a `ClassEditor` for the `ClassFile` passed as a parameter:

```
1   EDU.purdue.cs.bloat.editor.ClassEditor classEditor =
2                 bloatContext.editClass(classFile);
```

To create a `ClassEditor` from scratch is a little different. BLOAT supplies a constructor for the `ClassEditor` to do so however the API advises you not to use this. Instead you should use the method `newClass` provided by the `BloatContext`:

```
newClass(
   intmodifiers,
   java.lang.StringclassName,
   TypesuperType,
   Type[]interfaces);
```

The `newClass` method takes a modifier, a String, a Type, and a Type array. The modifier specifies whether the class will be public, private, or protected. The correct parameter to use for the modifier is a constant from the modifier interface found in the `EDU.purdue.cs.bloat.reflect` package:

```
EDU.purdue.cs.bloat.reflect.Modifier.PUBLIC
EDU.purdue.cs.bloat.reflect.Modifier.PRIVATE
EDU.purdue.cs.bloat.reflect.Modifier.PROTECTED
```

The String parameter specifies the name of the class being created. The Type passed to the constructor indicates from where the class extends i.e. java.lang.Object. Again this is specified by a constant from the `EDU.purdue.cs.bloat.editor` package:

```
EDU.purdue.cs.bloat.editor.Type.OBJECT
```

The Type array specifies all of the interfaces that the class is to implement. Again these can be found in the `EDU.purdue.cs.bloat.editor` package and include interface such as Serializable and Comparable. Here's a complete example of the use of this method:

```
1   EDU.purdue.cs.bloat.editor.ClassEditor classEditor =
2      bloatContext.newClass(
3         EDU.purdue.cs.bloat.reflect.Modifiers.SUPER,
4         "MyNewClass",
5         EDU.purdue.cs.bloat.editor.Type.OBJECT,
6         null);
```

Note that if you intend to implement an interface or override a method that is inherited you will have to create those methods from scratch as well. The section on the method editor will demonstrate method creation from scratch.

The `ClassEditor` provides manipulation of data such as the class modifiers, the constant pool, and the interfaces implemented by the class. All of these uses are pretty intuitive with the exception of constant pool manipulation. You should refer to `ChangeFieldNames.java`, provided in the directory examples-1.0, for an example of constant pool manipulation.

The `MethodEditor`

If you wish to edit a method in an existing class you will again need to refer to the `BloatContext`. The `BloatContext` provides a method, `editMethod`, to modify methods. The `editMethod` method takes a `MethodInfo` object as a parameter. You can retrieve this information from a `ClassEditor` with the method,`methods`, that returns an array of `MethodInfo` objects:

```
1   EDU.purdue.cs.bloat.reflect.MethodInfo[] methods =
2      classEditor.methods();
3   EDU.purdue.cs.editor.MethodEditor methodEditor =
4      bloatContext.editMethod(methods[0]);
```

This array contains the method information for all of the methods contained in the `ClassFile` related to the class editor including the constructor. Note that a `MethodInfo` object is really an interfaced used by a `Method` object and therefore you can cast a `MethodInfo` object to a `Method` object.

There is most likely a specific method that you wish to edit. You can retrieve a specific method by obtaining its name with the name method in the `MethodEditor` object. To create a new method for an existing class or a class you creating from scratch you must create a new `MethodEditor` object. Note that unlike creating a `ClassEditor` from scratch this is not done through the BLOAT context. To create a new `MethodEditor` use the constructor provided below:

```
MethodEditor(
   ClassEditoreditor,
   intmodifiers,
   TypereturnType,
   java.lang.StringmethodName,
   Type[]paramTypes,
   Type[]exceptionTypes);
```

The constructor first takes a `ClassEditor` for the class you wish to create a new method in. Then you must provide a modifier such as public, private, or protected. This can be specified by a constant in the `EDU.purdue.cs.reflect.Modifier` interface. The next parameter specifies the return type of the method. This can be specified using a constant from the `EDU.purdue.cs.editor.Type` interface. The String parameter is used to name the method you have chosen to create. Then the parameters of the method are specified by constants from the Type interface. The final parameter of the constructor specifies exceptions thrown by the method. Again this can be specified using constants found in the Type interface.

Once you have created a method you can begin adding instructions to the method. This done by calling the method `addInstruction` in the `MethodEditor`. The method takes an opcode constant as a parameter. These constants can be found in the `EDU.purdue.cs.bloat.editor.Opcode` interface. There is an opcode in this interface corresponding to each opcode in Java byte-code. It should be noted that valid Java byte-code must be entered or you will receive a Java runtime error.

The `FieldEditor`

The `FieldEditor` is used to edit attributes of a class. To edit fields of an existing class simply call the `editField` method of the `BloatContext`. This method takes a `FieldInfo` object. This object can be obtained from the `ClassEditor` by calling the fields method. This method returns an array of `FieldInfo` objects. This example demonstrates how to put these things together to create a `FieldEditor`:

```
1  EDU.purdue.cs.bloat.editor.FieldInfo[] fieldInfo =
2     classEditor.fields();
3  EDU.purdue.cs.bloat.editor.FieldEditor fieldEditor =
4     bloatContext.editField(fieldInfo[0]);
```

To create a new field you must use a constructor found in the `FieldEditor` class. The constructor is illustrated below:

```
FieldEditor(
   ClassEditoreditor,
   intmodifiers,
   java.lang.Classtype,
   java.lang.Stringname);
```

The constructor takes the `ClassEditor` for the class you wish to add an attribute to. Like the constructor to create a new method the constructor for the `FieldEditor` also takes a constant to specify a Modifier to specify whether the attribute is public, private, or protected. The constructor also takes a Type constant to indicate whether the attribute is an int, String, or other possibility. Finally the constructor accepts a String to name the attribute.

Committing Edits

One Rule To Follow

There is one rule to follow in BLOAT that will save you a lot of pain, after you edit commit. This seems pretty obvious but there are al lot of methods to `commit` edits in BLOAT. For instance the `ClassEditor`, `FieldEditor`, `MethodEditor`, and the `BloatContext` all have several commit methods. The rule can be applied as followed. If you make an edit with a `FieldEditor`, or any other editor, call the `commit` method of the editor when you have completed editing with that specific editor. The `BloatContext`'s `commit` method can be used to commit all edits to all class files the `BloatContext` has knowledge of, i.e. any class that you have loaded or created through the `BloatContext`.

Instruction Representation

Opcodes

The folks who created BLOAT were nice enough to create a constant to represent every opcode in the Java byte-code language. These opcodes can be found in the `EDU.purdue.cs.bloat.editor.Opcode` interface.

Instructions

You can create an instruction by simply passing an Instruction object an opcode constant form the `EDU.purdue.cs.bloat.editor.Opcode` interface.

```
1   EDU.purdue.cs.bloat.editor.Instruction instruction =
2       new EDU.purdue.cs.bloat.editor.Instruction(
3           EDU.purdue.cs.bloat.editor.Opcode.opc_astore);
```

There are several methods in the Instruction class that test whether the instruction the class represents is a jump or a switch as well as many other things. Each of these methods returns a boolean value based on their validity.

Each instruction can be add to a method or block in a control flow graph. You can do this by calling the `addInstruction` method for either of these objects and passing it an Instruction object.

Advanced Features

Inlining

Inlining is a fairly straight forward process. Simply create a `BloatContext` and load your classes as usual. You will now need to create an `Inline` object. The constructor for the object takes a `InlineContext` and an integer as a parameter. There are no special steps for obtaining an `InlineContext` the `BloatContext` already implements this interface. The integer determines the maximum number of instructions a method can grow to.

```
1   EDU.purdue.cs.bloat.inline.Inline inline =
2       new EDU.purdue.cs.bloat.inline.Inline(bloatContext,20);
```

To actually inline a method you must call the `inline` method of the Inline class. This method takes a `MethodEditor` as a parameter. Refer to the section on `MethodEditor`s if you do not already know how to create one. Once the method is called the inlining is complete:

```
1   EDU.purdue.cs.bloat.inline.Inline inline =
2       new EDU.purdue.cs.bloat.inline.Inline(bloatContext,20);
3   inline.inline(methodEditor);
```

There are some extra features that can be used during inlining. Using the method `maxCallDepth` you can specify with an integer the maximum number of nested method calls inlined. The inlining process can be rather time consuming. Bloat attempts to remedy this by providing a method, `setMaxInlineSize`, that sets the maximum size method that will be inlined. This size is set by passing an integer to the `setMaxInlineSize` method.You may also specify whether to inline methods that throw exceptions or not. This is simply set passing a boolean value to the method `setInlineExceptions`. Note that if you need to use one these features you must call one of these methods prior to calling the `inline` method.

```
1   int size = 25;
2   int depth = 3;
3   boolean flag = true;
4   inline.setMaxInlineSize(size);
5   inline.setMaxCallDepth(depth);
6   inline.setInlineExceptions(flag);
7   inline.inline(methodEditor);
```

Class Hierarchy

It maybe of some interest to understand or manipulate the class hierarchy of a class. The first step in this process is to load all classes in the hierarchy using your current `BloatContext`. The `BloatContext` must first have knowledge of all of the class relevant to the hierarchy. Once the classes are load you can extract information on the hierarchy by calling `getHierarchy` method found in the `BloatContext`. The method returns a `ClassHierarchy` object. The `ClassHierarchy` object allows you to view all of the classes and interfaces in a hierarchy as well as edit the hierarchy by adding or removing classes an interfaces.

```
1   // Add a class
2   EDU.purdue.cs.bloat.editor.ClassHierarchy classHierarchy =
3      bloatContext.getHeirarchy();
4   String className = "Test";
5   classHierarchy.addClassNamed(Test);
6   // extract the Type of classes in the hierarchy:
7   Collection classes =
8      classHierarchy.classes();
9   // extract Type of classes that implement a specific interface:
10  Collection classes =
11     classHierarchy.implementors(
12        EDU.purdue.cs.bloat.editor.Type.SERIALIZABLE);
13  // find interfaces a Type of object implements
14  Collection interfaces =
15     classHierarchy.interfaces(
16        EDU.purdue.cs.bloat.editor.Type.CLASS);
17  // extract subclasses of a Type of object
18  Collection subclasses =
19     classHierarchy.subclasses(EDU.purdue.cs.bloat.editor.Type.OBJECT);
```

The `ClassHierarchy` may also be used to determine what method will be invoked at any depth in the hierarchy. Therefore you can determine from where a method is inherited or overwritten.

```
// simulates dynamic dispatching:
MemberRef methodInvoked(
   Type receiver,
   NameAndType method);
```

The method takes two arguments a Type, which we are already familiar with from previous examples, and a `NameAndType`. The Type argument determines the receiver of the call. The `NameAndType` class simply

describes an a method by its name and descriptor. The name is simply a String and the descriptor describes the return type of the method using the Type class.

```
EDU.purdue.cs.bloat.editor.NameAndType(
   java.lang.String name,
   Type type);
```

The `MemberRef` returned by `methodInvoked` represents the method you are interested in and class from where it was invoked.

```
1   //Putting it all together:
2   EDU.purdue.cs.bloat.editor.ClassHierarchy classHierarchy =
3      bloatContext.getHeirarchy();
4   EDU.purdue.cs.bloat.editor.NameAndType nameAndType =
5      new EDU.purdue.cs.bloat.editor.NameAndType(
6         "test",
7         EDU.purdue.cs.bloat.editor.Type.STRING);
8   MemberRef memRef =
9      classHeirarchy.methodInvoked(
10        EDU.purdue.cs.bloat.editor.Type.OBJECT,
11        nameAndType);
12  System.out.println(memRef);
```

Using the `ClassHierarchy` you may also determine if a method has been overwritten. Note the previous example could be used to determine where it was overwritten. In order to determine if a method has been overwritten you must call the `methodIsOverridden` method:

```
boolean methodIsOverridden(
   Type classType,
   NameAndType method);
```

To use this method you must pass it the class Type of where the method resides and a `NameAndType` that identifies the method and its return type. The method will return a boolean value based on whether the method has been overwritten for that class Type.

```
1   //Putting it all together
2   EDU.purdue.cs.bloat.editor.ClassHierarchy classHierarchy =
3      bloatContext.getHeirarchy();
4   EDU.purdue.cs.bloat.editor.NameAndType nameAndType =
5      new EDU.purdue.cs.bloat.editor.NameAndType(
6         "test",
7         EDU.purdue.cs.bloat.editor.Type.STRING);
8   boolean value =
9      classHeirarchy.methodIsoverridden(
10        EDU.purdue.cs.bloat.editor.Type.OBJECT,
11        nameAndType);
```

Control Flow Graph (CFG)

Bloat allows a developer to create a control flow graph for any given method. To do so you simply need to construct a `FlowGraph` object. To create a `FlowGraph` you must provide a `MethodEditor` for the constructor the `FlowGraph`.

```
EDU.purdue.cs.bloat.cfg.FlowGraph(
   MethodEditor method);
```

To construct the `FlowGraph` object does not setup the CFG. You must call the initialize method of the `FlowGraph` object to create the graph. Once the graph is created you can extract the basic blocks using a variety of methods. Most of these methods produce a `List` or `Collection` of blocks that can be iterated across. There also several `print` methods to print the CFG in various orderings.

```
1  // printing cfg in different orders
2  EDU.purdue.cs.bloat.cfg.FlowGraph cfg =
3     new EDU.purdue.cs.bloat.cfg.FlowGraph(methodEditor);
4  cfg.preOrder();
5  cfg.postOrder();
6  cfg.print();
7  cfg.printGraph();
```

```
1  // extracting blocks and manipulating blocks:
2  EDU.purdue.cs.bloat.cfg.FlowGraph cfg =
3     new EDU.purdue.cs.bloat.cfg.FlowGraph(methodEditor);
4  //Blocks returned in trace order:
5  java.util.List list =
6     cfg.trace();
7  for(int i=0;i<list.size();i++){
8     EDU.purdue.cs.bloat.tree.Tree tree =
9        list.tree();
10    EDU.purdue.cs.bloat.editor.Instruction instruction =
11       new EDU.purdue.cs.bloat.editor.Instruction(
12          EDU.purdue.cs.bloat.editor.Opcode.ops_aaload);
13    tree.addInstruction(instruction);
14 }
```

The above example extracts the blocks from a CFG. It then accesses the Tree for each Block. A Tree represents the expression tree for each block. Using the tree you can view the instructions contained with in a Block as well as edit them by adding and deleting instructions. The example above demonstrates how to add an instruction.

Trees

Trees were previously introduced in the section on `FlowGraph`. `Tree` represents the expression tree of a basic block. It consists of an operand (expression) stack comprised of expressions and a list of statements contained in the block. The blocks, as previously mentioned, can be obtained from the `FlowGraph`. `Block` contains a method called `tree` that takes no arguments and returns a `Tree`. `Tree` allows you to manipulate the expression tree to obtain a `Tree`. Instructions may be added with the methods below.

```
1  // adds an instruction that does not change control flow
2  addInstruction(Instruction inst)
3  // adds an instruction to jump to another Block in the FlowGraph
4  addInstruction(Instruction inst, Block next)
5  // add an instruction such as ret or astore to invoke a subroutine
6  addInstruction(Instruction inst, Subroutine sub)
```

A branch in the CFG maybe labeled using the `addLabel(Label)` method. To create a label follow the example below. The index used in the example represents the labels in the instruction array. The boolean in the constructor specifies whether the `Label` starts a block or not. You may also add a comment to a `Label` by calling the `setComment` method which takes a string to represent the comment.

```
2  // The index represents represents a unique offset, .i.e. its offset
3  // in the instruction array
```

```
4   int index = 23;
5   EDU.purdue.cs.bloat.editor.Label label = new EDU.purdue.cs.bloat.editor.Label(index,false);
6   label.setComment("Hello");
7   System.out.println(label);
8   tree.addLabel(label);
```

An additional `Stmt` can also be added to a `Tree` using the method `addStmt(Stmt stmt)`. A `Stmt` is the super class of `JumpStmt`, `ExprStmt`, and `PhisStmt`. A `JmpStmt` is a super class of the following: `GotoStmt`,`IfStmt`,`JsrStmt`. A `GotoStmt` represents an unconditional branch to a basic block. The constructor takes a `Block` as an argument which represents the target of the `GotoStmt`. An `IfStmt` is a super class of statements in which some expression is evaluated and one of two branches is taken. Its subclasses are `IfCmpStmt` and `IfZeroStmt`. To create an `IfCmpStmt` the constructor takes 5 arguments. The first is an integer that represents the comparison operator the if statement. Next it takes a left `Expr` and right `Expr`. Finally it takes a `Block` to jump to when the if condition is true and then a `Block` to jump to if the condition is false. An `IfZeroStmt` evaluates an expression and executes one of its two branches depending on whether or not the expression evaluated to zero. To create a `IfZeroStmt` you simply use the same arguement sequence as supplied to a `IfCmpStmt`. You should take note that a `Stmt` can not be nested as `Expr` can. These are not all of the `Stmt` supplied by BLOAT. This is only an introduction into what they offer. `Stmt`s can be removed from a `Tree` using the method `removeStmt(Stmt stmt)`.

SSA (Value Graphs)

A `FlowGraph` object may also be used to create a value graph. The SSA graph (also called the value graph) represents the nesting of expression in a control flow graph. Each node in the SSA graph represents an expression. If the expression is a definition, the it is labeled with the variable it defines. Each node has directed edges to the nodes representing its operands.

`SSAGraph` is a representation of the definitions found in a CFG in the following form: Each node in the graph is an expression that defines a variable (a `VarExpr`, `PhiStmt`, or a `StackManipStmt`). Edges in the graph point to the nodes whose expressions define the operands of the expression in the source node.

This class is used primarily get the strongly connected components of the SSA graph in support of value numbering and induction variable analysis.

Nate (One of the Authors of BLOAT) warns: Do not modify the CFG while using the SSA graph! The effects of such modification are undefined and will probably lead to nasty things occurring.

```
1   EDU.purdue.cs.bloat.ssa.SSAGraph ssaGraph =
2      new EDU.purdue.cs.bloat.ssa.SSAGraph(cfg);
```

Type-Based Alias Analysis (TBBA)

Type-Based Alias Analysis (TBAA) is used to determine if one expression can alias another. An expression may alias another expression if there is the possibility that they refer to the same location in memory. BLOAT models expressions that may reference memory locations with `MemRefExprs`. There are two kinds of "access expressions" in Java: field accesses (`FieldExpr` and `StaticFieldExpr`) and array references (`ArrayRefExpr`). Access paths consist of one or more access expressions. For instance, `a.b[i].c` is an access expression.

TBAA uses the `FieldTypeDecl` relation to determine whether or not two access paths may refer to the same memory location.

A static method in the `EDU.purdue.cs.bloat.tbaa.TBAA` class may be used to do Type-Based Alias Analysis.

```
1    EDU.purdue.cs.bloat.tbaa.TBAA.canAlias(EditorContext context, Expr a, Expr b)
```

Where the `EditorContext` represents the `BloatContext` and the `Expr` parameters represent expressions. `EDU.purdue.cs.bloat.tree.Expr` is a an abstract superclass. Its known subclasses are listed below.

```
1   ArithExpr
2   ArrayLengthExpr
3   CallExpr
4   CastExpr
5   CatchExpr
6   CheckExpr
7   CondExpr
8   ConstantExpr
9   DefExpr
10  NegExpr
11  NewArrayExpr
12  NewExpr
13  NewMultiArrayExpr
14  ReturnAddressExpr
15  ShiftExpr
16  StoreExpr
```

Appendix B

Working with the SandMark Code-base

B.1 Coding standard

- Don't use tab characters.
- Indent by typing three (3) blanks.
- Put the left brace ({) on same line as preceding statement. I.e. format an if-statement like this:

```
if () {
   ...
}
```

  not like this:

```
if ()
{
  ...
}
```

  or like this:

```
if ()
   {
     ...
   }
```

- Don't use `import`-statements. Instead, always use fully qualified names. For example, say

```
java.util.Properties p = new java.util.Properties();
```

  instead of

```
import java.util.*;
...
Properties p = new Properties();
```

  In a large system like SandMark it is difficult to read the code when you don't know where the a particular name is declared. In SandMark we use deep package hierarchies to organize the code and always refer to every object using its fully qualified name. This also prevents name clashes.

- We make one exception: you are allowed to say `String` rather than `java.lang.String`!
- We use the following naming strategy:
  - Package names are short and in lower case. We favor deep package hierarchies over packages with many classes.
  - Class names are typically long and descriptive. They start with an uppercase letter.
  - Method names start with a lowercase letter.

B.2 Using CVS

CVS is a *source code control system*. It is used extensively in industry and in the open source community to organize source code that many people are working on simultaneously.

The basic idea is to have one master copy of the source code, residing in a special source code repository. Programmers check out the latest version of the code to their local account and work on it until they're ready to share their new code with their co-workers. They then check in the new code to the repository from which the rest of the team can download the new changes.

A programmer can have several versions of the code checked out at the same time: one at work, one at home, one on the laptop, etc. Furthermore, every change ever made to a file is stored in the repository allowing a programmer to check out "the version of the program from last Monday."

Installing CVS

If you are running Linux on your home machine (and why wouldn't you be?) you should have CVS installed already. Otherwise you can download CVS from here: `http://www.cvshome.org/downloads.html`. The manual is here: `http://www.cvshome.org/docs/manual`. ⌜`man cvs`⌝ gives you the basic information you need.

You also need to have `ssh` installed on your machine in order to communicate with our CVS server, `cvs.cs.arizona.edu`.

If you just intend to do your assignments on `lectura`, no installation is necessary.

Getting Started

Let's assume that your team consists of Alice and Bob whose `lectura` logins are `alice` and `bob` with the passwords `alice-pw` and `bob-pw`, respectively. The team name is `cs453bd`.

One team member (in our case Alice) should do the following to get the team's source repository set up:

```
> whoami
alice
> mkdir ass1            # Create a temorary directory.
> cd ass1
> setenv CVS_RSH ssh    # Or export CVS_RSH=ssh. Must always be set.
> cvs -d :ext:alice@cvs.cs.arizona.edu:/cvs/cvs/cs453/cs453bd import -m "New!" ass1 aaa bbb ass1
   alice@cvs.cs.arizona.edu's password: alice-pw
```

Now, Alice deletes the temporary directory:

```
> rmdir ass1
```

Checking out code

Everything should now be set up properly on the CVS server. Alice can check out the code (which so-far only consists of a single directory):

```
> cvs -d :ext:alice@cvs.cs.arizona.edu:/cvs/cvs/cs453/cs453bd checkout ass1
   alice@cvs.cs.arizona.edu's password: alice-pw
cvs server: Updating ass1
> ls ass1
CVS
```

Alice now wants to start programming. She creates a new C module in her CVS directory:

```
> cd ass1
/home/alice/ass1
> cat > interpreter.c
main() {
}
> cvs add -m "Started the project" interpreter.c
   alice@cvs.cs.arizona.edu's password: alice-pw
   cvs server: scheduling file 'interpreter.c' for addition
   cvs server: use 'cvs commit' to add this file permanently

> cvs commit -m "Finished first part of interpreter."
   cvs commit: Examining .
   alice@cvs.cs.arizona.edu's password:
   RCS file: /cvs/cvs/cs453/cs453bd/ass1/interpreter.c,v
   done
   Checking in interpreter.c;
   /cvs/cvs/cs453/cs453bd/ass1/interpreter.c,v  <--  interpreter.c
   initial revision: 1.1
   done
```

The `add` command told the CVS system that a new file is being created. The `commit` command actually uploaded the new file to the repository.

Now Alice realizes that she needs to add some more code the project:

```
> emacs interpreter.c
> cat interpreter.c
main() {
  int i;
}
> cvs commit -m "Added more code."
   cvs commit: Examining .
   alice@cvs.cs.arizona.edu's password: alice-pw
   Checking in interpreter.c;
   /cvs/cvs/cs453/cs453bd/ass1/interpreter.c,v  <--  interpreter.c
   new revision: 1.2; previous revision: 1.1
   done
```

OK, so what about Bob? Well, he decides he should also contribute to the project, so he checks out the source:

```
> cvs -d :ext:bob@cvs.cs.arizona.edu:/cvs/cvs/cs453/cs453bd checkout ass1
   bob@cvs.cs.arizona.edu's password: bob-pw
cvs server: Updating ass1
> cd ass1
> ls
CVS interpreter.c
> cat interpreter.c
main() {
```

```
  int i;
}
> emacs interpreter.c
> cat interpreter.c
main() {
  int i=5;
}
> cvs commit -m "Added more stuff to the project."
```

Alice has now gone back to her dorm-room where she wants to continue working on the project on her home computer. She has installed CVS and she has added

```
> setenv CVS_RSH ssh    # Or export CVS_RSH=ssh
```

to her `.cshrc` file to make sure that she runs this command every time. Now she can go ahead and check out the code again, this time on the home machine:

```
> cvs -d :ext:alice@cvs.cs.arizona.edu:/cvs/cvs/cs453/cs453bd checkout ass1
   alice@cvs.cs.arizona.edu's password: alice-pw
> cat ass1/interpreter.c
main() {
  int i=5;
}
```

Notice that she got the code that Bob checked in to CVS!

Alice can continue working on the code from home. When she's done for the day she uses the `commit` command to submit her changes to the cvs database.

Updating

The next day Bob is getting ready to work on the project again. In case Alice has made some changes to the code, he runs the `update` command:

```
> cvs update -d
```

Any files that have changed since the last time Bob worked on the project will be downloaded from the server. Bob makes his edits, then runs `commit` when he is done to upload the changes to the repository.

Deleting files

If Alice needs to delete a file she runs the CVS `rm` command:

```
> rm interpreter.c
> cvs rm interpreter.c
> cvs commit
```

Note that you have to delete the file before you can run the `cvs rm` command.

Summary

These are the most common CVS commands:

cvs add file Add a new file to the project. The file will not actually be uploaded to the repository until you run the `commit` command.

cvs rm file Remove a file from the project. The file will not actually be removed from the repository until you run the `commit` command.

cvs commit Update the repository with any changed files.

cvs update -d Download any changed files to your local machine.
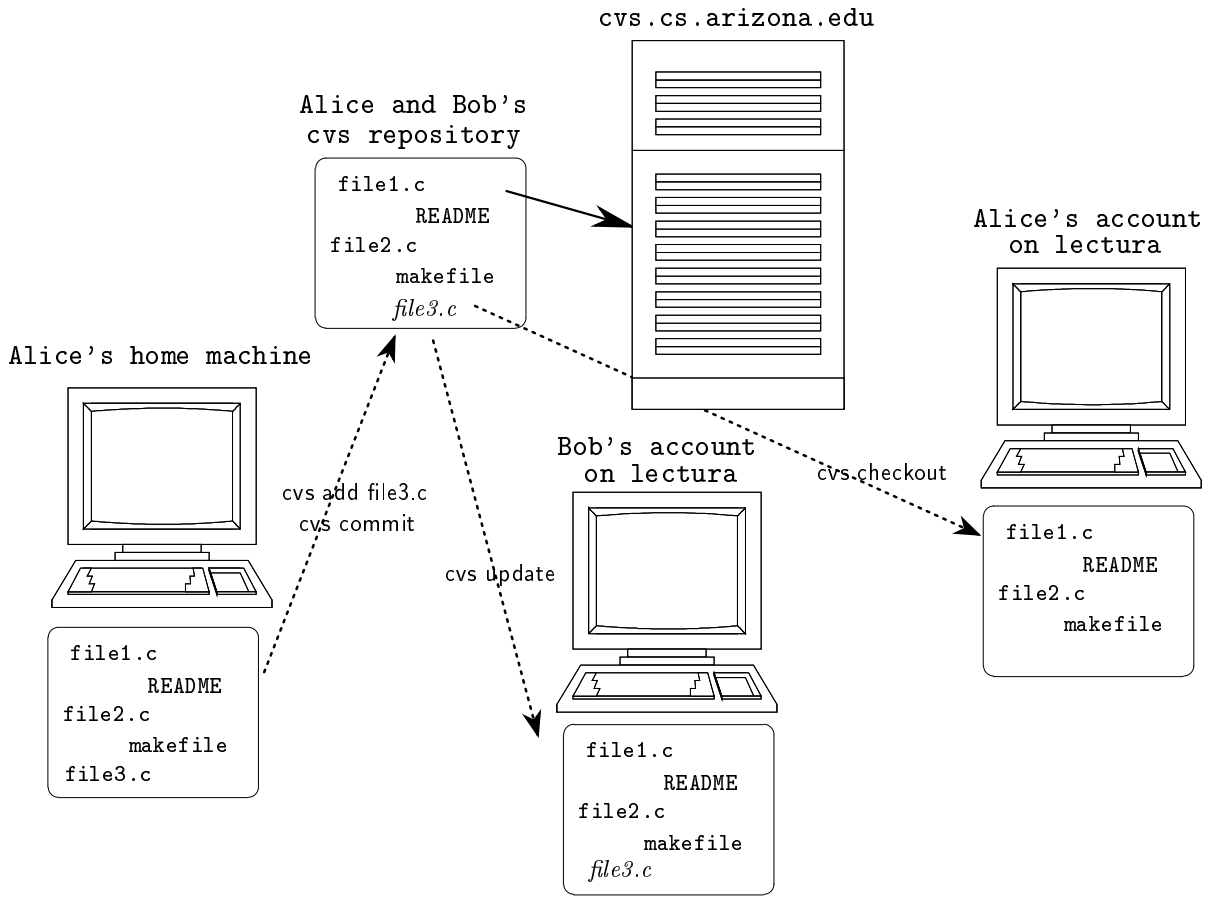
The figure below describes a typical situation. Alice and Bob have three versions of the code checked out: two on their lectura accounts, and one version in Alice's home machine. Alice adds a new file `file3.c` and checks it in to the repository. To see the new file, Bob has to run the `update` command.