The Major Mutation Framework

Version 1.3.0

December 17, 2016

Contents

1 Overview

MAJOR is a complete mutation analysis framework that divides the mutation analysis process into two individual, consecutive steps:

1. Generate and embed mutants during compilation.
2. Run the actual mutation analysis (e.g., to assess test-suite quality).

For the first step, MAJOR provides a lightweight mutator, which is integrated in the openjdk Java compiler. For the second step, MAJOR provides a default analysis back-end that extends Apache Ant's JUnit task.

1.1 Installation

This section describes how to install MAJOR for use from the command line or Ant. The installation is simple — all you need is included in the MAJOR release package!

- Download the MAJOR framework from `http://mutation-testing.org/major.zip`.
- Unzip `major.zip` to create the `major` directory.
- Optionally, update your environment and prepend MAJOR's `bin` directory to your execution path (`PATH` variable).

The `major` directory provides the following content:

```
major
├── bin ..... Executables for MAJOR's components
│   ├── ant
│   ├── javac
│   └── mmlc
├── config ..... Archive and sources of MAJOR's driver
├── doc ..... The manual of the current version
│   └── major.pdf
├── example ..... Examples for using MAJOR with Ant and standalone
│   ├── ant
│   │   ├── ...
│   │   └── run.sh
│   ├── standalone
│   │   ├── ...
│   │   └── run.sh
│   └── runAll.sh
├── lib ..... All libraries used by MAJOR
└── mml ..... Example MML files
```

1.2 How to get started

Verify that you are using MAJOR's compiler by running `javac -version`. The output should be the following:

```
major$ javac -version
javac 1.7.0-Major-v1.3.0
```

Suppose you want to mutate a class `MyClass.java`. The following command mutates and compiles `MyClass.java` using all mutation operators:

```
major$ javac -XMutator:ALL MyClass.java
#Generated Mutants: 90 (28 ms)
```

Note that `javac` must refer to MAJOR's compiler, which always prints the number of generated mutants when the `-XMutator` flag is enabled. Additionally, MAJOR's compiler produces a log file (`mutants.log`) for all generated mutants.

All generated mutants are embedded in the compiled class files. The `example` directory provides two examples on how to use mutants to perform mutation analysis on test suites:

- `ant`: mutation analysis using Apache Ant.
- `standalone`: mutation analysis using MAJOR's driver standalone.

Execute `runAll.sh` within the `example` directory to run all examples or `run.sh` in a subdirectory for a particular example. Section 2 provides a step by step tutorial on how to use MAJOR for a project using Apache Ant, and the subsequent sections describe MAJOR's components and configuration options in detail:

- Section 3 provides details about MAJOR's compiler.
- Section 4 describes MAJOR's DSL (MML).
- Section 5 provides details about MAJOR's default mutation analysis back-end.

2 Step by Step Tutorial

This sections provides a step by step tutorial on how to use MAJOR for:

- Configure the mutant generation with MML scripts (Section 2.1).
- Generate mutants with MAJOR's compiler (Section 2.2).
- Run mutation analysis with MAJOR's back-end (Section 2.3).

You can find all files and the triangle program, which are used in this tutorial, in the `example` and `mml` directories (see Section 1.1).

2.1 Prepare and compile a Mml script

MAJOR's domain specific language (MML) supports a detailed specification of the mutation process. Suppose only return statements, relational operators, and conditional operators shall be mutated within the method `classify` of the class `triangle.Triangle`. The following MML script (`tutorial.mml`) expresses these requirements:

```
1  targetOp{
2      // Define the replacements for ROR
3      BIN(>)->{>=,!=,FALSE};
4      BIN(<)->{<=,!=,FALSE};
5      BIN(>=)->{>,==,TRUE};
6      BIN(<=)->{<,==,TRUE};
7      BIN(==)->{<=,>=,FALSE,LHS,RHS};
8      BIN(!=)->{<,>,TRUE,LHS,RHS};
9      // Define the replacements for COR
10     BIN(&&)->{==,LHS,RHS,FALSE};
11     BIN(||)->{!=,LHS,RHS,TRUE};
12     // Define the type of statement that STD should delete
13     DEL(RETURN);
14
15     // Enable the STD, COR, and ROR mutation operators
16     STD;
17     COR;
18     ROR;
19 }
20 // Call the defined operator group for the target method
21 targetOp<"triangle.Triangle::classify(int,int,int)">;
```

Listing 2.1: MML script that generates COR, ROR, and STD mutants, targeting only the `classify` method in the class `triangle.Triangle`.

Section 4 provides a detailed description of the syntax and capabilities of the domain specific language MML. The MAJOR framework provides a compiler (`mmlc`) that compiles MML scripts into a binary representation. Given the MML script `tutorial.mml`, the `mmlc` compiler is invoked with the following command:

```
major$ mmlc tutorial.mml tutorial.mml.bin
```

Note that the second argument is optional — if omitted, the compiler will add `.bin` to the name of the provided script file, by default.

2.2 Generate mutants

To generate mutants based on the compiled MML script `tutorial.mml.bin` (see Section 2.1), the compiled script has to be passed as an argument to MAJOR's compiler.

2.2.1 Compiling from the command line

Use the `-XMutator` option to mutate and compile from the command line:

```
major$ javac -XMutator=tutorial.mml.bin -d bin src/triangle/Triangle.java
#Generated Mutants: 86 (144 ms)
```

2.2.2 Compiling with Apache Ant

If the source files shall be compiled using Apache Ant, the `compile` target of the corresponding `build.xml` file needs to be adapted to use MAJOR's compiler and to provide the necessary compiler option (See Section 3.4 for further details):

```
<property name="mutOp" value=":NONE"/>
<target name="compile" depends="init" description="Compile">
    <javac srcdir="src" destdir="bin" debug="yes"
           fork="yes" executable="pathToMajor/javac">
        <compilerarg value="-XMutator${mutOp}"/>
    </javac>
</target>
```

Given the compiled `tutorial.mml.bin` script and the adapted `build.xml` file, use the following command to mutate and compile the source files:

```
major$ ant -DmutOp="=tutorial.mml.bin" compile


compile:
    [javac] Compiling 1 source file to bin
    [javac] #Generated Mutants: 86 (105 ms)

BUILD SUCCESSFUL
Total time: 1 second
```

2.2.3 Inspecting generated mutants

If mutation has been enabled (i.e., the `-XMutator` option is used), MAJOR's compiler reports the number of generated mutants. Additionally, it produces the log file `mutants.log` that contains detailed information about the generated mutants (see Section 3.2.1 for a description of the format). The following example shows the log entries for the first 3 generated mutants:

```
major$ head -3 mutants.log
1:ROR:<=(int,int):<(int,int):triangle.Triangle@classify(int,int,int):11:a <= 0 |==> a < 0
2:ROR:<=(int,int):==(int,int):triangle.Triangle@classify(int,int,int):11:a <= 0 |==> a == 0
3:ROR:<=(int,int):TRUE(int,int):triangle.Triangle@classify(int,int,int):11:a <= 0 |==> true
```

MAJOR also supports the export of generated mutants to individual source files — see 3.2.2 for more details.

2.3 Analyze mutants

The `build.xml` file has to provide a suitable `mutation.test` target to use MAJOR's mutation analysis back-end, which performs the mutation analysis for a given test suite. The following `mutation.test` target enables the mutation analysis and exports the results to `summary.csv`, `results.csv`, and `killed.csv` (see Section 5 for further details):

```
<target name="mutation.test" description="Run mutation analysis">
    <echo message="Running mutation analysis ..."/>
    <junit printsummary="false"
           showoutput="false"
           mutationAnalysis="true"
           summaryFile="summary.csv"
           resultFile="results.csv"
           killDetailsFile="killed.csv">

           <classpath path="bin"/>
           <batchtest fork="false">
               <fileset dir="test">
                   <include name="**/*Test*.java"/>
               </fileset>
           </batchtest>
    </junit>
</target>
```

Using MAJOR's version of `ant`, the following command invokes the `mutation.test` target:

```
major$ ant mutation.test

mutation.test:
     [echo] Running mutation analysis ...
    [junit] MAJOR: Mutation analysis enabled
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Run 1 ordered test to verify independence
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Preprocessing time: 0.06 seconds
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Mutants generated: 86
    [junit] MAJOR: Mutants covered:   86 (100.00%)
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Export test map to (testMap.csv)
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Run mutation analysis with 1 individual test
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: 1/1 - triangle.test.TestSuite (4ms / 86):
    [junit] MAJOR: 494 (76 / 86 / 86) -> AVG-RTPM: 5ms
    [junit] MAJOR: Mutants killed / live: 76 (76-0-0) / 10
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Summary:
    [junit] MAJOR:
    [junit] MAJOR: Analysis time:  0.5 seconds
    [junit] MAJOR: Mutation score: 88.37% (88.37%)
    [junit] MAJOR: Mutants killed / live: 76 (76-0-0) / 10
    [junit] MAJOR: Mutant executions: 86
    [junit] MAJOR: ------------------------------------------------------------
    [junit] MAJOR: Export summary of results (to summary.csv)
    [junit] MAJOR: Export run-time results (to results.csv)
    [junit] MAJOR: Export mutant kill details (to killed.csv)

BUILD SUCCESSFUL
Total time: 0 seconds
```

As configured in the `build.xml` file, the results of the mutation analysis are exported to the files `summary.csv`, `killed.csv`, and `results.csv`, which provide the following information:

- `summary.csv`: Summary of mutation analysis results.
- `killed.csv`: The reason why a mutant was killed — i.e., assertion failure, exception, or timeout.
- `results.csv`: Detailed runtime information and mutation analysis results for each executed test.

3 The Mutation Compiler (Major-Javac)

MAJOR extends the *OpenJDK* Java compiler and implements conditional mutation as an optional transformation of the abstract syntac tree (AST). In order to generate mutants, this transformation has to be enabled by setting the compiler option `-XMutator` — if this flag is not set, then the compiler works exactly as if it were unmodified. The compile-time configuration of conditional mutation and the necessary runtime driver are externalized to avoid dependencies and to provide a non-invasive tool. This means that MAJOR's compiler can be used as a compiler replacement in any Java-based development environment.

3.1 Configuration

MAJOR extends the non-standardized `-X` options to avoid potential conflicts with future releases of the Java compiler. To use the mutation capabilities of MAJOR's compiler, the conditional mutation transformation has to be generally enabled at compile-time using the compiler option `-XMutator`. MAJOR's compiler supports (1) compiler sub-options and (2) mutation scripts (use `javac -X` to see a description of all configuration options):

(1) `javac -XMutator:<sub-options>`

(2) `javac -XMutator=<mml filename>`

If the mutation step is enabled, MAJOR's compiler prints the number of generated mutants at the end of the compilation process and produces the log file `mutants.log`, which contains detailed information about each generated and embedded mutant.

3.1.1 Compiler options

MAJOR's compiler provides wildcards and a list of valid sub-options, which correspond to the names of the available mutation operators. For instance, the following three commands enable (1) all operators, using the wildcard ALL, (2) all but one operator (-LVR), and (3) a custom subset of operators, namely AOR, ROR, and STD:

(1) `javac -XMutator:ALL ...`

(2) `javac -XMutator:ALL,-LVR ...`

(3) `javac -XMutator:AOR,ROR,STD ...`

Table 3.1 summarizes the mutation operators that are provided by MAJOR's compiler.

3.1.2 Mutation scripts

Instead of using compiler options, MAJOR's compiler can interpret mutation scripts written in its domain specific language MML. These MML scripts enable a detailed definition and

Table 3.1: Mutation operators implemented in MAJOR.

Mutation operator		Example
AOR	(Arithmetic Operator Replacement)	`a + b` ⟼ `a - b`
LOR	(Logical Operator Replacement)	`a ^ b` ⟼ `a | b`
COR[1]	(Conditional Operator Replacement)	`a || b` ⟼ `a && b`
ROR	(Relational Operator Replacement)	`a == b` ⟼ `a >= b`
SOR	(Shift Operator Replacement)	`a >> b` ⟼ `a << b`
ORU	(Operator Replacement Unary)	`-a` ⟼ `~a`
EVR	(Expression Value Replacement) Replaces an expression (in an otherwise unmutated statement) with a default value.	`return a` ⟼ `return 0` `int a = b` ⟼ `int a = 0`

LVR	(Literal Value Replacement)
	Replaces a literal with a default value:
	• A numerical literal is replaced with a positive number, a negative number, and zero.
	• A boolean literal is replaced with its logical complement.
	• A String literal is replaced with the empty String.
		`0` ⟼ `1`
		`1` ⟼ `-1`
		`1` ⟼ `0`
		`true` ⟼ `false`
		`false` ⟼ `true`
		`"Hello"` ⟼ `""`

STD	(STatement Deletion)
	Deletes (omits) a single statement:
	• `return` statement
	• `break` statement
	• `continue` statement
	• Method call
	• Assignment
	• Pre/post increment
	• Pre/post decrement
		`return a` ⟼ `<no-op>`
		`break` ⟼ `<no-op>`
		`continue` ⟼ `<no-op>`
		`foo(a,b)` ⟼ `<no-op>`
		`a = b` ⟼ `<no-op>`
		`++a` ⟼ `<no-op>`
		`--a` ⟼ `<no-op>`

[1]Also mutates atomic boolean conditions to true and false (e.g., if(flag) or if(isSet())).

a flexible application of mutation operators. For example, the replacement list for every operator in an operator group can be specified and mutations can be enabled or disabled for certain packages, classes, or methods. Within the following example, the mutation process is controlled by the definitions of the compiled script file `myscript.mml.bin`:

- `javac -XMutator="pathToFile/myscript.mml.bin" ...`

Note that MAJOR's compiler interprets pre-compiled script files. Use the script compiler `mmlc` to syntactically and semantically check, and compile a MML script file. MAJOR's domain specific language MML is described in detail in the subsequent Section 4.

3.2 Logging and exporting generated mutants

3.2.1 Log file for generated mutants

MAJOR's compiler generates the log file `mutants.log`, which provides detailed information about the generated mutants and uses a colon (`:`) as separator. The log file contains one row per generated mutant, where each row in turn contains 7 columns with the following information:

1. Mutants' unique number (id)
2. Name of the applied mutation operator
3. Original operator symbol
4. Replacement operator symbol
5. Fully qualified name of the mutated method
6. Line number in original source file
7. Visualization of the applied transformation (`from |==> to`)

The following example gives the log entry for a ROR mutation that has the mutant id 11 and is generated for the method `classify` (line number 18) of the class `Triangle`:

```
11:ROR:<=(int,int):<(int,int):Triangle@classify:18:a <= 0 |==> a < 0
```

3.2.2 Source-code export of generated mutants

MAJOR also supports the export of each generated mutant to an individual source file — this feature is disabled by default. If enabled, MAJOR duplicates the original source file for each mutant, injects the mutant in the copy, and exports the resulting faulty copy. MAJOR reads the following two properties that control the export of generated mutants (default values are given in parentheses):

- `-J-Dmajor.export.mutants=[true|false] (false)`
- `-J-Dmajor.export.directory=<directory> (./mutants)`

MAJOR automatically creates the export directory and parent directories if necessary.

Note that, if you are mutating a large code base, exporting all mutants to individual source files increases the compilation time and requires significantly more disk space than the log file.

3.3 Driver class

MAJOR references an external driver at runtime to gain access to a mutant identifier (`M_NO`) and a method that monitors mutation coverage (`COVERED`). Listing 3.1 shows an example of a simple driver class that provides both the mutant identifier and the mutation coverage method. Note that the mutant identifier and the coverage method must be implemented in a static context to avoid any overhead caused by polymorphism and instantiation.

```
1  public class Config {
2      public static int M_NO=0;
3      public static Set<Integer> covSet = new TreeSet<Integer>();
4
5      // Record coverage information
6      public static boolean COVERED(int from, int to) {
7          synchronized (covSet) {
8              for (int i=from; i<=to; ++i) {
9                  covSet.add(i);
10             }
11         }
12         return false;
13     }
14     // Reset the coverage information
15     public static void reset() {
16         synchronized (covSet) {
17             covSet.clear();
18         }
19     }
20     // Get (copied) list of all covered mutants
21     public static List<Integer> getCoverageList() {
22         synchronized (covSet) {
23             return new ArrayList<Integer>(covSet);
24         }
25     }
26 }
```

Listing 3.1: Driver class providing the mutant identifier `M_NO` and coverage method `COVERED`.

The archive and source files of the default driver implementation is provided in the `config` directory. Note that the driver class does **not** have to be available on the classpath during compilation. MAJOR does not try to resolve the driver class at compile-time but instead assumes that the mutant identifier and the coverage method will be provided by the driver class at runtime. Thus, MAJOR's compiler is non-invasive and the mutants can be generated without having a driver class available during compilation.

3.4 Integration into apache Ant's build.xml

MAJOR's compiler can be used standalone, but also in build systems, such as Apache Ant. Consider, for example, the following `compile` target in an Apache Ant build.xml file:

```
<target name="compile" depends="init" description="Compile">
    <javac srcdir="src"
          destdir="bin">
    </javac>
</target>
```

To use MAJOR's compiler without any further changes to your environment, add the following 3 options to the `compile` target:

```
<property name="mutOp" value=":NONE"/>
<target name="compile" depends="init" description="Compile">
    <javac srcdir="src"
          destdir="bin"

             fork="yes"
       executable="pathToMajor/bin/javac">
       <compilerarg value="-XMutator${mutOp}"/>
    </javac>
</target>
```

There is no need to duplicate the entire target since MAJOR's compiler can also be used for regular compilation. The following three commands illustrate how the `compile` target shown above can be used to: (1) compile without mutation, (2) compile with mutation using compiler options, and (3) compile with mutation using a MML script:

(1) `ant compile`

(2) `ant -DmutOp=":ALL" compile`

(3) `ant -DmutOp="=pathToFile/myscript.mml.bin" compile`

Note that the `mutOp` property provides a default value (`:NONE`) if this property is not set on the command line.

4 The Major Mutation Language (Major-Mml)

MAJOR is designed to support a wide variety of configurations by means of its own domain specific language, called MML. Generally, a MML script contains a sequence of an arbitrary number of statements, where a statement represents one of the following entities:

- Variable definition
- Replacement definition
- Definition of statement types for the STD operator
- Definition of literal types for the LVR operator
- Definition of a mutation operator group
- Invocation of a mutation operator (group)
- Line comment

While the first five statements are terminated by a semicolon, an operator definition is encapsulated by curly braces and a line comment is terminated by the end-of-line.

4.1 Statement scopes

MML provides statement scopes for replacement definitions and operator invocations to support the mutation of a certain package, class, or method within a program. Figure 4.1 depicts the definition of a statement scope, which can cover software units at different levels of granularity — from a specific method up to an entire package. Note that a statement scope is optional as indicated by the first rule of Figure 4.1. If no statement scope is provided, the corresponding replacement definition or operator call is applied to the root package. The scope's corresponding entity, that is package, class, or method, is determined by means of its fully qualified name, which is referred to as flatname. Such a flatname can be either provided within delimiters (quotes) or by means of a variable identified by IDENT.

Figure 4.2 shows the grammar rules for assembling a flatname. The naming conventions for valid identifiers (IDENT) are based on those of the Java programming language due to the fact that a flatname identifies a certain entity within a Java program. The following four examples show valid flatnames for a package, a class, a set of overloaded methods, and a particular method:

- `"java.lang"`
- `"java.lang.String"`

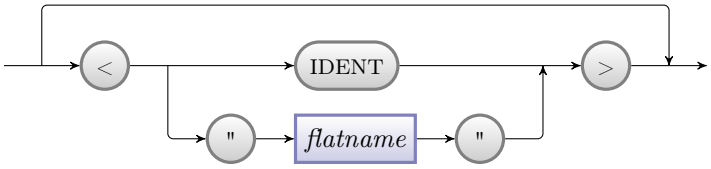

Figure 4.1: Syntax diagram for the definition of a statement scope.

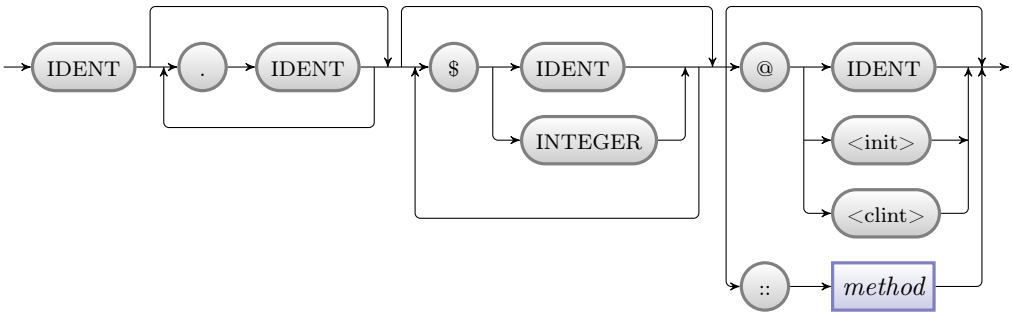

Figure 4.2: Syntax diagram for the definition of a flatname.

- `"java.lang.String@substring"`
- `"java.lang.String::substring(int,int)"`

Note that the syntax definition of a flatname also supports the identification of innerclasses and constructors, consistent with the naming conventions of Java. For Example, the subsequent definitions address an inner class, a constructor, and a static class initializer:

- `"foo.Bar$InnerClass"`
- `"foo.Bar@<init>"`
- `"foo.Bar@<clinit>"`

4.2 Overriding and extending definitions

In principle, mutation operators can be enabled (`+`), which is the default if the flag is omitted, or disabled (`-`) and this behavior can be defined for each scope. In the following example, the AOR mutation operator is generally enabled for the package `org` but, at the same time, disabled for the class `Foo` within this package:

```
+AOR<"org">;
-AOR<"org.Foo">;
```

Note that the flag for enabling or disabling operators is optional — the default flag (`+`) for enabling operators improves readability but can be omitted.

With regard to replacement definitions, there are two different possibilities: Individual replacements can be added (`+`) to an existing list or the entire replacement list can be overridden (`!`), where the latter represents the default case if this optional flag is omitted. The following example illustrates this feature, where the general definition of replacements for the package `org` is extended for the class `Foo` but overriden for the class `Bar`. The replacement lists that are effectively applied to the package and classes are given in comments.

```
 BIN(*)<"org">     -> {+,/};    // * -> {+,/}
+BIN(*)<"org.Foo"> -> {%};      // * -> {+,/,%}
!BIN(*)<"org.Bar"> -> {-};      // * -> {-}
```

4.3 Operator groups

To prevent code duplication due to the repetition of equal definitions for several scopes (i.e., the same replacements or enabled mutation operators for several packages, classes, or methods), MML provides the possibility to declare own operator groups. Such a group may in turn contain any statement that is valid in the context of the MML, except for a call of another operator group. An operator group is defined by means of a unique identifier and its statements are enclosed by curly braces, as shown in the following example:

```
myGroup {
    BIN(*) -> {+,/};
    AOR;
}
```

4.4 Script examples

Listing 4.1 shows a simple example of a mutation script that includes the following tasks:

- Define specific replacement lists for AOR and ROR
- Invoke the AOR and ROR operators on reduced lists
- Invoke the LVR operator without restrictions

The more enhanced script in Listing 4.2 exploits the scoping capabilities of MML in line 8 and 13-20, and takes, additionally, advantage of the possibility to define a variable in line 11 to avoid code duplication in the subsequent scope declarations. Both features are useful if only a certain package, class, or method shall be mutated in a hierarchical software system.
Finally, the example in Listing 4.3 visualizes the grouping feature, which is useful if the same group of operations (replacement definitions or mutation operator invocations) shall be applied to several packages, classes, or methods.

```
// Define own replacement list for AOR
BIN(*) -> {/,%};
BIN(/) -> {*,%};
BIN(%) -> {*,/};

// Define own replacement list for ROR
BIN(>)  -> {<=,!=,==};
BIN(==) -> {<,!=,>};

// Define types of literals that should be mutated by the LVR operator.
// Literal type is one of {BOOLEAN, NUMBER, STRING}.
LIT(NUMBER);
LIT(BOOLEAN);

// Enable and invoke mutation operators
AOR;
ROR;
LVR;
```

Listing 4.1: Mml script that 1) defines replacements for the AOR and ROR mutation operators, 2) defines the types of the literals that should be mutated by the LVR mutation operator, and 3) enables AOR, ROR, and LVR on the root node.

```
// Definitions for the root node
BIN(>=)->{TRUE,>,==};
BIN(<=)->{TRUE,<,==};
BIN(!=)->{TRUE,<,> };
LIT(NUMBER);
LVR;

// Definition for the package org
ROR<"org">;

// Variable definition for the class Foo
foo="org.x.y.z.Foo";

// Scoping for replacement lists
BIN(&&)<foo>->{LHS,RHS,==,FALSE};
BIN(||)<foo>->{LHS,RHS,!=,TRUE };

// Scoping for mutation operators
-LVR<foo>;
ROR<foo>;
COR<foo>;
```

Listing 4.2: Enhanced Mml script that uses scoping and a variable definition.

```
myOp{
    // Definitions for the operator group
    BIN(>=)->{TRUE,>,==};
    BIN(<=)->{TRUE,<,==};
    BIN(!=)->{TRUE,<,> };
    BIN(&&)->{LHS,RHS,==,FALSE};
    BIN(||)->{LHS,RHS,!=,TRUE };
    // Mutation operators enabled in this group
    ROR;
    COR;
}

// Calls of the defined operator group
myOp<"org">;
myOp<"de">;
myOp<"com">;
```

Listing 4.3: Mml script that defines a mutation operator (myOp) and applies this operator to different scopes (i.e., the packages org, de, and com).

5 The Mutation Analysis Back-end (Major-Ant)

MAJOR provides a default back-end for mutation analysis, which extends the Apache Ant junit task. Therefore, this back-end can be used to evaluate existing JUnit tests. Note that MAJOR does currently not support forking a JVM when executing JUnit tests, meaning that the `fork` option must not be set to `true` — forking will be supported in a future release.

5.1 Setting up a mutation analysis target

Most software projects that are build with Apache Ant provide a `test` target, which executes the corresponding unit tests. Even if no such target exists, it can be easily set up to execute a set of given unit tests. The following code snippet shows an exemplary `test` target (See http://ant.apache.org/manual/Tasks/junit.html for a detailed description of the options used in the `junit` task):

```
<target name="test" description="Run all unit test cases">
    <junit printsummary="true"
             showoutput="true"
          haltonfailure="true">

        <formatter type="plain" usefile="true"/>
        <classpath path="bin"/>
        <batchtest fork="no">
            <fileset dir="test">
                <include name="**/*Test*.java"/>
            </fileset>
        </batchtest>
    </junit>
</target>
```

To enable mutation analysis in MAJOR's enhanced version of the junit task, the option `mutationAnalysis` has to be set to `true`. For the sake of clarity, it is advisable to duplicate an existing `test` target, instead of parameterizing it, and to create a new target, e.g., called `mutation.test` (See Section 5.3 for recommended configurations):

```
<target name="mutation.test" description="Run mutation analysis">
    <junit printsummary="false"
             showoutput="false"
          haltonfailure="true"

       mutationAnalysis="true"
            summaryFile="summary.csv"
             resultFile="results.csv"
        killDetailsFile="killed.csv">

        <classpath path="bin"/>
        <batchtest fork="no">
            <fileset dir="test">
                <include name="**/*Test*.java"/>
            </fileset>
        </batchtest>
    </junit>
</target>
```

5.2 Configuration options for mutation analysis

MAJOR enhances the junit task with additional options to control the mutation analysis process. The available, additional, options are summarized in Table 5.1.

5.3 Performance optimization

During the mutation analysis process, the provided JUnit tests are repeatedly executed. For performance reasons, consider the following advices when setting up the mutation analysis target for a JUnit test suite:

- Turn off logging output (options `showsummary`, `showoutput`, etc.)
- Do not use result formatters (nested task `formatter`, especially the `usefile` option)

For performance reasons, especially due to frequent class loading and thread executions, the following JVM options are recommended:

- `-XX:ReservedCodeCacheSize=128M`
- `-XX:MaxPermSize=256M`

Table 5.1: Additional configuration options for Major-Ant's JUnit task

	Description	Values	Default
mutationAnalysis	Enable mutation analysis	`[true|false]`	`false`
mutantsLogFile	Mutants log file produced my the mutation compiler	`<String>`	`mutants.log`
coverage	Enable mutation coverage	`[true|false]`	`true`
timeoutFactor	Base timeout factor for test runtime	`<int>`	`8`
sort	Enable sort of test cases	`[original|random|sort_classes|sort_methods]`	`original`
excludeFailingTests	Exclude all failing tests (if haltonfailure is set to false)	`[true|false]`	`true`
excludeFile	Exclude mutants whose ids are listed in this file (1 id per row)	`<String>`	`null`
summaryFile	Export summary of results to this file (csv)	`<String>`	`summary.csv`
resultFile	Export detailed runtime information to this file (csv)	`<String>`	`null`
killDetailsFile	Export kill details for each mutant to this file (csv)	`<String>`	`null`
filterDetailsFile	Export filtering details for each mutant to this file (csv)	`<String>`	`null`
exportCovMap	Export mutation coverage map	`[true|false]`	`false`
covMapFile	File name for mutation coverage map (csv)	`<String>`	`covMap.csv`
exportKillMap*	Export mutation kill map	`[true|false]`	`false`
killMapFile	File name for mutation kill map	`<String>`	`killMap.csv`
testMapFile	File name for mapping of test id to test name (csv)	`<String>`	`testMap.csv`

*Note: this option leads to the execution of every test on every covered mutant!