

The Checker Framework: Custom pluggable types for Java

<http://types.cs.washington.edu/checker-framework/>

Version 1.4.0-b1 (1 Sep 2012)

For the impatient: Section 1.2 (page 10) describes how to **install and use** pluggable type-checkers.

Contents

1	Introduction	9
1.1	How it works: Pluggable types	9
1.2	Installation	10
1.3	Example use: detecting a null pointer bug	11
2	Using a checker	12
2.1	Writing annotations	12
2.1.1	Distributing your annotated project	13
2.2	Running a checker	13
2.2.1	Summary of command-line options	13
2.2.2	Checker auto-discovery	14
2.3	What the checker guarantees	14
2.4	Tips about writing annotations	15
2.4.1	How to get started annotating legacy code	15
2.4.2	Do not annotate local variables unless necessary	16
2.4.3	Annotations indicate normal behavior	16
2.4.4	Subclasses must respect superclass annotations	16
2.4.5	Annotations on constructor invocations	17
2.4.6	When to use (and not use) type qualifiers	18
2.4.7	What to do if a checker issues a warning about your code	18
3	Nullness checker	19
3.1	What the Nullness checker checks	19
3.2	Nullness annotations	20
3.2.1	Nullness qualifiers	20
3.2.2	Nullness method annotations	20
3.2.3	Rawness qualifiers	21
3.2.4	Map key qualifiers	22
3.3	Writing nullness annotations	22
3.3.1	Implicit qualifiers	22
3.3.2	Default annotation	22
3.3.3	Conditional nullness	22
3.3.4	Nullness and arrays	23
3.3.5	Inference of @NonNull and @Nullable annotations	23
3.4	Suppressing nullness warnings	24
3.4.1	Suppressing warnings with assertions and method calls	24
3.4.2	Suppressing warnings on nullness-checking routines and defensive programming	25
3.5	@Raw annotation for partially-initialized objects	26
3.5.1	Rawness qualifiers	27
3.5.2	How an object becomes non-raw	27

3.5.3	More details about rawness checking	28
3.6	Map key annotations	28
3.7	Additional details	29
3.8	Examples	29
3.8.1	Tiny examples	29
3.8.2	Annotated library	29
3.9	Tips for getting started	29
3.10	Other tools for nullness checking	30
3.10.1	Which tool is right for you?	30
3.10.2	Incompatibility note about FindBugs @Nullable	31
4	Interning checker	33
4.1	Interning annotations	34
4.2	Annotating your code with @Interned	34
4.2.1	Implicit qualifiers	34
4.3	What the Interning checker checks	34
4.4	Examples	35
4.5	Other interning annotations	35
5	IGJ immutability checker	36
5.1	IGJ and Mutability	36
5.2	IGJ Annotations	36
5.3	What the IGJ checker checks	37
5.4	Implicit and default qualifiers	37
5.5	Annotation IGJ Dialect	37
5.5.1	Semantic Changes	38
5.5.2	Syntax Changes	38
5.5.3	Templating Over Immutability: @I	38
5.6	Iterators and their abstract state	39
5.7	Examples	39
6	Javari immutability checker	40
6.1	Javari annotations	40
6.2	Writing Javari annotations	41
6.2.1	Implicit qualifiers	41
6.2.2	Inference of Javari annotations	41
6.3	What the Javari checker checks	41
6.4	Iterators and their abstract state	41
6.5	Examples	41
7	Lock checker	42
7.1	Lock annotations	42
7.1.1	Examples	42
7.1.2	Discussion of @Holding	43
7.1.3	Relationship to annotations in <i>Java Concurrency in Practice</i>	44
8	Fake Enum checker	45
8.1	Fake enum annotations	45
8.2	What the Fenum checker checks	46
8.3	Running the Fenum checker	46
8.4	Suppressing warnings	46
8.5	Example	47

8.6	References	47
9	Tainting checker	48
9.1	Tainting annotations	48
9.2	Tips on writing @Untainted annotations	48
9.3	@Tainted and @Untainted can be used for many purposes	49
10	Linear checker for preventing aliasing	50
10.1	Linear annotations	50
10.2	Limitations	51
11	Regex checker for regular expression syntax	52
11.1	Regex annotations	52
11.2	Annotating your code with @Regex	52
11.2.1	Implicit qualifiers	52
11.2.2	Capturing groups	53
11.2.3	Concatenation of partial regular expressions	54
11.2.4	Testing whether a string is a regular expression	54
11.2.5	Suppressing warnings	54
12	Property file checker	56
12.1	Generic property file checker	56
12.2	Internationalization checker	57
12.2.1	Internationalization annotations	57
12.2.2	Running the Internationalization Checker	57
12.3	Compiler Message Key checker	57
13	Signature checker for string representations of types	59
13.1	Signature annotations	59
13.2	What the Signature Checker checks	60
14	Units checker	61
14.1	Units annotations	61
14.2	Extending the Units Checker	62
14.3	What the Units Checker checks	62
14.4	Running the Units Checker	63
14.5	Suppressing warnings	63
14.6	References	63
15	Basic checker	64
15.1	Using the Basic checker	64
15.2	Basic checker example	64
16	Typestate checker	67
16.1	Comparison to flow-sensitive type refinement	67
17	Third-party checkers	68
17.1	Units and dimensions checker	68
17.2	Thread locality checker	68
17.3	Safety-Critical Java checker	68
17.4	Generic Universe Types checker	68
17.5	EnerJ checker	68

18 Generics and polymorphism	69
18.1 Generics (parametric polymorphism or type polymorphism)	69
18.1.1 Restricting instantiation of a generic class	69
18.1.2 A qualifier on a type parameter is like two bounds	70
18.1.3 Examples of qualifiers on a type parameter	70
18.1.4 Defaults for bounds	71
18.1.5 Type annotations on a use of a generic type variable	71
18.1.6 Covariant type parameters	72
18.2 Qualifier polymorphism	72
18.2.1 Examples of using polymorphic qualifiers	73
18.2.2 Relationship to subtyping and generics	73
18.2.3 Using multiple polymorphic qualifiers in a method signature	73
18.2.4 Using a single polymorphic qualifier on an element type	74
18.2.5 The <code>@PolyAll</code> qualifier applies to every type system	74
19 Advanced type system features	76
19.1 The effective qualifier on a type (defaults and inference)	76
19.1.1 Default qualifier for unannotated types	76
19.1.2 Automatic type refinement (flow-sensitive type qualifier inference)	78
19.1.3 Fields and flow sensitivity analysis	79
19.1.4 Inherited defaults	79
19.2 Writing Java expressions as annotation arguments	80
19.3 Unused fields and dependent types	80
19.3.1 Unused fields	81
19.3.2 Dependent types	81
20 Handling warnings and legacy code	83
20.1 Checking partially-annotated programs: handling unannotated code	83
20.2 Suppressing warnings	84
20.2.1 <code>@SuppressWarnings</code> annotation	84
20.2.2 <code>-AskipUses</code> command-line option	85
20.2.3 <code>-AskipDefs</code> command-line option	85
20.2.4 <code>-Alint</code> command-line option	85
20.2.5 No <code>-processor</code> command-line option	85
20.2.6 Checker-specific mechanisms	86
20.3 Backward compatibility with earlier versions of Java	86
20.3.1 Annotations in comments	86
20.3.2 Implicit import statements	87
20.3.3 Migrating away from annotations in comments	87
20.3.4 Annotations in Java 5 <code>.class</code> files	88
21 Annotating libraries	89
21.1 Choosing between stub files and annotated <code>.class</code> files	89
21.2 Using stub classes	90
21.2.1 Creating a stub file	90
21.2.2 Using a stub file	90
21.2.3 Stub file format	91
21.2.4 Limitations	91
21.3 Using distributed annotated JDKs	91
21.4 Troubleshooting/debugging annotated libraries	92

22	How to create a new checker	93
22.1	Relationship of the Checker Framework to other tools	93
22.2	The parts of a checker	94
22.3	Annotations: Type qualifiers and hierarchy	94
22.3.1	Declaratively defining the qualifier and type hierarchy	95
22.3.2	Procedurally defining the qualifier and type hierarchy	95
22.3.3	Defining a default annotation	96
22.3.4	Completeness of the type hierarchy	96
22.4	Type factory: Implicit annotations	97
22.4.1	Declaratively specifying implicit annotations	97
22.4.2	Procedurally specifying implicit annotations	97
22.4.3	Flow-sensitive type qualifier inference	98
22.5	Visitor: Type rules	98
22.5.1	AST traversal	99
22.5.2	Avoid hardcoding	99
22.6	The checker class: Compiler interface	99
22.6.1	Bundling multiple checkers	100
22.6.2	Providing command-line options	100
22.7	Testing framework	101
22.8	Debugging options	101
22.9	javac implementation survival guide	102
22.9.1	Checker access to compiler information	102
22.9.2	How a checker fits in the compiler as an annotation processor	103
23	Integration with external tools	104
23.1	Javac Compiler	104
23.1.1	Unix/Linux/macOS installation	104
23.1.2	Windows installation	105
23.2	Ant task	106
23.2.1	Explanation	106
23.3	Maven plugin	107
23.4	Gradle	108
23.5	IntelliJ IDEA	109
23.6	Eclipse	109
23.7	tIDE	109
23.8	Type inference tools	110
23.8.1	Varieties of type inference	110
23.8.2	Type inference to annotate a program	110
24	Frequently Asked Questions (FAQs)	111
24.1	Are type annotations easy to read and write?	111
24.2	Will my code become cluttered with type annotations?	112
24.3	I don't make type errors, so would pluggable type checking help me?	112
24.4	What should I do if a checker issues a warning about my code?	112
24.5	Can a pluggable type-checker give an absolute guarantee of correctness?	112
24.6	How do I make compilation succeed even if a checker issues errors?	113
24.7	Will using the Checker Framework slow down my program? Will it slow down the compiler?	113
24.8	How can I do run-time monitoring of properties that were not statically checked?	113
24.9	How do I get started annotating an existing program?	114
24.10	How do I shorten the command line when invoking a checker?	114
24.11	When should I use type qualifiers, and when should I use subclasses?	114
24.12	How do I create a new checker?	114

24.13	Why is there no declarative syntax for writing type rules?	114
24.14	Why not just use a bug detector (like FindBugs)?	114
24.15	How does pluggable type-checking compare with JML?	115
24.16	Why shouldn't a qualifier apply to both types and declarations?	115
24.17	What is the meaning of array annotations such as @NonNull Object @Nullable []?	116
24.18	Why are the type parameters to List and Map annotated as @NonNull?	116
24.19	Is the Checker Framework an official part of Java?	117
25	Troubleshooting and getting help	118
25.1	Common problems and solutions	118
25.1.1	Unable to run the checker, or checker crashes	118
25.1.2	Unexpected type-checking results	120
25.1.3	Unable to build the checker, or to run programs	120
25.2	How to report problems (bug reporting)	120
25.3	Building from source	121
25.3.1	Obtain the source	121
25.3.2	Build the Type Annotations compiler	121
25.3.3	Build the Annotation File Utilities	121
25.3.4	Build the Checker Framework	122
25.3.5	Build the Checker Framework manual (this document)	122
25.4	Learning more	122
25.5	Comparison to other tools	122
25.6	Credits and changelog	123

Chapter 1

Introduction

The Checker Framework enhances Java’s type system to make it more powerful and useful. This lets software developers detect and prevent errors in their Java programs.

The Checker Framework comes with checkers for specific types of errors:

1. Nullness checker for null pointer errors (see Chapter 3, page 19)
2. Interning checker for errors in equality testing and interning (see Chapter 4, page 33)
3. IGJ checker for mutation errors (incorrect side effects), based on the IGJ type system (see Chapter 5, page 36)
4. Javari checker for mutation errors (incorrect side effects), based on the Javari type system (see Chapter 6, page 40)
5. Lock checker for concurrency and lock errors, inspired by the Java Concurrency in Practice (JCIP) annotations (see Chapter 7, page 42)
6. Fake enum checker to allow type-safe fake enum patterns (see Chapter 8, page 45)
7. Tainting checker for trust and security errors (see Chapter 9, page 48)
8. Linear checker to control aliasing and prevent re-use (see Chapter 10, page 50)
9. Regex checker to prevent use of syntactically invalid regular expressions (see Chapter 11, page 52)
10. Property file checker to ensure that valid keys are used for property files and resource bundles (see Chapter 12, page 56). Also includes a checker that code is properly internationalized.
11. Signature string checker to ensure that the string representation of a type is properly used, for example in `Class.forName` (see Chapter 13, page 59). Also includes a checker that code is properly internationalized.
12. Units checker to ensure operations are performed on correct units of measurement (see Chapter 14, page 61)
13. Basic checker for customized checking without writing any code (see Chapter 15, page 64)
14. Typestate checker to ensure operations are performed on objects that are in the right state, such as only opened files being read (see Chapter 16, page 67)
15. Third-party checkers that are distributed separately from the Checker Framework (see Chapter 17, page 68)

These checkers are easy to use and are invoked as arguments to `javac`.

The Checker Framework also enables you to write new checkers of your own; see Chapters 15 and 22.

1.1 How it works: Pluggable types

The Checker Framework supports adding pluggable type systems to the Java language in a backward-compatible way. Java’s built-in typechecker finds and prevents many errors — but it doesn’t find and prevent *enough* errors. The Checker Framework lets you run an additional typechecker as a plug-in to the `javac` compiler. Your code stays completely backward-compatible: your code compiles with any Java compiler, it runs on any JVM, and your coworkers don’t have to use the enhanced type system if they don’t want to. You can check only part of your program. Type inference tools exist to help you annotate your code.

A type system designer uses the Checker Framework to define type qualifiers and their semantics, and a compiler plug-in (a “checker”) enforces the semantics. Programmers can write the type qualifiers in their programs and use the plug-in to detect or prevent errors. The Checker Framework is useful both to programmers who wish to write error-free code, and to type system designers who wish to evaluate and deploy their type systems.

This document uses the terms “checker”, “checker plugin”, “type-checking compiler plugin”, and “annotation processor” as synonyms.

1.2 Installation

This section describes how to install the binary release of the Checker Framework. The binary release contains everything that you need, both to run checkers and to write your own checkers. As an alternative, you can build the latest development version from source (Section 25.3, page 121).

Requirement: You must have **JDK 6** or later installed. You can get JDK 6 from Oracle or elsewhere. If you are using Apple Mac OS X, you can use Apple’s implementation, SoyLatte, or the OpenJDK.

The installation process is simple!

1. Download the Checker Framework distribution (<http://types.cs.washington.edu/checker-framework/current/checker-framework-1.7.0-javadoc.jar>).
2. Unzip it to create a `checker-framework` directory.
3. Optionally, update your execution path or create an alias.

When doing pluggable type-checking, you need to use the “Type Annotations compiler”, an updated version of the OpenJDK `javac` compiler that understands type annotations. The Type Annotations compiler is backward-compatible, so using it as your Java compiler should have no negative consequences.

You can do this in three ways. You can use any one of them. However, if you are using the Windows command shell, you must use the last one. The instructions assume a `CHECKERS` environment variable that is set to the `.../checker-framework/checkers` directory, but you can use an absolute path if you prefer.

- Add directory `$(CHECKERS)/binary` to your path, *before* any other directory that contains a `javac` executable. Now, whenever you run `javac`, you will use the updated compiler. If you are using the bash shell, a way to do this is to add the following to your `~/.bashrc` file:

```
export PATH=$(CHECKERS)/binary:${PATH}
```

- Whenever this document tells you to run `javac`, you can instead run `$(CHECKERS)/binary/javac`. You can simplify this by introducing an alias. Then, whenever this document tells you to run `javac`, instead use that alias. Here is the syntax for your `~/.bashrc` file:

```
alias javacheck='$(CHECKERS)/binary/javac'
```

- Whenever this document tells you to run `javac`, instead run `java` (not `javac`), but pass an extra `-Xbootclasspath/p:` and `-jar` argument, as in:

```
java -Xbootclasspath/p:$(CHECKERS)/binary/jsr308-all.jar -jar $(CHECKERS)/binary/jsr308-all.jar ...
```

You can simplify this by introducing an alias. Then, whenever this document tells you to run `javac`, instead use that alias. For example:

```
# Unix
```

```
alias javacheck='java -Xbootclasspath/p:$(CHECKERS)/binary/jsr308-all.jar -jar $(CHECKERS)/binary/jsr308-all.jar ...'
```

```
# Windows
```

```
doskey javacheck=java -Xbootclasspath/p:%CHECKERS%\binary\jsr308-all.jar -jar %CHECKERS%\binary\jsr308-all.jar ...
```

To ensure that it was installed properly, run `javac -version` (possibly using the full pathname to `javac` or the alias, if you did not add the Type Annotations `javac` to your path).

The output should be:

```
javac 1.7.0-jsr308-1.4.0-b1
```

That's all there is to it! Now you are ready to start using the checkers.

Section 1.3 walks you through a simple example. More detailed instructions for using a checker appear in Chapter 2.

1.3 Example use: detecting a null pointer bug

To run a checker on a source file, just run `javac` as usual, passing the `-processor` flag. (You can also use an IDE or other build tool; see Chapter 23.)

For instance, if you usually run the compiler like this:

```
javac Foo.java Bar.java
```

then you will instead use the command line:

```
javac -processor ProcessorName Foo.java Bar.java
```

but take note that the `javac` command must refer to the Type Annotations compiler (see Section 1.2).

If you usually do your coding within an IDE, you will need to configure the IDE. This manual contains instructions for Ant (Section 23.2), Maven (Section 23.3), IntelliJ IDEA (Section 23.5), Eclipse (Section 23.6), and tIDE (Section 23.7). Otherwise, see your IDE documentation for details.

1. Let's consider this very simple Java class. One local variable is annotated as `NonNull`, indicating that `ref` must be a reference to a non-null object. Save the file as `GetStarted.java`.

```
import checkers.nullness.quals.*;

public class GetStarted {
    void sample() {
        @NonNull Object ref = new Object();
    }
}
```

2. Run the nullness checker on the class. Either run this command:

```
javac -processor checkers.nullness.NullnessChecker GetStarted.java
```

or compile from within your IDE, which you have customized to use the JSR 308 compiler and to pass the extra arguments.

The compilation should complete without any errors.

3. Let's introduce an error now. Modify `ref`'s assignment to:

```
@NonNull Object ref = null;
```

4. Run the nullness checker again, just as before. This run should emit the following error:

```
GetStarted.java:5: incompatible types.
found   : @Nullable <nulltype>
required: @NonNull Object
    @NonNull Object ref = null;
        ^
1 error
```

The type qualifiers (e.g., `@NonNull`) are permitted anywhere that would write a type, including generics and casts; see Section 2.1.

```
@Interned String intern() { ... } // return value
int compareTo(@NonNull String other) { ... } // parameter
@NonNull List<@Interned String> messages; // non-null list of interned Strings
```

Chapter 2

Using a checker

A pluggable type-checker enables you to detect certain bugs in your code, or to prove that they are not present. The verification happens at compile time.

Finding bugs, or verifying their absence, with a checker plugin is a two-step process, whose steps are described in Sections 2.1 and 2.2.

1. The programmer writes annotations, such as `@NonNull` and `@Interned`, that specify additional information about Java types. (Or, the programmer uses an inference tool to automatically insert annotations in his code: see Sections 3.3.5 and 6.2.2.) It is possible to annotate only part of your code: see Section 20.1.
2. The checker reports whether the program contains any erroneous code — that is, code that is inconsistent with the annotations.

This chapter is structured as follows:

- Section 2.1: How to write annotations
- Section 2.2: How to run a checker
- Section 2.3: What the checker guarantees
- Section 2.4: Tips about writing annotations

Additional topics that apply to all checkers are covered later in the manual:

- Chapter 19: Advanced type system features
- Chapter 20: Handling warnings and legacy code
- Chapter 21: Annotating libraries
- Chapter 22: How to create a new checker
- Chapter 23: Integration with external tools

2.1 Writing annotations

The syntax of type annotations in Java is specified by JSR 308 [Ern08]. Ordinary Java permits annotations on declarations. JSR 308 permits annotations anywhere that you would write a type, including generics and casts. You can also write annotations to indicate type qualifiers for array levels and receivers. Here are a few examples:

```
@Interned String intern() { ... }           // return value
int compareTo(@NonNull String other) { ... } // parameter
String toString(@ReadOnly MyClass this) { ... } // receiver ("this" parameter)
@NonNull List<@Interned String> messages;    // generics: non-null list of interned Strings
@Interned String @NonNull [] messages;      // arrays: non-null array of interned Strings
myDate = (@ReadOnly Date) readonlyObject;   // cast
```

You can also write the annotations within comments, as in `List</*@NonNull*/ String>`. The Type Annotations compiler, which is distributed with the Checker Framework, will still process the annotations. However, your code will remain compilable by people who are not using the Type Annotations compiler. For more details, see Section 20.3.1.

2.1.1 Distributing your annotated project

If your code contains annotations, then your code has a dependency on the annotation declarations. People who want to compile or run your code may need declarations of the annotations on their classpath.

- To perform pluggable type-checking, all of the Checker Framework (which also contains the annotation declarations) is needed.
- To compile the code:
 - If you wrote annotations in comments (see Section 20.3.1) and/or used implicit import statements (see Section 20.3.2), then the code can be compiled by any Java compiler, without needing declarations of the annotations.
 - Otherwise, compiling the code requires a declaration of the annotations. These appear in the full Checker Framework. Additionally, the Checker Framework distribution .zip file contains a small jar file, `checkers-quals.jar`, that only contains the definitions of the distributed qualifiers, without any support for type-checking.
- To run the code:
 - If you compiled the code without using the annotation declarations, then no annotation declarations are needed.
 - If you compiled the code using the annotation declarations, then users may need to have the annotation declarations on their classpath.

A simple rule of thumb is as follows. When distributing your source code, you may wish to include either the Checker Framework jar file or the `checkers-quals.jar` file. When distributing compiled binaries, you may wish to compile them without using the annotations, or include the contents of `checkers-quals.jar` in your distribution.

2.2 Running a checker

To run a checker plugin, run the compiler `javac` as usual, but pass the `-processor plugin_class` command-line option. (You can run a checker from within your favorite IDE or build system. See Chapter 23 for details about Ant (Section 23.2), Maven (Section 23.3), IntelliJ IDEA (Section 23.5), Eclipse (Section 23.6), and tIDE (Section 23.7), and about customizing other IDEs and build tools.) Remember that you must be using the Type Annotations version of `javac`, which you already installed (see Section 1.2).

Two concrete examples (using the Nullness checker) are:

```
javac -processor checkers.nullness.NullnessChecker MyFile.java
javac -processor checkers.nullness.NullnessChecker -Xbootclasspath/p:checkers/jdk/jdk.jar MyFile.java
```

For a discussion of the `-Xbootclasspath/p` argument, see Section 21.3.

The checker is run only on any Java file that `javac` compiles. This includes all Java files specified on the command line (or created by another annotation processor). It may also include other of your Java files (but not if a more recent `.class` file exists). Even when the checker does not analyze a class (say, the class was already compiled, or source code is not available), it does check the *uses* of those classes in the source code being compiled.

You can always compile the code without the `-processor` command-line option, but in that case no checking of the type annotations is performed. The annotations are still written to the resulting `.class` files, however.

2.2.1 Summary of command-line options

You can pass command-line arguments to a checker via `javac`'s standard `-A` option (“A” stands for “annotation”). All of the distributed checkers support the following command-line options:

- `-AskipUses` Suppress all errors and warnings at all uses of a given class; see Section 20.2
- `-AskipDefs` Suppress all errors and warnings within the definition of a given class; see Section 20.2
- `-Astubs` List of stub files or directories; see Section 21.2.2
- `-Alint` Enable or disable optional checks; see Section 20.2.4
- `-Awarns` Treat checker errors as warnings. If you use this, you may wish to also supply `-Xmaxwarns 10000`, because by default `javac` prints at most 100 warnings.
- `-Afilenames`, `-Anomsgtext`, `-Ashowchecks`, `-AprintErrorStack`, `-AprintAllQualifiers`, `-Aignorejdkastub` Aids for testing or debugging a checker; see Section 22.8

Some checkers support additional options, such as `-Aquals` for the Basic Checker to check; see Chapter 15.

Here are some standard `javac` command-line options that you may find useful. Many of them contain the word “processor”, because in `javac` jargon, a checker is a type of “annotation processor”.

- `-processor` Names the checker to be run; see Section 2.2
- `-processorpath` Indicates where to search for the checker; should also contain any qualifiers used by the Basic Checker; see Section 15.2
- `-proc:{none,only}` Controls whether checking happens; `-proc:none` means to skip checking; `-proc:only` means to do only checking, without any subsequent compilation; see Section 2.2.2
- `-Xbootclasspath/p:` Indicates where to find the annotated JDK classes; see Section 21.3
- `-implicit:class` Suppresses warnings about implicitly compiled files (not named on the command line); see Section 23.2
- `-XDTA:spacesincomments` parse annotation comments even when they contain spaces; applicable only to the Type Annotations compiler; see Section 20.3.1
- `-J` Supply an argument to the JVM that is running `javac`; example: `-J-Djsr308_imports=checkers.nullness.quals.*`; see Section 20.3.2

2.2.2 Checker auto-discovery

“Auto-discovery” makes the `javac` compiler always run a checker plugin, even if you do not explicitly pass the `-processor` command-line option. This can make your command line shorter, and ensures that your code is checked even if you forget the command-line option.

To enable auto-discovery, place a configuration file named `META-INF/services/javac.annotation.processing.Processor` in your classpath. The file contains the names of the checker plugins to be used, listed one per line. For instance, to run the Nullness and the Interning checkers automatically, the configuration file should contain:

```
checkers.nullness.NullnessChecker
checkers.interning.InterningChecker
```

You can disable this auto-discovery mechanism by passing the `-proc:none` command-line option to `javac`, which disables all annotation processing including all pluggable type-checking.

2.3 What the checker guarantees

A checker can guarantee that a particular property holds throughout the code. For example, the Nullness checker (Chapter 3) guarantees that every expression whose type is a `@NonNull` type never evaluates to null. The Interning checker (Chapter 4) guarantees that every expression whose type is an `@Interned` type evaluates to an interned value. The checker makes its guarantee by examining every part of your program and verifying that no part of the program violates the guarantee.

There are some limitations to the guarantee.

- A compiler plugin can check only those parts of your program that you run it on. If you compile some parts of your program without running the checker, then there is no guarantee that the entire program satisfies the property being checked. Some examples of un-checked code are:

- Code compiled without the `-processor` switch, including any external library supplied as a `.class` file.
- Code compiled with the `-AskipUses` or `-AskipDefs` properties (see Section 20.2).
- Suppression of warnings, such as via the `@SuppressWarnings` annotation (see Section 20.2).
- Native methods (because the implementation is not Java code, it cannot be checked).

In each of these cases, any *use* of the code is checked — for example, a call to a native method must be compatible with any annotations on the native method’s signature. However, the annotations on the un-checked code are trusted; there is no verification that the implementation of the native method satisfies the annotations.

- Reflection can violate the Java type system, and the checkers are not sophisticated enough to reason about the possible effects of reflection. Similarly, deserialization and cloning can create objects that could not result from normal constructor calls, and that therefore may violate the property being checked.
- Your code should pass the Java compiler without errors or warnings. In particular, your code should use generic types, with no uses of raw types. Misuse of generics, including casting away generic types, can cause other errors to be missed.
- The Checker Framework does not yet support annotations on intersection types (see JLS §4.9). As a result, checkers cannot provide guarantees about intersection types.
- Specific checkers may have other limitations; see their documentation for details.

A checker can be useful in finding bugs or in verifying part of a program, even if the checker is unable to verify the correctness of an entire program.

In order to avoid a flood of unhelpful warnings, many of the checkers avoid issuing the same warning multiple times. For example, in this code:

```
@Nullable Object x = ...;
x.toString();           // warning
x.toString();           // no warning
```

In this case, the second call to `toString` cannot possibly throw a null pointer warning — `x` is non-null if control flows to the second statement. In other cases, a checker avoids issuing later warnings with the same cause even when later code in a method might also fail. This does not affect the soundness guarantee, but a user may need to examine more warnings after fixing the first ones identified. (More often, at least in our experience to date, a single fix corrects all the warnings.)

If you find that a checker fails to issue a warning that it should, then please report a bug (see Section 25.2).

2.4 Tips about writing annotations

2.4.1 How to get started annotating legacy code

Annotating an entire existing program may seem like a daunting task. But, if you approach it systematically and do a little bit at a time, you will find that it is manageable.

You should start with a property that matters to you, to achieve the best benefits. It is easiest to add annotations if you know the code or the code contains documentation; you will find that you spend most of your time understanding the code, and very little time actually writing annotations or running the checker.

Don’t get discouraged if you see many type-checker warnings at first. Often, adding just a few missing annotations will eliminate many warnings, and you’ll be surprised how fast the process goes overall.

It is best to annotate one package at a time, and to annotate the entire package so that you don’t forget any classes (failing to annotate a class can lead to unexpected results). Start as close to the leaves of the call tree as possible, such as with libraries — that is, start with methods/classes/packages that have few dependences on other code or, equivalently, start with code that a lot of your other code depends on. The reason for this is that it is easiest to annotate a class if the code it calls has already been annotated.

For each class, read its Javadoc. For instance, if you are adding annotations for the Nullness Checker (Section 3), then you can search the documentation for “null” and then add `@Nullable` anywhere appropriate. Do not annotate the method bodies yet — first, get the signatures and fields annotated. The only reason to even *read* the method bodies

yet is to determine signature annotations for undocumented methods — for example, if the method returns `null`, you know its return type should be annotated `@Nullable`, and a parameter that is compared against `null` may need to be annotated `@Nullable`. If you are only annotating signatures (say, for a library you do not maintain and do not wish to check), you are now done.

If you wish to check the implementation, then after the signatures are annotated, run the checker. Then, add method body annotations (usually, few are necessary), fix bugs in code, and add annotations to signatures where necessary. If signature annotations are necessary, then you may want to fix the documentation that did not indicate the property; but this isn't strictly necessary, since the annotations that you wrote provide that documentation.

You may wonder about the effect of adding a given annotation — how many other annotations it will require, or whether it conflicts with other code. Suppose you have added an annotation to a method parameter. You could manually examine all callees. A better way can be to save the checker output before adding the annotation, and to compare it to the checker output after adding the annotation. This helps you to focus on the specific consequences of your change.

Also see Chapter 20, which tells you what to do when you are unable to eliminate checker warnings.

2.4.2 Do not annotate local variables unless necessary

The checker infers annotations for local variables (see Section 19.1.2). Usually, you only need to annotate fields and method signatures. After doing those, you can add annotations inside method bodies if the checker is unable to infer the correct annotation, if you need to suppress a warning (see Section 20.2), etc.

2.4.3 Annotations indicate normal behavior

You should use annotations to indicate *normal* behavior. The annotations indicate all the values that you *want* to flow to reference — not every value that might possibly flow there if your program has a bug.

Many methods are guaranteed to throw an exception if they are passed `null` as an argument. Examples include

```
java.lang.Double.valueOf(String)
java.lang.String.contains(CharSequence)
org.junit.Assert.assertNotNull(Object)
com.google.common.base.Preconditions.checkNotNull(Object)
```

`@Nullable` (see Section 3.2) might seem like a reasonable annotation for the parameter, for two reasons. First, `null` is a legal argument with a well-defined semantics: throw an exception. Second, `@Nullable` describes a possible program execution: it might be possible for `null` to flow there, if your program has a bug.

However, it is never useful for a programmer to pass `null`. It is the programmer's intention that `null` never flows there. If `null` does flow there, the program will not continue normally.

Therefore, you should mark such parameters as `@NonNull`, indicating the intended use of the method. When you use the `@NonNull` annotation, the checker is able to issue compile-time warnings about possible run-time exceptions, which is its purpose. Marking the parameter as `@Nullable` would suppress such warnings, which is undesirable.

2.4.4 Subclasses must respect superclass annotations

An annotation indicates a guarantee that a client can depend upon. A subclass is not permitted to *weaken* the contract; for example, if a method accepts `null` as an argument, then every overriding definition must also accept `null`. A subclass is permitted to *strengthen* the contract; for example, if a method does *not* accept `null` as an argument, then an overriding definition is permitted to accept `null`.

As a bad example, consider an erroneous `@Nullable` annotation at line 141 of `com/google/common/collect/Multiset.java`, version r78:

```
101 public interface Multiset<E> extends Collection<E> {
...

```



```

122  /**
123   * Adds a number of occurrences of an element to this multiset.
...
129   * @param element the element to add occurrences of; may be {@code null} only
130   *       if explicitly allowed by the implementation
...
137   * @throws NullPointerException if {@code element} is null and this
138   *       implementation does not permit null elements. Note that if {@code
139   *       occurrences} is zero, the implementation may opt to return normally.
140   */
141   int add(@Nullable E element, int occurrences);

```

There exist implementations of `Multiset` that permit null elements, and implementations of `Multiset` that do not permit null elements. A client with a variable `Multiset ms` does not know which variety of `Multiset` `ms` refers to. However, the `@Nullable` annotation promises that `ms.add(null, 1)` is permissible. (Recall from Section 2.4.3 that annotations should indicate normal behavior.)

If parameter `element` on line 141 were to be annotated, the correct annotation would be `@NonNull`. Suppose a client has a reference to same `Multiset ms`. The only way the client can be sure not to throw an exception is to pass only non-null elements to `ms.add()`. A particular class that implements `Multiset` could declare `add` to take a `@Nullable` parameter. That still satisfies the original contract. It strengthens the contract by promising even more: a client with such a reference can pass any non-null value to `add()`, and may also pass `null`.

However, the best annotation for line 141 is no annotation at all. The reason is that each implementation of the `Multiset` interface should specify its own nullness properties when it specifies the type parameter for `Multiset`. For example, two clients could be written as

```

class MyNullPermittingMultiset implements Multiset<@Nullable Object> { ... }
class MyNullProhibitingMultiset implements Multiset<@NonNull Object> { ... }

```

or, more generally, as

```

class MyNullPermittingMultiset<E extends @Nullable Object> implements Multiset<E> { ... }
class MyNullProhibitingMultiset<E extends @NonNull Object> implements Multiset<E> { ... }

```

Then, the specification is more informative, and the Checker Framework is able to do more precise checking, than if line 141 has an annotation.

It is a pleasant feature of the Checker Framework that in many cases, no annotations at all are needed on type parameters such as `E` in `MultiSet`.

2.4.5 Annotations on constructor invocations

In the checkers distributed with the Checker Framework, an annotation on a constructor invocation is equivalent to a cast on a constructor result. That is, the following two expressions have identical semantics: one is just shorthand for the other.

```

new @ReadOnly Date()
(@ReadOnly Date) new Date()

```

However, you should rarely have to use this. The Checker Framework will determine the qualifier on the result, based on the “return value” annotation on the constructor definition. The “return value” annotation appears before the constructor name, for example:

```

class MyClass {
    @ReadOnly MyClass() { ... }
}

```

In general, you should only use an annotation on a constructor invocation when you know that the cast is guaranteed to succeed. An example from the IGJ checker (Chapter 5) is `new @Immutable MyClass()` or `new @Mutable MyClass()`, where you know that every other reference to the class is annotated `@ReadOnly`.

2.4.6 When to use (and not use) type qualifiers

For some programming tasks, you can use either a Java subclass or a type qualifier. For instance, suppose that your code currently uses `String` to represent an address. You could create a new `Address` class and refactor your code to use it, or you could create a `@Address` annotation and apply it to some uses of `String` in your code. If both of these are truly possible, then it is probably more foolproof to use the Java class. We do not encourage you to use type qualifiers as a poor substitute for classes. However, sometimes type qualifiers are a better choice.

Using a new class may make your code incompatible with existing libraries or clients. Brian Goetz expands on this issues in an article on the pseudo-typedef antipattern [Goe06]. Even if compatibility is not a concern, a code change may introduce bugs, whereas adding annotations does not change the run-time behavior. It is possible to add annotations to existing code, including code you do not maintain or cannot change. It is possible to annotate primitive types without converting them to wrappers, which would make the code both uglier and slower.

Type qualifiers can be applied to any type, including final classes that cannot be subclassed.

Type qualifiers permit you to remove operations, with a compile-time guarantee. An example is mutating methods that are forbidden by immutable types (see Chapters 5 and 6). More generally, type qualifiers permit creating a new supertype, not just a subtype, of an existing Java type.

A final reason is efficiency. Type qualifiers can be more efficient, since there is no run-time representation such as a wrapper or a separate class, nor introduction of dynamic dispatch for methods that could otherwise be statically dispatched.

2.4.7 What to do if a checker issues a warning about your code

When you first run a type-checker on your code, it is likely to issue warnings or errors. For each warning, try to understand why the checker issues it. (For example, if you are using the Nullness checker (Chapter 3, page 19), try to understand why it cannot prove that no null pointer exception ever occurs.) The reason will sometimes be an actually possible null dereference, sometimes be a weakness of the annotations, and sometimes be a weakness of the checker. You will need to examine your code, and possibly write test cases, to understand the reason.

If there is an actual possible null dereference, then fix your code to prevent that crash.

If there is a weakness in the annotations, then improve the annotations. For example, continuing the Nullness Checker example, if a particular variable is annotated as `@Nullable` but it actually never contains `null` at run time, then change the annotation to `@NonNull`. The weakness might be in the annotations in your code, or in the annotations in a library that your code calls. Another possible problem is that a library is unannotated (see Chapter 21, page 89).

If there is a weakness in the checker, then your code is safe — it never suffers the specific run-time error — but the checker cannot prove this fact. This is most often because the checker is not omniscient, and some tricky coding paradigms are beyond its analysis capabilities; in this case, you should suppress the warning (see Chapter 20.2, page 84). In other cases, the problem is a bug in the checker; in this case, please report the bug (see Chapter 25.2, page 120).

Chapter 3

Nullness checker

If the Nullness checker issues no warnings for a given program, then running that program will never throw a null pointer exception. This guarantee enables a programmer to prevent errors from occurring when a program is run. See Section 3.1 for more details about the guarantee and what is checked.

The most important annotations supported by the Nullness Checker are `@NonNull` and `@Nullable`. `@NonNull` is rarely written, because it is the default. All of the annotations are explained in Section 3.2.

To run the Nullness Checker, supply the `-processor checkers.nullness.NullnessChecker` command-line option to `javac`. For examples, see Section 3.8.

3.1 What the Nullness checker checks

The checker issues a warning in these cases:

1. When an expression of non-`@NonNull` type is dereferenced, because it might cause a null pointer exception. Dereferences occur not only when a field is accessed, but when an array is indexed, an exception is thrown, a lock is taken in a synchronized block, and more. For a complete description of all checks performed by the Nullness checker, see the Javadoc for `NullnessVisitor`.
2. When an expression of `@NonNull` type might become null, because it is a misuse of the type: the null value could flow to a dereference that the checker does not warn about.
As a special case of an `@NonNull` type becoming null, the checker also warns whenever a field of `@NonNull` type is not initialized in a constructor. Also see the discussion of the `-Alint=uninitialized` command-line option below.

This example illustrates the programming errors that the checker detects:

```
@Nullable Object obj; // might be null
@NonNull Object nobj; // never null
...
obj.toString()        // checker warning: dereference might cause null pointer exception
nobj = obj;           // checker warning: nobj may become null
if (nobj == null)     // checker warning: redundant test
```

Parameter passing and return values are checked analogously to assignments.

The Nullness Checker also checks the correctness, and correct use, of rawness annotations for checking initialization (see Section 3.5) and of map key annotations (see Section 3.6).

The checker performs additional checks if certain `-Alint` command-line options are provided. (See Section 20.2.4 for more details about the `-Alint` command-line option.)

1. If you supply the `-Alint=nulltest` command-line option, then the checker warns when a null check is performed against a value that is guaranteed to be non-null, as in `("m" == null)`. Such a check is unnecessary and might indicate a programmer error or misunderstanding. The lint option is disabled by default because sometimes such checks are part of ordinary defensive programming.
2. If you supply the `-Alint=uninitialized` command-line option, then the checker warns if a constructor fails to initialize any field, including `Nullable` types and primitive types. Such a warning is unrelated to whether your code might throw a null pointer exception. However, you might want to enable this warning because it is better code style to supply an explicit initializer, even if there is a default value such as `0` or `false`. This command-line option does not affect the Nullness Checker's tests that fields of `NonNull` type are initialized — such initialization is mandatory, not optional.

3.2 Nullness annotations

The Nullness checker uses three separate type hierarchies: one for nullness, one for rawness (Section 3.5), and one for map keys (Section 3.6) The Nullness checker has four varieties of annotations: nullness type qualifiers, nullness method annotations, rawness type qualifiers, and map key type qualifiers.

3.2.1 Nullness qualifiers

The nullness hierarchy contains these qualifiers:

@Nullable indicates a type that includes the null value. For example, the type `Boolean` is nullable: a variable of type `Boolean` always has one of the values `TRUE`, `FALSE`, or `null`.

@NonNull indicates a type that does not include the null value. The type `boolean` is non-null; a variable of type `boolean` always has one of the values `true` or `false`. The type `@NonNull Boolean` is also non-null: a variable of type `@NonNull Boolean` always has one of the values `TRUE` or `FALSE` — never `null`. Dereferencing an expression of non-null type can never cause a null pointer exception.

The `@NonNull` annotation is rarely written in a program, because it is the default (see Section 3.3.2).

@PolyNull indicates qualifier polymorphism. For a description of `@PolyNull`, see Section 18.2.

@LazyNonNull indicates a reference that may be `null`, but if it ever becomes non-null, then it never becomes `null` again. This is appropriate for lazily-initialized fields, among other uses. When the variable is read, its type is treated as `Nullable`, but when the variable is assigned, its type is treated as `NonNull`.

Because the Nullness checker works intraprocedurally (it analyzes one method at a time), when a `LazyNonNull` field is first read within a method, the field cannot be assumed to be non-null. The benefit of `LazyNonNull` over `Nullable` is its different interaction with flow-sensitive type qualifier refinement (Section 19.1.2). After a check of a `LazyNonNull` field, all subsequent accesses *within that method* can be assumed to be `NonNull`, even after arbitrary external method calls that have access to the given field.

Figure 3.1 shows part of the type hierarchy for the Nullness type system. (The annotations exist only at compile time; at run time, Java has no multiple inheritance.)

3.2.2 Nullness method annotations

The Nullness checker supports several annotations that specify method behavior. These are declaration annotations, not type annotations: they apply to the method itself rather than to some particular type.

@NonNullOnEntry indicates a method precondition: The annotated method expects the specified variables (typically field references) to be non-null when the method is invoked.

@AssertNonNullAfter

@AssertNonNullIfTrue




Figure 3.1: Partial type hierarchy for the Nullness type system. Java’s `Object` is expressed as `@Nullable Object`. Programmers can omit most type qualifiers, because the default annotation (Section 3.3.2) is usually correct. Also shown is the type hierarchy for rawness (Section 3.5), which indicates whether initialization has completed. The two type hierarchies are independent but inter-related. The Nullness Checker verifies both of these, as well as another type hierarchy for map keys (Section 3.2.4).

@AssertNonNullIfFalse indicates a method postcondition. With `@AssertNonNullAfter`, the given expressions are non-null after the method returns; this is useful for a method that initializes a field, for example. With `@AssertNonNullIfTrue` and `@AssertNonNullIfFalse`, if the annotated method returns the given boolean value (true or false), then the given expressions are non-null. See Section 3.3.3 and the Javadoc for examples of their use.

@Pure indicates that the method has no (visible) side effects. Furthermore, if the method is called multiple times with the same arguments, then it returns the same result.

For example, consider the following declaration and uses:

```
@Nullable Object getField(Object arg) { ... }

...
if (x.getField(y) != null) {
    x.getField(y).toString();
}
```

Ordinarily, the Nullness Checker would issue a warning regarding the `toString()` call, because the receiver `x.getField(y)` might be null, according to the `@Nullable` annotation in the declaration of `getField`. If you change the declaration of `getField` to

```
@Pure @Nullable Object getField(Object arg) { ... }
```

then the Nullness Checker issues no warnings, because it can reason that the two invocations `x.getField(y)` have the same value, and therefore that `x.getField(y)` is non-null within the then branch of the if statement.

If a method is pure, then it would be legal to annotate its receiver and every parameter as `@ReadOnly`, in the IGJ (Chapter 5) or Javari (Chapter 6) type systems. The reverse is not true, both because the method might side-effect a global variable and because the method might not be deterministic.

@AssertParametersNonNull is used for suppressing warnings, in very rare cases. See the Javadoc for details.

3.2.3 Rawness qualifiers

The Nullness Checker supports rawness annotations that indicate whether an object is fully initialized — that is, whether its fields have all been assigned.

@Raw
@NonRaw
@PolyRaw

Use of these annotations can help you to type-check more code. Figure 3.1 shows its type hierarchy. For details, see Section 3.5.

3.2.4 Map key qualifiers

The Nullness Checker supports a map key annotation, `@KeyFor` that indicates whether a value is a key for a given map — that is, whether `map.containsKey(value)` would evaluate to `true`.

`@KeyFor`

Use of this annotation can help you to type-check more code. For details, see Section 3.6.

3.3 Writing nullness annotations

3.3.1 Implicit qualifiers

As described in Section 19.1, the Nullness checker adds implicit qualifiers, reducing the number of annotations that must appear in your code. For example, enum types are implicitly non-null, so you never need to write `@NonNull MyEnumType`.

For a complete description of all implicit nullness qualifiers, see the Javadoc for `NullnessAnnotatedTypeFactory`.

3.3.2 Default annotation

Unannotated references are treated as if they had a default annotation, using the NNEL (non-null except locals) rule described below. A user may choose a different rule for defaults using the `@DefaultQualifier` annotation; see Section 19.1.1.

Here are three possible default rules you may wish to use. Other rules are possible but are not as useful.

- `@Nullable`: Unannotated types are regarded as possibly-null, or nullable. This default is backward-compatible with Java, which permits any reference to be null. You can activate this default by writing a `@DefaultQualifier(Nullable.class)` annotation on a class or method declaration.
- `@NonNull`: Unannotated types are treated as non-null. You can activate this default via the `@DefaultQualifier(NonNull.class)` annotation.
- Non-null except locals (NNEL): Unannotated types are treated as `@NonNull`, *except* that the unannotated raw type of a local variable is treated as `@Nullable`. (Any generic arguments to a local variable still default to `@NonNull`.) This is the standard behavior. You can explicitly activate this default via the `@DefaultQualifier(value=NonNull.class, locations={DefaultLocation.ALL_EXCEPT_LOCALS})` annotation.

The NNEL default leads to the smallest number of explicit annotations in your code [PAC⁺08]. It is what we recommend. If you do not explicitly specify a different default, then NNEL is the default.

3.3.3 Conditional nullness

The Nullness Checker supports a form of conditional nullness types, via the `@AssertNonNullIfTrue` and `@AssertNonNullIfFalse` method annotations. The annotation on a method declares that some expressions are non-null, if the method returns true (false, respectively).

Consider `java.io.File`. Method `File.listFiles()` may return null, but is specified to return a non-null value if `File.isDirectory()` is true. The same holds for method `File.list()`. You could declare this relationship in the following way (this particular example is already done for you in the annotated JDK that comes with the Checker Framework):

```
class File {  
  
    @AssertNonNullIfTrue({"list()", "listFiles()"})  
    public boolean isDirectory() { ... }  
}
```

```

    public File @Nullable [] listFiles();
}

```

A client that checks that a `File` reference is indeed that of a directory, can then de-reference `File.isDirectory` safely without any nullness check.

```

static void analyze(File file) {
    if (file.isDirectory()) {
        for (File child : file.listFiles()) { // no possible null dereference
            analyze(child);
        }
    } else {
        ... analyze file ...
    }
}

```

The argument to `@AssertNonNullIfTrue` and `@AssertNonNullIfFalse` is a Java expression, including method calls (as shown above), method formal parameters, fields, etc.; for details, see Section 19.2. More examples of the use of these annotations appear in the Javadoc for `@AssertNonNullIfTrue` and `@AssertNonNullIfFalse`.

3.3.4 Nullness and arrays

The components of a newly created object of reference type are all null. Only after initialization can the array actually be considered to contain non-null components. Therefore, the following is not allowed:

```

@NonNull Object [] oa = new @NonNull Object[10]; // error

```

Instead, one creates a nullable or lazy-nonnull array, initializes each component, and then assigns the result to a non-null array:

```

@LazyNonNull Object [] temp = new @LazyNonNull Object[10];
for (int i = 0; i < temp.length; ++i) {
    temp[i] = new Object();
}
@SuppressWarnings("nullness")
@NonNull Object [] oa = temp;

```

Note that the checker is currently not powerful enough to ensure that each array component was initialized. Therefore, the last assignment needs to be trusted.

3.3.5 Inference of `@NonNull` and `@Nullable` annotations

It can be tedious to write annotations in your code. Tools exist that can automatically infer annotations and insert them in your source code. (This is different than type qualifier refinement for local variables (Section 19.1.2), which infers a more specific type for local variables and uses them during type-checking but does not insert them in your source code. Type qualifier refinement is always enabled, no matter how annotations on signatures got inserted in your source code.)

Your choice of tool depends on what default annotation (see Section 3.3.2) your code uses. You only need one of these tools.

- Inference of `@Nullable`: If your code uses the standard NNEL (non-null-except-locals) default or the `NonNull` default, then use the `AnnotateNullable` tool of the Daikon invariant detector.
- Inference of `@NonNull`: If your code uses the `Nullable` default, use one of these tools:
 - Julia analyzer,
 - Nit: Nullability Inference Tool,
 - Non-null checker and inferencer of the JastAdd Extensible Compiler.

3.4 Suppressing nullness warnings

The Checker Framework supplies several ways to suppress warnings, most notably the `@SuppressWarnings("nullness")` annotation (see Section 20.2). An example use is

```
// might return null
@Nullable Object getObject() { ... }

void myMethod() {
    // The programmer knows that this particular call never returns null.
    @SuppressWarnings("nullness")
    @NonNull Object o2 = getObject();
}
```

The Nullness Checker supports an additional warning suppression key, `nullness:generic.argument`. Use of `@SuppressWarnings("nullness:generic.argument")` causes the Nullness Checker to suppress warnings related to misuse of generic type arguments. One use for this key is when a class is declared to take only `@NonNull` type arguments, but you want to instantiate the class with a `@Nullable` type argument, as in `List<@Nullable Object>`. For a more complete explanation of this example, see Section 24.18, page 116.

The Nullness Checker also permits you to use assertions or method calls to suppress warnings; see below.

3.4.1 Suppressing warnings with assertions and method calls

Occasionally, it is inconvenient or verbose to use the `@SuppressWarnings` annotation. For example, Java does not permit annotations such as `@SuppressWarnings` to appear on statements.

For situations when the `@SuppressWarnings` annotation is inconvenient, the Nullness Checker provides three additional ways to suppress warnings: via an `assert` statement, the `castNonNull` method, and the `@AssertParametersNonNull` annotation. These are appropriate when the Nullness Checker issues a warning, but the programmer knows for sure that the warning is a false positive, because the value cannot ever be null at run time.

1. Use an assertion. If the string “nullness” appears in the message body, then the Nullness Checker treats the assertion as suppressing a warning and assumes that the assertion always succeeds. For example, the checker assumes that no null pointer exception can occur in code such as

```
assert x != null : "@SuppressWarnings(nullness)";
... x.f ...
```

If the string “nullness” does not appear in the assertion message, then the Nullness Checker treats the assertion as being used for defensive programming, and it warns if the method might throw a nullness-related exception. A downside of putting the string in the assertion message is that if the assertion ever fails, then a user might see the string and be confused. But the string should only be used if the programmer has reasoned that the assertion can never fail.

2. Use the `NullnessUtils.castNonNull` method.

The Nullness Checker considers both the return value, and also the argument, to be non-null after the method call. Therefore, the `castNonNull` method can be used either as a cast expression or as a statement. The Nullness Checker issues no warnings in any of the following code:

```
// one way to use as a cast:
@NonNull String s = castNonNull(possiblyNull1);

// another way to use as a cast:
castNonNull(possiblyNull2).toString();

// one way to use as a statement:
castNonNull(possiblyNull3);
possiblyNull3.toString();`
```


The method also throws `AssertionError` if Java assertions are enabled and the argument is `null`. However, it is not intended for general defensive programming; see Section 3.4.2.

A potential disadvantage of using the `castNonNull` method is that your code becomes dependent on the Checker Framework at run time as well as at compile time. You can avoid this by copying the implementation of `castNonNull` into your own code, and possibly renaming it if you do not like the name. Be sure to retain the documentation that indicates that your copy is intended for use only to suppress warnings and not for defensive programming. See Section 3.4.2 for an explanation of the distinction.

3. Use the `@AssertParametersNonNull` annotation. It is used on `castNonNull`, and may be used on other methods with the same semantics; it should probably never be used in any other situation.

3.4.2 Suppressing warnings on nullness-checking routines and defensive programming

One way to suppress warnings in the Nullness Checker is to use method `castNonNull`. (Section 3.4.1 gives other techniques.)

This section explains why the Nullness Checker introduces a new method rather than re-using the `assert` statement (as in `assert x != null`) or an existing method such as:

```
org.junit.Assert.assertNotNull(Object)
com.google.common.base.Preconditions.checkNotNull(Object)
```

In each case, the assertion or method indicates an application invariant — a fact that should always be true. There are two distinct reasons a programmer may have written the invariant, depending on whether the programmer is 100% sure that the application invariant holds.

1. A programmer might write it as **defensive programming**. This causes the program to throw an exception, which is useful for debugging because it gives an earlier run-time indication of the error. A programmer would use an assertion in this way if the programmer is not 100% sure that the application invariant holds.
2. A programmer might write it to **suppress** false positive **warning messages** from a checker. A programmer would use an assertion this way if the programmer is 100% sure that the application invariant holds, and the reference can never be null at run time.

With assertions and existing methods like JUnit's `assertNotNull`, there is no way of knowing the programmer's intent in using the method. Different programmers or codebases may use them in different ways. Guessing wrong would make the Nullness Checker less useful, because it would either miss real errors or issue warnings where there is no real error. Also, different checking tools issue different false warnings that need to be suppressed, so warning suppression needs to be customized for each tool rather than inferred from general-purpose code.

As an example of using assertions for defensive programming, some style guides suggest using assertions or method calls to indicate nullness. A programmer might write

```
String s = ...
assert s != null;    // or: assertNotNull(s);    or: checkNotNull(s);
... Double.valueOf(s) ...
```

A programming error might cause `s` to be null, in which case the code would throw an exception at run time. If the assertion caused the Nullness Checker to assume that `s` is not null, then the Nullness Checker would issue no warning for this code. That would be undesirable, because the whole purpose of the Nullness Checker is to give a compile-time warning about possible run-time exceptions. Furthermore, if the programmer uses assertions for defensive programming systematically throughout the codebase, then many useful Nullness Checker warnings would be suppressed.

Because it is important to distinguish between the two uses of assertions (defensive programming vs. suppressing warnings), the Checker Framework introduces the `NullnessUtils.castNonNull` method. Unlike existing assertions and methods, `castNonNull` is intended only to suppress false warnings that are issued by the Nullness Checker, not for defensive programming.

If you know that a particular codebase uses a nullness-checking method not for defensive programming but to indicate facts that are guaranteed to be true (that is, these assertions will never fail at run time), then you can cause the Nullness Checker to suppress warnings related to them, just as it does for `castNonNull`. Annotate its definition just as `NullnessUtils.castNonNull` is annotated (see the source code for the Checker Framework). Also, be sure to document the intention in the method's Javadoc, so that programmers do not accidentally misuse it for defensive programming.

If you are annotating a codebase that already contains precondition checks, such as:

```
public String get(String key, String def) {
    checkNotNull(key, "key"); //NOI18N
    ...
}
```

then you should mark the appropriate parameter as `@NonNull` (which is the default). This will prevent the checker from issuing a warning about the `checkNotNull` call.

3.5 `@Raw` annotation for partially-initialized objects

An object is *raw* from the time that its constructor starts until its constructor finishes. This is relevant to the Nullness Checker because while the constructor is executing — that is, before initialization completes — a `@NonNull` field may be observed to be null, until that field is set. In particular, the Nullness Checker issues a warning for code like this:

```
public class MyClass {
    private @NonNull Object f;
    public MyClass(int x, int y) {
        // Error because constructor contains no assignment to this.f.
        // By the time the constructor exits, f must be initialized to a non-null value.
    }
    public MyClass(int x) {
        // Error because this.f is accessed before f is initialized.
        // At the beginning of the constructor's execution, accessing this.f
        // yields null, even though field f has a non-null type.
        this.f.toString();
    }
    public MyClass(int x, int y, int z) {
        m();
    }
    public void m() {
        // Error because this.f is accessed before f is initialized,
        // even though the access is not in a constructor.
        // When m is called from the constructor, accessing f yields null,
        // even though field f has a non-null type.
        this.f.toString();
    }
}
```

In general, code can depend that field `f` is not null, because the field is declared with a `@NonNull` type. However, this guarantee does not hold for a partially-initialized object.

The Nullness Checker uses the `@Raw` annotation to indicate that an object is not yet fully initialized — that is, not all its `@NonNull` fields have been assigned. Rawness is mostly relevant within the constructor, or for references to `this` that escape the constructor (say, by being stored in a field or passed to a method before initialization is complete). Use of rawness annotations is rare in most code.

Declarations	Expression	Expression's nullness type, or checker error
<pre>class C { @NonNull Object f; @Nullable Object g; ... }</pre>		
<code>@NonNull @NonRaw C a;</code>	a	@NonNull
	a.f	@NonNull
	a.g	@Nullable
<code>@NonNull @Raw C b;</code>	b	@NonNull
	b.f	@LazyNonNull
	b.g	@Nullable
<code>@Nullable @NonRaw C c;</code>	c	@Nullable
	c.f	error: deref of nullable
	c.g	error: deref of nullable
<code>@Nullable @Raw C d;</code>	d	@Nullable
	d.f	error: deref of nullable
	d.g	error: deref of nullable

Figure 3.2: Examples of the interaction between nullness and rawness. Declarations are shown at the left for reference, but the focus of the table is the expressions and their nullness type or error.

3.5.1 Rawness qualifiers

The rawness hierarchy is shown in Figure 3.1. The rawness hierarchy contains these qualifiers:

@Raw indicates a type that may contain a partially-initialized object. In a partially-initialized object, fields that are annotated as @NonNull may be null because the field has not yet been assigned. Within the constructor, `this` has @Raw type until all the @NonNull fields have been assigned. A partially-initialized object (`this` in a constructor) may be passed to a helper method or stored in a variable; if so, the method receiver, or the field, would have to be annotated as @Raw.

@NonRaw indicates a type that contains a fully-initialized object. `NonRaw` is the default, so there is little need for a programmer to write this explicitly.

@PolyRaw indicates qualifier polymorphism over rawness (see Section 18.2).

If a reference has @Raw type, then all of its @NonNull fields are treated as @LazyNonNull: when read, they are treated as being @Nullable, but when written, they are treated as being @NonNull.

The rawness hierarchy is orthogonal to the nullness hierarchy. It is legal for a reference to be @NonNull @Raw, @Nullable @Raw, @NonNull @NonRaw, or @Nullable @NonRaw. The nullness hierarchy tells you about the reference itself: might the reference be null? The rawness hierarchy tells you about the @NonNull fields in the referred-to object: might those fields be temporarily null in contravention of their type annotation? Figure 3.2 contains some examples.

3.5.2 How an object becomes non-raw

Within the constructor, `this` starts out with @Raw type. As soon as all of the @NonNull fields have been initialized, then `this` is treated as non-raw.

The Nullness checker issues an error if the constructor fails to initialize any @NonNull field. This ensures that the object is in a legal (non-raw) state by the time that the constructor exits. This is different than Java's test for definite assignment (see JLS ch.16), which does not apply to fields (except blank final ones, defined in JLS §4.12.4) because fields have a default value of null.

All @NonNull fields must either have a default in the field declaration, or be assigned in the constructor or in a helper method that the constructor calls. If your code initializes (some) fields in a helper method, you will need to annotate the helper method with an annotation such as `@AssertNonNullAfter({"field1", "field2"})` for all the fields

that the helper method assigns. It's a bit odd, but you use that same annotation, `@AssertNonNullAfter`, to indicate that a primitive field has its value set in a helper method, which is relevant when you supply the `-Alint=uninitialized` command-line option (see Section 2).

3.5.3 More details about rawness checking

Suppressing warnings You can suppress warnings related to partially-initialized objects with `@SuppressWarnings("rawness")`. Do not confuse this with the unrelated `@SuppressWarnings("rawtypes")` annotation for non-instantiated generic types!

Checking initialization of all fields, not just `@NonNull` ones When the `-Alint=uninitialized` command-line option is provided, then an object is considered raw until *all* its fields are assigned, not just the `@NonNull` ones. See Section 2.

The terminology “raw” The name “raw” comes from a research paper that proposed this approach [FL03]. A better name might have been “not yet initialized” or “partially initialized”, but the term “raw” is now well-known. The `@Raw` annotation has nothing to do with the raw types of Java Generics.

3.6 Map key annotations

Java's `Map.get` method always has the possibility to return null, if the key is not in the map. Thus, to guarantee that the value returned from `Map.get` is non-null, it is necessary that the map contains only non-null values, *and* the key is in the map. The `@KeyFor` annotation states the latter property.

If a type is annotated as `@KeyFor("m")`, then any value `v` with that type is a key in `Map m`. Another way of saying this is that the expression `m.containsKey(v)` evaluates to true.

You usually do not have to write `@KeyFor` explicitly, because the checker infers it based on usage patterns, such as calls to `containsKey` or iteration over a map's key set.

One usage pattern where you *do* have to write `@KeyFor` is for a user-managed collection that is a subset of the key set:

```
Map<String, Object> m;
Set<@KeyFor("m") String> matchingKeys; // keys that match some criterion
for (@KeyFor("m") String k : matchingKeys) {
    ... m.get(k) ... // known to be non-null
}
```

As with any annotation, use of the `@KeyFor` annotation may force you to slightly refactor your code. For example, this would be illegal:

```
Map<K, V> m;
Collection<@KeyFor("m") K> coll;
coll.add(x); // compiler error, because the @KeyFor annotation is violated
m.put(x, ...);
```

but this would be OK (no compiler error):

```
Collection<@KeyFor("m") K> coll;
m.put(x, ...);
coll.add(x);
```

Because the `@KeyFor` type hierarchy is independent from the nullness and rawness hierarchies, it uses a different warning suppression key. You can suppress warnings related to map keys with `@SuppressWarnings("keyfor")`.

3.7 Additional details

The Nullness Checker does some special checks in certain circumstances, in order to soundly reduce the number of warnings that it produces.

For example, a call to `System.getProperty(String)` can return null in general, but it will not return null if the argument is one of the built-in-keys listed in the documentation of `System.getProperties()`. The Nullness Checker is aware of this fact, so you do not have to suppress a warning for a call like `System.getProperty("line.separator")`. The warning is still issued for code like this:

```
final String s = "line.separator";
nonnullvar = System.getProperty(s);
```

though that case could be handled as well, if desired. (Suppression of the warning is, strictly speaking, not sound, because a library that your code calls, or your code itself, could perversely change the system properties; the Nullness Checker assumes this bizarre coding pattern does not happen.)

3.8 Examples

3.8.1 Tiny examples

To try the Nullness checker on a source file that uses the `@NonNull` qualifier, use the following command (where `javac` is the JSR 308 compiler that is distributed with the Checker Framework):

```
javac -processor checkers.nullness.NullnessChecker examples/NullnessExample.java
```

Compilation will complete without warnings.

To see the checker warn about incorrect usage of annotations (and therefore the possibility of a null pointer exception at run time), use the following command:

```
javac -processor checkers.nullness.NullnessChecker examples/NullnessExampleWithWarnings.java
```

The compiler will issue three warnings regarding violation of the semantics of `@NonNull`.

3.8.2 Annotated library

Some libraries that are annotated with nullness qualifiers are:

- The Nullness checker itself.
- The Plume-lib library. Run the command `make check-nullness`.
- The Daikon invariant detector. Run the command `make check-nullness`.

3.9 Tips for getting started

Here are some tips about getting started using the Nullness Checker on a legacy codebase. For more generic advice (not specific to the Nullness Checker), see Section 2.4.1.

Your goal is to add `@Nullable` annotations to the types of any variables that can be null. (The default is to assume that a variable is non-null unless it has a `@Nullable` annotation.) Then, you will run the Nullness Checker. Each of its errors indicates either a possible null pointer exception, or a wrong/missing annotation. When there are no more warnings from the checker, you are done!

We recommend that you start by searching the code for occurrences of `null` in the following locations; when you find one, write the corresponding annotation:

- in Javadoc: add `@Nullable` annotations to method signatures (parameters and return types).

- `return null`: add a `@Nullable` annotation to the return type of the given method.
- `param == null`: when a formal parameter is compared to `null`, then in most cases you can add a `@Nullable` annotation to the formal parameter's type
- `TypeName field = null;`: when a field is initialized to `null` in its declaration, then it needs either a `@Nullable` or a `@LazyNonNull` annotation. If the field is always set to a non-null value in the constructor, then you can just change the declaration to `TypeName field;`, without an initializer, and write no type annotation (because the default is `@NonNull`).
- declarations of `contains`, `containsKey`, `containsValue`, `equals`, `get`, `indexOf`, `lastIndexOf`, and `remove` (with `Object` as the argument type): change the argument type to `@Nullable Object`; for `remove`, also change the return type to `@Nullable Object`.

You should ignore all other occurrences of `null` within a method body. In particular, you (almost) never need to annotate local variables.

Only after this step should you run `ant` to invoke the Nullness Checker. The reason is that it is quicker to search for places to change than to repeatedly run the checker and fix the errors it tells you about, one at a time.

Here are some other tips:

- In any file where you write an annotation such as `@Nullable`, don't forget to add `import checkers.nullness.quals.*;`
- To indicate an array that can be null, write, for example: `int @Nullable []`.
By contrast, `@Nullable Object []` means a non-null array that contains possibly-null objects.
- If you know that a particular variable is definitely not null, but the Nullness Checker cannot figure it out, then you can tell it by writing an assertion (see Section 20.2):

```
assert var != null : "@SuppressWarnings(nullness)";
```
- To indicate that a routine returns the same value every time it is called, use `@Pure` (see Section 3.2.2).
- To indicate a method precondition (a contract stating the conditions under which a client is allowed to call it), you can use annotations such as `@NonNullOnEntry` (see Section 3.2.2).

3.10 Other tools for nullness checking

The Checker Framework's nullness annotations are similar to annotations used in IntelliJ IDEA, FindBugs, JML, the JSR 305 proposal, NetBeans, and other tools. Also see Section 25.5 for a comparison to other tools.

You might prefer to use the Checker Framework because it has a more powerful analysis that can warn you about more null pointer errors in your code.

If your code is already annotated with a different nullness annotation, you can reuse that effort by converting them to the Checker Framework's nullness annotations. Perform the refactoring described in Figure 3.3.

Alternately, the Checker Framework can process those other annotations (as well as its own, if they also appear in your program). The Checker Framework has its own definition of the annotations on the left side of Figure 3.3, so that they can be used as type qualifiers. The Checker Framework interprets them according to the right side of Figure 3.3.

The Checker Framework may issue more or fewer errors than another tool. This is expected, since each tool uses a different analysis. Remember that the Checker Framework aims at soundness: it aims to never miss a possible null dereference, while at the same time limiting false reports. Also, note FindBugs's non-standard meaning for `@Nullable` (Section 3.10.2).

Because some of the names are the same (`NonNull`, `Nullable`), you can import at most one of the annotations with conflicting names; the other(s) must be written out fully rather than imported.

3.10.1 Which tool is right for you?

Different tools are appropriate in different circumstances. Here is a brief comparison with FindBugs, but similar points apply to other tools.

The Checker Framework has a more powerful nullness analysis; FindBugs misses some real errors. However, FindBugs does not require you to annotate your code as thoroughly as the Checker Framework does. Depending on

com.sun.istack.NotNull
edu.umd.cs.findbugs.annotations.NotNull
javax.annotation.Nonnull
org.jetbrains.annotations.NotNull
org.netbeans.api.annotations.common.NotNull
org.jmlspecs.annotation.NotNull

⇒ checkers.nullness.quals.NotNull

com.sun.istack.Nullable
edu.umd.cs.findbugs.annotations.Nullable
edu.umd.cs.findbugs.annotations.CheckForNull
edu.umd.cs.findbugs.annotations.UnknownNullness
javax.annotation.Nullable
javax.annotation.CheckForNull
javax.validation.constraints.NotNull
org.jetbrains.annotations.Nullable
org.netbeans.api.annotations.common.CheckForNull
org.netbeans.api.annotations.common.NullAllowed
org.netbeans.api.annotations.common.NullUnknown
org.jmlspecs.annotation.Nullable

⇒ checkers.nullness.quals.Nullable

Figure 3.3: Refactoring for converting other nullness annotations to the Checker Framework’s annotations.

the importance of your code, you may desire: no nullness checking, the cursory checking of FindBugs, or the thorough checking of the Checker Framework. You might even want to ensure that both tools run, for example if your coworkers or some other organization are still using FindBugs. If you know that you will eventually want to use the Checker Framework, there is no point using FindBugs first; it is easier to go straight to using the Checker Framework.

FindBugs can find other errors in addition to nullness errors; here we focus on its nullness checks. Even if you use FindBugs for its other features, you may want to use the Checker Framework for analyses that can be expressed as pluggable type-checking, such as detecting nullness errors.

Regardless of whether you wish to use the FindBugs nullness analysis, you may continue running all of the other FindBugs analyses at the same time as the Checker Framework; there are no interactions among them.

If FindBugs (or any other tool) discovers a nullness error that the Checker Framework does not, please report it to us (see Section 25.2) so that we can enhance the Checker Framework.

3.10.2 Incompatibility note about FindBugs @Nullable

FindBugs has a non-standard definition of @Nullable. FindBugs’s treatment is not documented in its own Javadoc; it is different from the definition of @Nullable in every other tool for nullness analysis; it means the same thing as @NonNull when applied to a formal parameter; and it invariably surprises programmers. Thus, FindBugs’s @Nullable is detrimental rather than useful as documentation. In practice, your best bet is to not rely on FindBugs for nullness analysis, even if you find FindBugs useful for other purposes.

You can skip the rest of this section unless you wish to learn more details.

FindBugs suppresses all warnings at uses of a @Nullable variable. (You have to use @CheckForNull to indicate a nullable variable that FindBugs should check.) For example:

```
// declare getObject() to possibly return null
@Nullable Object getObject() { ... }

void myMethod() {
    @Nullable Object o = getObject();
    // FindBugs issues no warning about calling toString on a possibly-null reference!
```

```
    o.toString();  
}
```

The Checker Framework does not emulate this non-standard behavior of FindBugs, even if the code uses FindBugs annotations.

With FindBugs, you annotate a declaration, which suppresses checking at *all* client uses, even the places that you want to check. It is better to suppress warnings at only the specific client uses where the value is known to be non-null; the Checker Framework supports this, if you write `@SuppressWarnings` at the client uses. The Checker Framework also supports suppressing checking at all client uses, by writing a `@SuppressWarnings` annotation at the declaration site. Thus, the Checker Framework supports both use cases, whereas FindBugs supports only one and gives the programmer less flexibility.

In general, the Checker Framework will issue more warnings than FindBugs, and some of them may be about real bugs in your program. See Section 3.4 for information about suppressing nullness warnings.

(FindBugs made a poor choice of names. The choice of names should make a clear distinction between annotations that specify whether a reference is null, and annotations that suppress false warnings. The choice of names should also have been consistent for other tools, and intuitively clear to programmers. The FindBugs choices make the FindBugs annotations less helpful to people, and much less useful for other tools. As a separate issue, the FindBugs analysis is also very imprecise. For type-related analyses, it is best to stay away from the FindBugs nullness annotations, and use a more capable tool like the Checker Framework.)

Chapter 4

Interning checker

If the Interning checker issues no warnings for a given program, then all reference equality tests (i.e., all uses of “==”) are proper; that is, == is not misused where equals() should have been used instead.

Interning is a design pattern in which the same object is used whenever two different objects would be considered equal. Interning is also known as canonicalization or hash-consing, and it is related to the flyweight design pattern. Interning has two benefits: it can save memory, and it can speed up testing for equality by permitting use of ==.

The Interning checker prevents two types of errors in your code. First, == should be used only on interned values; using == on non-interned values can result in subtle bugs. For example:

```
Integer x = new Integer(22);
Integer y = new Integer(22);
System.out.println(x == y); // prints false!
```

The Interning checker helps programmers to prevent such bugs. Second, the Interning checker also helps to prevent performance problems that result from failure to use interning. (See Section 2.3 for caveats to the checker’s guarantees.)

Interning is such an important design pattern that Java builds it in for strings. Every string literal in the program is guaranteed to be interned (JLS §3.10.5), and the String.intern() method performs interning for strings that are computed at run time. Users can also write their own interning methods for other types.

It is a proper optimization to use ==, rather than equals(), whenever the comparison is guaranteed to produce the same result — that is, whenever the comparison is never provided with two different objects for which equals() would return true. Here are three reasons that this property could hold:

1. Interning. A factory method ensures that, globally, no two different interned objects are equals() to one another. (In some cases other, non-interned objects of the class might be equals() to one another; in other cases, every object of the class is interned.) Interned objects should always be immutable.
2. Global control flow. The program’s control flow is such that the constructor for class C is called a limited number of times, and with specific values that ensure the results are not equals() to one another. Objects of class C can always be compared with ==. Such objects may be mutable or immutable.
3. Local control flow. Even though not all objects of the given type may be compared with ==, the specific objects that can reach a given comparison may be. For example, suppose that an array contains no duplicates. Then testing to find the index of a given element that is known to be in the array can use ==.

To eliminate Interning Checker warnings, you will need to annotate your code regarding all legal uses of ==. Thus, the Interning Checker could also have been called the Reference Equality Checker. In the future, the checker will include annotations that target the non-interning cases above, but for now you need to use @Interned, @UsesObjectEquals (which handles a surprising number of cases), and/or @SuppressWarnings.

To run the Interning Checker, supply the -processor checkers.interning.InterningChecker command-line option to javac. For examples, see Section 4.4.




Figure 4.1: Type hierarchy for the Interning type system.

4.1 Interning annotations

These qualifiers are part of the Interning type system:

@Interned indicates a type that includes only interned values (no non-interned values).

@PolyInterned indicates qualifier polymorphism. For a description of `@PolyInterned`, see Section 18.2.

@UsesObjectEquals is a class (not type) annotation that indicates that this class's `equals` method is the same as that of `Object`. In other words, neither this class nor any of its superclasses overrides the `equals` method. Since `Object.equals` uses reference equality, this means that for such a class, `==` and `equals` are equivalent, and so the Interning Checker does not issue warnings for either one.

4.2 Annotating your code with @Interned

In order to perform checking, you must annotate your code with the `@Interned` type annotation, which indicates a type for the canonical representation of an object:

```
String s1 = ...; // type is (uninterned) "String"
@Interned String s2 = ...; // Java type is "String", but checker treats it as "Interned String"
```

The type system enforced by the checker plugin ensures that only interned values can be assigned to `s2`.

To specify that *all* objects of a given type are interned, annotate the class declaration:

```
public @Interned class MyInternedClass { ... }
```

This is equivalent to annotating every use of `MyInternedClass`, in a declaration or elsewhere. For example, enum classes are implicitly so annotated.

4.2.1 Implicit qualifiers

As described in Section 19.1, the Interning checker adds implicit qualifiers, reducing the number of annotations that must appear in your code. For example, `String` literals and the null literal are always considered interned, and object creation expressions (using `new`) are never considered `@Interned` unless they are annotated as such, as in

```
@Interned Double internedDoubleZero = new @Interned Double(0); // canonical representation for Double zero
```

For a complete description of all implicit interning qualifiers, see the Javadoc for `InterningAnnotatedTypeFactory`.

4.3 What the Interning checker checks

Objects of an `@Interned` type may be safely compared using the `==` operator.

The checker issues a warning in two cases:

1. When a reference (in)equality operator (`==` or `!=`) has an operand of non-`@Interned` type.
2. When a non-`@Interned` type is used where an `@Interned` type is expected.

`com.sun.istack.Interned` ⇒ `checkers.interning.qual.Interned`

Figure 4.2: Refactoring for converting interning annotations from other tools to the Checker Framework.

This example shows both sorts of problems:

```
        Object obj;
@Interned Object iobj;
...
if (obj == iobj) { ... } // checker warning: reference equality test is unsafe
iobj = obj;             // checker warning: iobj's referent may no longer be interned
```

The checker also issues a warning when `.equals` is used where `==` could be safely used. You can disable this behavior via the `javac -Alint` command-line option, like so: `-Alint=-dotequals`.

For a complete description of all checks performed by the checker, see the Javadoc for `InterningVisitor`.

You can also restrict which types the checker should examine and type-check, using the `-Acheckclass` option. For example, to find only the interning errors related to uses of `String`, you can pass `-Acheckclass=java.lang.String`. The Interning checker always checks all subclasses and superclasses of the given class.

4.4 Examples

To try the Interning checker on a source file that uses the `@Interned` qualifier, use the following command (where `javac` is the JSR 308 compiler that is distributed with the Checker Framework):

```
javac -processor checkers.interning.InterningChecker examples/InterningExample.java
```

Compilation will complete without warnings.

To see the checker warn about incorrect usage of annotations, use the following command:

```
javac -processor checkers.interning.InterningChecker examples/InterningExampleWithWarnings.java
```

The compiler will issue a warning regarding violation of the semantics of `@Interned`.

The Daikon invariant detector (<http://groups.csail.mit.edu/pag/daikon/>) is also annotated with `@Interned`. From directory `java`, run `make check-interning`.

4.5 Other interning annotations

The Checker Framework's interning annotations are similar to annotations used elsewhere.

If your code is already annotated with a different interning annotation, you can reuse that effort by converting them to the Checker Framework's nullness annotations. Perform the refactoring described in Figure 4.2.

Alternately, the Checker Framework can process those other annotations (as well as its own, if they also appear in your program). The Checker Framework has its own definition of the annotations on the left side of Figure 4.2, so that they can be used as type qualifiers. The Checker Framework interprets them according to the right side of Figure 4.2.

Chapter 5

IGJ immutability checker

IGJ is a Java language extension that helps programmers to avoid mutation errors (unintended side effects). If the IGJ checker issues no warnings for a given program, then that program will never change objects that should not be changed. This guarantee enables a programmer to detect and prevent mutation-related errors. (See Section 2.3 for caveats to the guarantee.)

To run the IGJ Checker, supply the `-processor checkers.igj.IGJChecker` command-line option to `javac`. For examples, see Section 5.7.

5.1 IGJ and Mutability

IGJ [ZPA⁺07] permits a programmer to express that a particular object should never be modified via any reference (object immutability), or that a reference should never be used to modify its referent (reference immutability). Once a programmer has expressed these facts, an automatic checker analyzes the code to either locate mutability bugs or to guarantee that the code contains no such bugs.

To learn more details of the IGJ language and type system, please see the ESEC/FSE 2007 paper “Object and reference immutability using Java generics” [ZPA⁺07]. The IGJ checker supports Annotation IGJ (Section 5.5), which is a slightly different dialect of IGJ than that described in the ESEC/FSE paper.

5.2 IGJ Annotations

Each object is either immutable (it can never be modified) or mutable (it can be modified). The following qualifiers are part of the IGJ type system.

@Immutable An immutable reference always refers to an immutable object. Neither the reference, nor any aliasing reference, may modify the object.

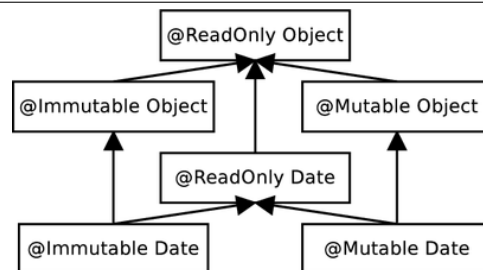


Figure 5.1: Type hierarchy for three of IGJ’s type qualifiers.

@Mutable A mutable reference refers to a mutable object. The reference, or some aliasing mutable reference, may modify the object.

@ReadOnly A readonly reference cannot be used to modify its referent. The referent may be an immutable or a mutable object. In other words, it is possible for the referent to change via an aliasing mutable reference, even though the referent cannot be changed via the readonly reference.

@Assignable The annotated field may be re-assigned regardless of the immutability of the enclosing class or object instance.

@AssignsFields is similar to **@Mutable**, but permits only limited mutation — assignment of fields — and is intended for use by constructor helper methods.

@I simulates mutability overloading or the template behavior of generics. It can be applied to classes, methods, and parameters. See Section 5.5.3.

For additional details, see [ZPA⁺07].

5.3 What the IGJ checker checks

The IGJ checker issues an error whenever mutation happens through a readonly reference, when fields of a readonly reference which are not explicitly marked with **@Assignable** are reassigned, or when a readonly reference is assigned to a mutable variable. The checker also emits a warning when casts increase the mutability access of a reference.

5.4 Implicit and default qualifiers

As described in Section 19.1, the IGJ checker adds implicit qualifiers, reducing the number of annotations that must appear in your code.

For a complete description of all implicit IGJ qualifiers, see the Javadoc for `IGJAnnotatedTypeFactory`.

The default annotation (for types that are unannotated and not given an implicit qualifier) is as follows:

- **@Mutable** for almost all references. This is backward-compatible with Java, since Java permits any reference to be mutated.
- **@ReadOnly** for local variables. This qualifier may be refined by flow-sensitive local type refinement (see Section 19.1.2).
- **@ReadOnly** for type parameter and wildcard bounds. For example,

```
interface List<T extends Object> { ... }
```

is defaulted to

```
interface List<T extends @ReadOnly Object> { ... }
```

This default is not backward-compatible — that is, you may have to explicitly add **@Mutable** annotations to some type parameter bounds in order to make unannotated Java code type-check under IGJ. However, this reduces the number of annotations you must write overall (since most variables of generic type are in fact not modified), and permits more client code to type-check (otherwise a client could not write `List<@ReadOnly Date>`).

5.5 Annotation IGJ Dialect

The IGJ checker supports the Annotation IGJ dialect of IGJ. The syntax of Annotation IGJ is based on type annotations.

The syntax of the original IGJ dialect [ZPA⁺07] was based on Java 5's generics and annotation mechanisms. The original IGJ dialect was not backward-compatible with Java (either syntactically or semantically). The dialect of IGJ checked by the IGJ checker corrects these problems.

The differences between the Annotation IGJ dialect and the original IGJ dialect are as follows.

5.5.1 Semantic Changes

- Annotation IGJ does not permit covariant changes in generic type arguments, for backward compatibility with Java. In ordinary Java, types with different generic type arguments, such as `Vector<Integer>` and `Vector<Number>`, have no subtype relationship, even if the arguments (`Integer` and `Number`) do. The original IGJ dialect changed the Java subtyping rules to permit safely varying a type argument covariantly in certain circumstances. For example,

```
Vector<Mutable, Integer> <: Vector<ReadOnly, Integer>
                          <: Vector<ReadOnly, Number>
                          <: Vector<ReadOnly, Object>
```

is valid in IGJ, but in Annotation IGJ, only

```
@Mutable Vector<Integer> <: @ReadOnly Vector<Integer>
```

holds and the other two subtype relations do not hold

```
@ReadOnly Vector<Integer> </: @ReadOnly Vector<Number>
                          </: @ReadOnly Vector<Object>
```

- Annotation IGJ supports array immutability. The original IGJ dialect did not permit the (im)mutability of array elements to be specified, because the generics syntax used by the original IGJ dialect cannot be applied to array elements.

5.5.2 Syntax Changes

- Immutability is specified through type annotations [Ern08] (Section 5.2), not through a combination of generics and annotations. Use of type annotations makes Annotation IGJ backward compatible with Java syntax.
- Templating over Immutability: The annotation `@I(id)` is used to template over immutability. See Section 5.5.3.

5.5.3 Templating Over Immutability: @I

`@I` is a template annotation over IGJ Immutability annotations. It acts similarly to type variables in Java's generic types, and the name `@I` mimics the standard `<I>` type variable name used in code written in the original IGJ dialect. The annotation value string is used to distinguish between multiple instances of `@I` — in the generics-based original dialect, these would be expressed as two type variables `<I>` and `<J>`.

Usage on classes A class declaration annotated with `@I` can then be used with any IGJ Immutability annotation. The actual immutability that `@I` is resolved to dictates the immutability type for all the non-static appearances of `@I` with the same value as the class declaration.

Example:

```
@I
public class FileDescriptor {
    private @Immutable Date creationData;
    private @I Date lastModData;

    public @I Date getLastModDate(@ReadOnly FileDescriptor this) { }
}

...
void useFileDescriptor() {
    @Mutable FileDescriptor file =
        new @Mutable FileDescriptor(...);
    ...
}
```

```

    @Mutable Data date = file.getLastModDate();
}

```

In the last example, `@I` was resolved to `@Mutable` for the instance file.

Usage on methods For example, it could be used for method parameters, return values, and the actual IGJ immutability value would be resolved based on the method invocation.

For example, the below method `getMidpoint` returns a `Point` with the same immutability type as the passed parameters if `p1` and `p2` match in immutability, otherwise `@I` is resolved to `@ReadOnly`:

```

static @I Point getMidpoint(@I Point p1, @I Point p2) { ... }

```

The `@I` annotation value distinguishes between `@I` declarations. So, the below method `findUnion` returns a collection of the same immutability type as the *first* collection parameter:

```

static <E> @I("First") Collection<E> findUnion(@I("First") Collection<E> coll,
                                              @I("Second") Collection<E> col2) { ... }

```

5.6 Iterators and their abstract state

This section explains why the receiver of `Iterator.next()` is annotated as `@ReadOnly`.

An iterator conceptually has two pieces of state:

1. the underlying collection
2. an index into that collection (indicating the next object to be returned)

We choose to exclude the index from the abstract state of the iterator. That is, a change to the index does not count as a mutation of the iterator itself.

Changes to the underlying collection are more important and interesting, and unintentional changes are much more likely to lead to important errors. Therefore, this choice about the iterator's abstract state appears to be more useful than other choices. For example, if the iterator's abstract state included both the underlying collection and the index, then there would be no way to express, or check, that `Iterator.next` does not change the underlying collection.

5.7 Examples

To try the IGJ checker on a source file that uses the IGJ qualifier, use the following command (where `javac` is the JSR 308 compiler that is distributed with the Checker Framework).

```

javac -processor checkers.igj.IGJChecker examples/IGJExample.java

```

The IGJ checker itself is also annotated with IGJ annotations.

Chapter 6

Javari immutability checker

Javari [TE05, QTE08] is a Java language extension that helps programmers to avoid mutation errors that result from unintended side effects. If the Javari checker issues no warnings for a given program, then that program will never change objects that should not be changed. This guarantee enables a programmer to detect and prevent mutation-related errors. (See Section 2.3 for caveats to the guarantee.) The Javari webpage (<http://types.cs.washington.edu/javari/>) contains papers that explain the Javari language and type system. By contrast to those papers, the Javari checker uses an annotation-based dialect of the Javari language.

The Javarifier tool infers Javari types for an existing program; see Section 6.2.2.

Also consider the IGJ checker (Chapter 5). The IGJ type system is more expressive than that of Javari, and the IGJ checker is a bit more robust. However, IGJ lacks a type inference tool such as Javarifier.

To run the Javari Checker, supply the `-processor checkers.javari.JavariChecker` command-line option to `javac`. For examples, see Section 6.5.

6.1 Javari annotations

The following six annotations make up the Javari type system.

@ReadOnly indicates a type that provides only read-only access. A reference of this type may not be used to modify its referent, but aliasing references to that object might change it.

@Mutable indicates a mutable type.

@Assignable is a field annotation, not a type qualifier. It indicates that the given field may always be assigned, no matter what the type of the reference used to access the field.

@QReadOnly corresponds to Javari’s “? readonly” for wildcard types. An example of its use is `List<@QReadOnly Date>`. It allows only the operations which are allowed for both readonly and mutable types.

@PolyRead (previously named `@RoMaybe`) specifies polymorphism over mutability; it simulates mutability overloading. It can be applied to methods and parameters. See Section 18.2 and the `@PolyRead` Javadoc for more details.




Figure 6.1: Type hierarchy for Javari’s ReadOnly type qualifier.

`@ThisMutable` means that the mutability of the field is the same as that of the reference that contains it. `@ThisMutable` is the default on fields, and does not make sense to write elsewhere. Therefore, `@ThisMutable` should never appear in a program.

6.2 Writing Javari annotations

6.2.1 Implicit qualifiers

As described in Section 19.1, the Javari checker adds implicit qualifiers, reducing the number of annotations that must appear in your code.

For a complete description of all implicit Javari qualifiers, see the Javadoc for `JavariAnnotatedTypeFactory`.

6.2.2 Inference of Javari annotations

It can be tedious to write annotations in your code. The Javarifier tool (<http://types.cs.washington.edu/javari/javarifier/>) infers Javari types for an existing program. It automatically inserts Javari annotations in your Java program or in `.class` files.

This has two benefits: it relieves the programmer of the tedium of writing annotations (though the programmer can always refine the inferred annotations), and it annotates libraries, permitting checking of programs that use those libraries.

6.3 What the Javari checker checks

The checker issues an error whenever mutation happens through a readonly reference, when fields of a readonly reference which are not explicitly marked with `@Assignable` are reassigned, or when a readonly expression is assigned to a mutable variable. The checker also emits a warning when casts increase the mutability access of a reference.

6.4 Iterators and their abstract state

For an explanation of why the receiver of `Iterator.next()` is annotated as `@ReadOnly`, see Section 5.6.

6.5 Examples

To try the Javari checker on a source file that uses the Javari qualifier, use the following command (where `javac` is the JSR 308 compiler that is distributed with the Checker Framework). Alternately, you may specify just one of the test files.

```
javac -processor checkers.javari.JavariChecker tests/javari/*.java
```

The compiler should issue the errors and warnings (if any) specified in the `.out` files with same name.

To run the test suite for the Javari checker, use `ant javari-tests`.

The Javari checker itself is also annotated with Javari annotations.

Chapter 7

Lock checker

The Lock checker prevents certain kinds of concurrency errors. If the Lock checker issues no warnings for a given program, then the program holds the appropriate lock every time that it accesses a variable.

Note: This does *not* mean that your program has *no* concurrency errors. (You might have forgotten to annotate that a particular variable should only be accessed when a lock is held. You might release and re-acquire the lock, when correctness requires you to hold it throughout a computation. And, there are other concurrency errors that cannot, or should not, be solved with locks.) However, ensuring that your program obeys its locking discipline is an easy and effective way to eliminate a common and important class of errors.

To run the Lock Checker, supply the `-processor checkers.lock.LockChecker` command-line option to `javac`.

7.1 Lock annotations

The Lock checker uses two annotations. One is a type qualifier, and the other is a method annotation.

@GuardedBy indicates a type whose value may be accessed only when the given lock is held. See the `GuardedBy` Javadoc for an explanation of the argument. The lock acquisition and the value access may be arbitrarily far in the future; or, if the value is never accessed, the lock never need be held.

@Holding is a method annotation (not a type qualifier). It indicates that when the method is called, the given lock must be held by the caller. In other words, the given lock is already held at the time the method is called.

7.1.1 Examples

Most often, field values are annotated with `@GuardedBy`, but other uses are possible.

A return value may be annotated with `@GuardedBy`:

```
@GuardedBy("MyClass.myLock") Object myMethod() { ... }

// reassignments without holding the lock are OK.
@GuardedBy("MyClass.myLock") Object x = myMethod();
@GuardedBy("MyClass.myLock") Object y = x;
Object z = x; // ILLEGAL (assuming no lock inference),
              // because z can be freely accessed.
x.toString() // ILLEGAL because the lock is not held
synchronized(MyClass.myLock) {
    y.toString(); // OK: the lock is held
}
```

A parameter may be annotated with `@GuardedBy`:

```

void helper1(@GuardedBy("MyClass.myLock") Object a) {
    a.toString(); // ILLEGAL: the lock is not held
    synchronized(MyClass.myLock) {
        a.toString(); // OK: the lock is held
    }
}
@Holding("MyClass.myLock")
void helper2(@GuardedBy("MyClass.myLock") Object b) {
    b.toString(); // OK: the lock is held
}
void helper3(Object c) {
    c.toString(); // OK: no lock constraints
}
void helper4(@GuardedBy("MyClass.myLock") Object d) {
    d.toString(); // ILLEGAL: the lock is not held
}
void myMethod2(@GuardedBy("MyClass.myLock") Object e) {
    helper1(e); // OK to pass to another routine without holding the lock
    e.toString(); // ILLEGAL: the lock is not held
    synchronized (MyClass.myLock) {
        helper2(e);
        helper3(e);
        helper4(e); // OK, but helper4's body still does not type-check
    }
}

```

7.1.2 Discussion of @Holding

A programmer might choose to use the `@Holding` method annotation in two different ways: to specify a higher-level protocol, or to summarize intended usage. Both of these approaches are useful, and the Lock checker supports both.

Higher-level synchronization protocol `@Holding` can specify a higher-level synchronization protocol that is not expressible as locks over Java objects. By requiring locks to be held, you can create higher-level protocol primitives without giving up the benefits of the annotations and checking of them.

Method summary that simplifies reasoning `@Holding` can be a method summary that simplifies reasoning. In this case, the `@Holding` doesn't necessarily introduce a new correctness constraint; the program might be correct even if the lock were acquired later in the body of the method or in a method it calls, so long as the lock is acquired before accessing the data it protects.

Rather, here `@Holding` expresses a fact about execution: when execution reaches this point, the following locks are already held. This fact enables people and tools to reason intra- rather than inter-procedurally.

In Java, it is always legal to re-acquire a lock that is already held, and the re-acquisition always works. Thus, whenever you write

```

@Holding("myLock")
void myMethod() {
    ...
}

```

it would be equivalent, from the point of view of which locks are held during the body, to write

```

void myMethod() {
    synchronized (myLock) { // no-op: re-acquire a lock that is already held
        ...
    }
}

```

The advantages of the `@Holding` annotation include:

- The annotation documents the fact that the lock is intended to already be held.
- The Lock Checker enforces that the lock is held when the method is called, rather than masking a programmer error by silently re-acquiring the lock.
- The `synchronized` statement can deadlock if, due to a programmer error, the lock is not already held. The Lock Checker prevents this type of error.
- The annotation has no run-time overhead. Even if the lock re-acquisition succeeds, it still consumes time.

7.1.3 Relationship to annotations in *Java Concurrency in Practice*

The book *Java Concurrency in Practice* [GPB⁺06] defines a `@GuardedBy` annotation that is the inspiration for ours. The book's `@GuardedBy` serves two related purposes:

- When applied to a field, it means that the given lock must be held when accessing the field. The lock acquisition and the field access may be arbitrarily far in the future.
- When applied to a method, it means that the given lock must be held by the caller at the time that the method is called — in other words, at the time that execution passes the `@GuardedBy` annotation.

One rationale for reusing the annotation name for both purposes in JCIP is that there are fewer annotations to learn. Another rationale is that both variables and methods are “members” that can be “accessed”; variables can be accessed by reading or writing them (`putfield`, `getfield`), and methods can be accessed by calling them (`invokevirtual`, `invokeinterface`). In both cases, `@GuardedBy` creates preconditions for accessing so-annotated members. This informal intuition is inappropriate for a tool that requires precise semantics.

The Lock checker renames the method annotation to `@Holding`, and it generalizes the `@GuardedBy` annotation into a type qualifier that can apply not just to a field but to an arbitrary type (including the type of a parameter, return value, local variable, generic type parameter, etc.). This makes the annotations more expressive and also more amenable to automated checking. It also accommodates the distinct (though related) meanings of the two annotations.

Chapter 8

Fake Enum checker

Java's `enum` keyword lets you define an enumeration type: a finite set of distinct values that are related to one another but are disjoint from all other types, including other enumerations. Before enums were added to Java, there were two ways to encode an enumeration, both of which are error-prone:

the fake enum pattern a set of `int` or `String` constants (as often found in older C code).

the typesafe enum pattern a class with private constructor.

Sometimes you need to use the fake enum pattern, rather than a real enum or the typesafe enum pattern. One reason is backward-compatibility. A public API that predates Java's `enum` keyword may use `int` constants; it cannot be changed, because doing so would break existing clients. For example, Java's JDK still uses `int` constants in the AWT and Swing frameworks. Another reason is performance, especially in environments with limited resources. Use of an `int` instead of an object can reduce code size, memory requirements, and run time.

In cases when code has to use the fake enum pattern, the fake enum (Fenum) checker gives the same safety guarantees as a true enumeration type. The developer can introduce new types that are distinct from all values of the base type and from all other fake enums. Fenums can be introduced for primitive types as well as for reference types.

Figure 8.1 shows part of the type hierarchy for the Fenum type system.

8.1 Fake enum annotations

The checker supports two ways to introduce a new fake enum (fenum):

1. Introduce your own specialized fenum annotation with code like this in file `MyFenum.java`:

```
package myprojectquals;  
  
import java.lang.annotation.*;
```




Figure 8.1: Partial type hierarchy for the Fenum type system. There are two forms of fake enumeration annotations — above, illustrated by `@Fenum("A")` and `@FenumC`. See section 8.1 for descriptions of how to introduce both types of fenums. The type qualifiers in gray (`@FenumTop`, `@FenumUnqualified`, and `@Bottom`) should never be written in source code; they are used internally by the type system.

```
import checkersquals.SubtypeOf;
import checkersquals.TypeQualifier;
```

```
@Documented
@Retention(RetentionPolicy.RUNTIME)
@TypeQualifier
@SubtypeOf( { FenumTop.class } )
public @interface MyFenum {}
```

You only need to adapt the italicized package, annotation, and file names in the example.

2. Use the provided `@Fenum` annotation, which takes a `String` argument to distinguish different fenums. For example, `@Fenum("A")` and `@Fenum("B")` are two distinct fenums.

The first approach allows you to define a short, meaningful name suitable for your project, whereas the second approach allows quick prototyping.

8.2 What the Fenum checker checks

The Fenum checker ensures that unrelated types are not mixed. All types with a particular fenum annotation, or `@Fenum(...)` with a particular `String` argument, are disjoint from all unannotated types and all types with a different fenum annotation or `String` argument.

The checker forbids method calls on fenum types and ensures that only compatible fenum types are used in comparisons and arithmetic operations (if applicable to the annotated type).

It is the programmer's responsibility to ensure that fields with a fenum type are properly initialized before use. Otherwise, one might observe a `null` reference or zero value in the field of a fenum type. (The Nullness checker (Chapter 3, page 19) can prevent failure to initialize a reference variable.)

8.3 Running the Fenum checker

The Fenum checker can be invoked by running the following commands.

- If you define your own annotation, provide the name of the annotation using the `-Aquals` option:

```
javac -processor checkers.fenum.FenumChecker
      -Aquals=myproject.quals.MyFenum MyFile.java ...
```

- If your code uses the `@Fenum` annotation, you do not need the `-Aquals` option:

```
javac -processor checkers.fenum.FenumChecker MyFile.java ...
```

8.4 Suppressing warnings

One example of when you need to suppress warnings is when you initialize the fenum constants to literal values. To remove this warning message, add a `@SuppressWarnings` annotation to either the field or class declaration, for example:

```
@SuppressWarnings("fenum:assignment.type.incompatible")
class MyConsts {
    public static final @Fenum("A") int ACONST1 = 1;
    public static final @Fenum("A") int ACONST2 = 2;
}
```

8.5 Example

The following example introduces two fenums in class `TestStatic` and then performs a few typical operations.

```
@SuppressWarnings("fenum:assignment.type.incompatible") // for initialization
public class TestStatic {
    public static final @Fenum("A") int ACONST1 = 1;
    public static final @Fenum("A") int ACONST2 = 2;

    public static final @Fenum("B") int BCONST1 = 4;
    public static final @Fenum("B") int BCONST2 = 5;
}

class FenumUser {
    @Fenum("A") int state1 = TestStatic.ACONST1; // ok
    @Fenum("B") int state2 = TestStatic.ACONST1; // Incompatible fenums forbidden!

    void fenumArg(@Fenum("A") int p) {}

    void foo() {
        state1 = 4; // Direct use of value forbidden!
        state1 = TestStatic.BCONST1; // Incompatible fenums forbidden!
        state1 = TestStatic.ACONST2; // ok

        fenumArg(5); // Direct use of value forbidden!
        fenumArg(TestStatic.BCONST1); // Incompatible fenums forbidden!
        fenumArg(TestStatic.ACONST1); // ok
    }
}
```

8.6 References

- **Java Language Specification on enums:**
<http://docs.oracle.com/javase/specs/jls/se7/html/jls-8.html#jls-8.9>
- **Tutorial trail on enums:**
<http://docs.oracle.com/javase/tutorial/java/javaOO/enum.html>
- **Typesafe enum pattern:**
<http://www.oracle.com/technetwork/java/page1-139488.html>
- **Java Tip 122: Beware of Java typesafe enumerations:**
<http://www.javaworld.com/javaworld/jwatips/jw-jwatip122.html>

Chapter 9

Tainting checker

The tainting checker prevents certain kinds of trust errors. A *tainted*, or untrusted, value is one that comes from an arbitrary, possibly malicious source, such as user input or unvalidated data. In certain parts of your application, using a tainted value can compromise the application’s integrity, causing it to crash, corrupt data, leak private data, etc.

For example, a user-supplied pointer, handle, or map key should be validated before being dereferenced. As another example, a user-supplied string should not be concatenated into a SQL query, lest the program be subject to a SQL injection attack. A location in your program where malicious data could do damage is called a *sensitive sink*.

A program must “sanitize” or “untaint” an untrusted value before using it at a sensitive sink. There are two general ways to untaint a value: by checking that it is innocuous/legal (e.g., it contains no characters that can be interpreted as SQL commands when pasted into a string context), or by transforming the value to be legal (e.g., quoting all the characters that can be interpreted as SQL commands). A correct program must use one of these two techniques so that tainted values never flow to a sensitive sink. The Tainting Checker ensures that your program does so.

If the Tainting Checker issues no warning for a given program, then no tainted value ever flows to a sensitive sink. However, your program is not necessarily free from all trust errors. As a simple example, you might have forgotten to annotate a sensitive sink as requiring an untainted type, or you might have forgotten to annotate untrusted data as having a tainted type.

To run the Tainting Checker, supply the `-processor checkers.tainting.TaintingChecker` command-line option to `javac`.

9.1 Tainting annotations

The Tainting type system uses the following annotations:

- `@Untainted` indicates a type that includes only untainted, trusted values.
- `@Tainted` indicates a type that may include only tainted, untrusted values. `@Tainted` is a supertype of `@Untainted`.
- `@PolyTainted` is a qualifier that is polymorphic over tainting (see Section 18.2).

9.2 Tips on writing `@Untainted` annotations

Most programs are designed with a boundary that surrounds sensitive computations, separating them from untrusted values. Outside this boundary, the program may manipulate malicious values, but no malicious values ever pass the boundary to be operated upon by sensitive computations.

In some programs, the area outside the boundary is very small: values are sanitized as soon as they are received from an external source. In other programs, the area inside the boundary is very small: values are sanitized only immediately before being used at a sensitive sink. Either approach can work, so long as every possibly-tainted value is sanitized before it reaches a sensitive sink.

Once you determine the boundary, annotating your program is easy: put `@Tainted` outside the boundary, `@Untainted` inside, and `@SuppressWarnings("tainting")` at the validation or sanitization routines that are used at the boundary.

The Tainting Checker's standard default qualifier is `@Tainted` (see Section 19.1.1 for overriding this default). This is the safest default, and the one that should be used for all code outside the boundary (for example, code that reads user input). You can set the default qualifier to `@Untainted` in code that may contain sensitive sinks.

The Tainting Checker does not know the intended semantics of your program, so it cannot warn you if you mis-annotate a sensitive sink as taking `@Tainted` data, or if you mis-annotate external data as `@Untainted`. So long as you correctly annotate the sensitive sinks and the places that untrusted data is read, the Tainting Checker will ensure that all your other annotations are correct and that no undesired information flows exist.

As an example, suppose that you wish to prevent SQL injection attacks. You would start by annotating the `Statement` class to indicate that the `execute` operations may only operate on untainted queries (Chapter 21 describes how to annotate external libraries):

```
public boolean execute(@Untainted String sql) throws SQLException;
public boolean executeUpdate(@Untainted String sql) throws SQLException;
```

9.3 `@Tainted` and `@Untainted` can be used for many purposes

The `@Tainted` and `@Untainted` annotations have only minimal built-in semantics. In fact, the Tainting Checker provides only a small amount of functionality beyond the Basic Checker (Section 15). This lack of hard-coded behavior means that the annotations can serve many different purposes. Here are just a few examples:

- Prevent SQL injection attacks: `@Tainted` is external input, `@Untainted` has been checked for SQL syntax.
- Prevent cross-site scripting attacks: `@Tainted` is external input, `@Untainted` has been checked for JavaScript syntax.
- Prevent information leakage: `@Tainted` is secret data, `@Untainted` may be displayed to a user.

In each case, you need to annotate the appropriate untainting/sanitization routines. This is similar to the `@Encrypted` annotation (Section 15.2), where the cryptographic functions are beyond the reasoning abilities of the type system. In each case, the type system verifies most of your code, and the `@SuppressWarnings` annotations indicate the few places where human attention is needed.

If you want more specialized semantics, or you want to annotate multiple types of tainting in a single program, then you can copy the definition of the Tainting Checker to create a new annotation and checker with a more specific name and semantics. See Chapter 22 for more details.

Chapter 10

Linear checker for preventing aliasing

The Linear Checker implements type-checking for a linear type system. A linear type system prevents aliasing: there is only one (usable) reference to a given object at any time. Once a reference appears on the right-hand side of an assignment, it may not be used any more. The same rule applies for pseudo-assignments such as procedure argument-passing (including as the receiver) or return.

One way of thinking about this is that a reference can only be used once, after which it is “used up”. This property is checked statically at compile time. The single-use property only applies to use in an assignment, which makes a new reference to the object; ordinary field dereferencing does not use up a reference.

By forbidding aliasing, a linear type system can prevent problems such as unexpected modification (by an alias), or ineffectual modification (after a reference has already been passed to, and used by, other code).

To run the Linear Checker, supply the `-processor checkers.Linear.LinearChecker` command-line option to `javac`.

Figure 10.1 gives an example of the Linear Checker’s rules.

10.1 Linear annotations

The linear type system uses one user-visible annotation: `@Linear`. The annotation indicates a type for which each value may only have a single reference — equivalently, may only be used once on the right-hand side of an assignment.

The full qualifier hierarchy for the linear type system includes three types:

- `@UsedUp` is the type of references whose object has been assigned to another reference. The reference may not be used in any way, including having its fields dereferenced, being tested for equality with `==`, or being assigned to another reference. Users never need to write this qualifier.
- `@Linear` is the type of references that have no aliases, and that may be dereferenced at most once in the future. The type of `new T()` is `@Linear T` (the analysis does not account for the slim possibility that an alias to `this` escapes the constructor).
- `@NonLinear` is the type of references that may be dereferenced, and aliases made, as many times as desired. This is the default, so users only need to write `@NonLinear` if they change the default.

`@UsedUp` is a supertype of `@NonLinear`, which is a supertype of `@Linear`.

This hierarchy makes an assignment like

```
@Linear Object l = new Object();
@NonLinear Object nl = l;
@NonLinear Object nl2 = nl;
```

legal. In other words, the fact that an object is referenced by a `@Linear` type means that there is only one usable reference to it *now*, not that there will *never* be multiple usable references to it. (The latter guarantee would be possible to enforce, but it is not what the Linear Checker does.)

```

class Pair {
    Object a;
    Object b;
    public String toString() {
        return "<" + String.valueOf(a) + ", " + String.valueOf(b) + ">";
    }
}

void print(@Linear Object arg) {
    System.out.println(arg);
}

@Linear Pair printAndReturn(@Linear Pair arg) {
    System.out.println(arg.a);
    System.out.println(arg.b);    // OK: field dereferencing does not use up the reference arg
    return arg;
}

@Linear Object m(Object o, @Linear Pair lp) {
    @Linear Object lo2 = o;        // ERROR: aliases may exist
    @Linear Pair lp3 = lp;
    @Linear Pair lp4 = lp;        // ERROR: reference lp was already used
    lp3.a;
    lp3.b;                        // OK: field dereferencing does not use up the reference
    print(lp3);
    print(lp3);                  // ERROR: reference lp3 was already used
    lp3.a;                       // ERROR: reference lp3 was already used
    @Linear Pair lp4 = new Pair(...);
    lp4.toString();
    lp4.toString();              // ERROR: reference lp4 was already used
    lp4 = new Pair();            // OK to reassign to a used-up reference
    // If you need a value back after passing it to a procedure, that
    // procedure must return it to you.
    lp4 = printAndReturn(lp4);
    if (...) {
        print(lp4);
    }
    if (...) {
        return lp4;              // ERROR: reference lp4 may have been used
    } else {
        return new Object();
    }
}

```

Figure 10.1: Example of Linear Checker rules.

10.2 Limitations

The `@Linear` annotation is supported and checked only on method parameters (including the receiver), return types, and local variables. Supporting `@Linear` on fields would require a sophisticated alias analysis or type system, and is future work.

No annotated libraries are provided for linear types. Most libraries would not be able to use linear types in their purest form. For example, you cannot put a linearly-typed object in a hashtable, because hashtable insertion calls `hashCode`; `hashCode` uses up the reference and does not return the object, even though it does not retain any pointers to the object. For similar reasons, a collection of linearly-typed objects could not be sorted or searched.

Our lightweight implementation is intended for use in the parts of your program where errors relating to aliasing and object reuse are most likely. You can use manual reasoning (and possibly an unchecked cast or warning suppression) when objects enter or exit those portions of your program, or when that portion of your program uses an unannotated library.

Chapter 11

Regex checker for regular expression syntax

The Regex Checker prevents, at compile-time, use of syntactically invalid regular expressions and access of invalid capturing groups.

A regular expression, or `regex`, is a pattern for matching certain strings of text. In Java, a programmer writes a regular expression as a string. At run time, the string is “compiled” into an efficient internal form (`Pattern`) that is used for text-matching. Regular expression in Java also have capturing groups, which are delimited by parentheses and allow for extraction from text.

The syntax of regular expressions is complex, so it is easy to make a mistake. It is also easy to accidentally use a regex feature from another language that is not supported by Java (see section “Comparison to Perl 5” in the `Pattern` Javadoc). Ordinarily, the programmer does not learn of these errors until run time. The Regex checker warns about these problems at compile time.

To run the Regex Checker, supply the `-processor checkers.regex.RegexChecker` command-line option to `javac`.

11.1 Regex annotations

These qualifiers make up the Regex type system:

@Regex indicates valid regular expression `Strings`. This qualifier takes an optional parameter of at the least the number of capturing groups in the regular expression. If not provided, the parameter defaults to 0.

@PolyRegex indicates qualifier polymorphism. For a description of `@PolyRegex`, see Section 18.2.

The subtyping hierarchy of the Regex checker’s qualifiers is shown in Figure 11.1.

11.2 Annotating your code with @Regex

11.2.1 Implicit qualifiers

As described in Section 19.1, the Regex checker adds implicit qualifiers, reducing the number of annotations that must appear in your code. The checker implicitly adds the `Regex` qualifier with the parameter set to the correct number of capturing groups to any `String` literal that is a valid regex. The Regex checker allows the `null` literal to be assigned to any type qualified with the `Regex` qualifier.

In rare cases, you will need to override the implicit qualifiers, most notably because of type inference for type parameters. For example, the following code does not type-check

```
List<String> list1 = Arrays.asList("a", "b", "c");
```

because of invariant subtyping for type parameters, which is the same reason that




Figure 11.1: The subtyping relationship of the Regex checker’s qualifiers. Because the parameter to a `@Regex` qualifier is at least the number of capturing groups in a regular expression, a `@Regex` qualifier with more capturing groups is a subtype of a `@Regex` qualifier with fewer capturing groups. Qualifiers in gray are used internally by the type system but should never be written by a programmer.

```
List<Object> objlist = Arrays.asList(2.718, 3.14);
```

does not type-check in ordinary Java.

Depending on your intention, you should rewrite your code in one of the following ways (alternately, you can just suppress the warning):

```
List<@Regex String> list2 = Arrays.asList("a", "b", "c");
```

```
List<String> list3 = Arrays.<String>asList("a", "b", "c");
```

```
List<? extends String> list4 = Arrays.asList("a", "b", "c");
```

For the latter two variants, you might want to add a comment explaining that you used `<String>asList()` instead of just `aslist()`, or `List<? extends String>` instead of just `List<String>`, because of invariant subtyping for type parameters and `@Regex` type inference. That will keep some other developer from undoing your change because it looks silly.

11.2.2 Capturing groups

The Regex checker validates that a legal capturing group number is passed to `Matcher`’s `group`, `start` and `end` methods. To do this, the type of `Matcher` must be qualified with a `@Regex` annotation with the number of capturing groups in the regular expression. This is handled implicitly by the Regex checker for local variables (see Section 19.1.2), but you may need to add `@Regex` annotations with a capturing group count to `Pattern` and `Matcher` fields and parameters.

```

public @Regex String parenthesize(@Regex String regex) {
    return "(" + regex + " "; // Even though the parentheses are not @Regex Strings, the checker still treats
                               // the whole expression as a @Regex String
}

```

Figure 11.2: An example of the Regex Checker’s support for concatenation of non-regular expression Strings to produce valid regular expression Strings.

```

String regex = getRegexFromUser();
if (! RegexUtil.isRegex(regex)) {
    throw new RuntimeException("Error parsing regex " + regex, RegexUtil.regexException(regex));
    // or: System.out.println("Error parsing regex \"" + regex + "\": " + RegexUtil.regexError(regex));
}
// The following line suppresses a Regex Checker warning and is only necessary until the
// Regex Checker supports flow-sensitivity, after which time it can be removed from the code.
regex = RegexUtil.asRegex(regex); // @SuppressWarnings("regex") // flow-sensitivity
Pattern p = Pattern.compile(regex);

```

Figure 11.3: Example use of RegexUtil methods.

11.2.3 Concatenation of partial regular expressions

The Regex Checker supports concatenation of non-regular expression Strings that produce valid regular expression Strings. For an example see Figure 11.2.

11.2.4 Testing whether a string is a regular expression

Sometimes, the Regex Checker cannot infer whether a particular expression is a regular expression — and sometimes your code cannot either! In these cases, you can use the `isRegex` method to perform such a test, and other helper methods to provide useful error messages. A common use is for user-provided regular expressions (such as ones passed on the command-line). Figure 11.3 gives an example of the intended use of the `RegexUtil` methods.

RegexUtil.isRegex returns `true` if its argument is a valid regular expression.

RegexUtil.regexError returns a `String` error message if its argument is not a valid regular expression, or `null` if its argument is a valid regular expression.

RegexUtil.regexException returns the `PatternSyntaxException` that `Pattern.compile(String)` throws when compiling an invalid regular expression. It returns `null` if its argument is a valid regular expression.

An additional version of each of these methods is also provided that takes an additional group count parameter. The `RegexUtil.isRegex` method verifies that the argument has at least the given number of groups. The `RegexUtil.regexError` and `RegexUtil.regexException` methods return a `String` error message and `PatternSyntaxException`, respectively, detailing why the given `String` is not a syntactically valid regular expression with at least the given number of capturing groups.

A potential disadvantage of using these methods is that your code becomes dependent on the Checker Framework at run time as well as at compile time. You can avoid this by copying the implementation of these methods into your own code.

11.2.5 Suppressing warnings

If you are positive that a particular string that is being used as a regular expression is syntactically valid, but the Regex Checker cannot conclude this and issues a warning about possible use of an invalid regular expression, then you can use the `RegexUtil.asRegex` method to suppress the warning.

You can think of this method as a cast: it returns its argument unchanged, but with the type `@Regex String` if it is a valid regular expression. It throws an `Error` if its argument is not a valid regular expression, but you should only use it when you are sure it will not throw an error.

There is an additional `RegexUtil.asRegex` method that takes a capturing group parameter. This method works the same as described above, but returns a `@Regex String` with the parameter on the annotation set to the value of the capturing group parameter passed to the method.

This method is mainly a workaround until the `Regex Checker` supports flow-sensitivity (see Section 19.1.2) and should be used rarely once the `Regex Checker` supports flow-sensitivity.

Chapter 12

Property file checker

The property file checker ensures that a property file or resource bundle (both of which act like maps from keys to values) is only accessed with valid keys. Accesses without a valid key either return `null` or a default value, which can lead to a `NullPointerException` or hard-to-trace behavior. The property file checker (Section 12.1, page 56) ensures that the used keys are found in the corresponding property file or resource bundle.

We also provide two specialized checkers. An internationalization checker (Section 12.2, page 57) verifies that code is properly internationalized. A compiler message key checker (Section 12.3, page 57) verifies that compiler message keys used in the Checker Framework are declared in a property file; This is an example of a simple specialization of the property file checker, and the Checker Framework source code shows how it is used.

It is easy to customize the property key checker for other related purposes. Take a look at the source code of the compiler message key checker and adapt it for your purposes.

12.1 Generic property file checker

The generic property file checker ensures that a resource key is located in a specified property file or resource bundle.

The annotation `@PropertyKey` indicates that the qualified `String` is a valid key found in the property file or resource bundle. You do not need to annotate `String` literals. The checker looks up every `String` literal in the specified property file or resource bundle, and adds annotations as appropriate.

If you pass a `String` variable to be eventually used as a key, you also need to annotate all these variables with `@PropertyKey`.

The checker can be invoked by running the following command:

```
javac -processor checkers.propkey.PropertyKeyChecker
      -Abundlenames=MyResource MyFile.java ...
```

You must specify the resources, which map keys to strings. The checker supports two types of resource: resource bundles and property files. You can specify one or both of the following two command-line options:

1. `-Abundlenames=resource_name`
resource_name is the name of the resource to be used with `ResourceBundle.getBundle()`. The checker uses the default `Locale` and `ClassLoader` in the compilation system. (For a tutorial about `ResourceBundles`, see <http://docs.oracle.com/javase/tutorial/ui/features/i18n.html>.) Multiple resource bundle names are separated by colons `':'`.
2. `-Apropfiles=prop_file`
prop_file is the name of a properties file that maps keys to values. The file format is described in the Javadoc for `Properties.load()`. Multiple files are separated by colons `':'`.

12.2 Internationalization checker

The Internationalization Checker verifies that your code is properly internationalized. Internationalization is the process of adapting software to different languages and locales. Internationalization is sometimes called localization (though the terms are not identical), and is sometimes called `i18n` (because the word starts with “i”, ends with “n”, and has 18 characters in between; localization is similarly sometimes abbreviated as `l10n`).

The checker focuses on one aspect of internationalization: user-visible strings should be presented in the user’s own language, such as English, French, or German. This is achieved by looking up keys in a localization resource, which maps keys to user-visible strings. For instance, one version of a resource might map `"CANCEL_STRING"` to `"Cancel"`, and another version of the same resource might map `"CANCEL_STRING"` to `"Abbrechen"`.

There are other aspects to localization, such as formatting of dates (`3/5` vs. `5/3` for March 5), that the checker does not check.

The Internationalization Checker verifies these two properties:

1. Any user-visible text should be obtained from a localization resource. For example, `String` literals should not be output to the user.
2. When looking up keys in a localization resource, the key should exist in that resource. This check catches incorrect or misspelled localization keys.

12.2.1 Internationalization annotations

The Internationalization Checker supports two annotations:

1. `@Localized`: indicates that the qualified `String` is a message that has been localized and/or formatted with respect to the used locale.
2. `@LocalizableKey`: indicates that the qualified `String` or `Object` is a valid key found in the localization resource. This annotation is a specialization of the `@PropertyKey` annotation, that gets checked by the generic property key checker.

You may need to add the `@Localized` annotation to more methods in the JDK or other libraries, or in your own code.

12.2.2 Running the Internationalization Checker

The Internationalization Checker can be invoked by running the following command:

```
javac -processor checkers.i18n.I18nChecker -Abundlenames=MyResource MyFile.java ...
```

You must specify the localization resource, which maps keys to user-visible strings. Like the generic property key checker, the internationalization checker supports two types of localization resource: `ResourceBundles` using the `-Abundlenames=resource_name` option or property files using the `-Apropfiles=prop_file` option.

12.3 Compiler Message Key checker

The Checker Framework uses compiler message keys to output error messages. These keys are substituted by localized strings for user-visible error messages. Using keys instead of the localized strings in the source code enables easier testing, as the expected error keys can stay unchanged while the localized strings can still be modified. We use the compiler message key checker to ensure that all internal keys are correctly localized. Instead of using the property file checker, we use a specialized checker, giving us more precise documentation of the intended use of `Strings`.

The single annotation used by this checker is `@CompilerMessageKey`. The Checker Framework is completely annotated; for example, class `checkers.source.Result` uses `@CompilerMessageKey` in methods `failure` and `warning`. For most users of the Checker Framework there will be no need to annotate any `Strings`, as the checker looks up all `String` literals and adds annotations as appropriate.

The compiler message key checker can be invoked by running the following command:

```
javac -processor checkers.compilermsgs.CompilerMessagesChecker  
      -Apropfiles=messages.properties MyFile.java ...
```

You must specify the resource, which maps compiler message keys to user-visible strings. The checker supports the same options as the generic property key checker. Within the Checker Framework we only use property files, so the `-Apropfiles=prop_file` option should be used.

Chapter 13

Signature checker for string representations of types

The Signature String Checker, or Signature Checker for short, verifies that string representations of types and signatures are used correctly.

Java defines multiple different string representations, and it is easy to misuse them or to miss bugs during testing. Using the wrong string format leads to a run-time exception or an incorrect result. This is a particular problem for fully qualified and binary names, which are nearly the same — they differ only for nested classes and arrays.

13.1 Signature annotations

Java defines three main formats for the string representation of a type. There is an annotation for each of these representations, plus one more. Figure 13.1 shows how they are related.

@FullyQualifiedName A *fully qualified name* (JLS §6.7), such as `package.Outer.Inner`, is used in Java code and in messages to the user.

@BinaryName A *binary name* (JLS §13.1), such as `package.Outer$Inner`, is the representation of a type in its own `.class` file.

@FieldDescriptor A *field descriptor* (JVMS §4.3.2), such as `Lpackage/Outer$Inner;`, is used in a `.class` file’s constant pool, for example to refer to other types; it abbreviates primitives and arrays, and uses internal form (JVMS §4.2) for class names.

@ClassGetName The type representation used by the `Class.getName()`, `Class.forName(String)`, and `Class.forName(String, boolean, ClassLoader)` methods. This format is: for any non-array type, the binary name; and for any array type, a format like the `@link FieldDescriptor` field descriptor, but using “.” where the field descriptor uses “/”.

@SourceName A source name is a string that is a valid fully qualified name *and* a valid binary name. A programmer should never or rarely use this — you should know how you intend to use a given variable. The checker




Figure 13.1: Partial type hierarchy for the Signature type system, showing string representations of a Java type. Programmers only need to write the boldfaced qualifiers, in the second row; qualifiers below those are included to improve the internal handling of String literals.

infers it for literal strings such as "package.MyClass" that are valid in both formats, and you might occasionally see it in an error message. Likewise, you might see other types such as `SourceNameForNonArray`, `BinaryNameForNonArray`, and `FieldDescriptorForArray`, but you generally should not use them either.

Java also defines other string formats for a type: simple names (JLS §6.2), qualified names (JLS §6.2), and canonical names (JLS §6.7). The Signature Checker does not include annotations for these.

Here are examples of the supported formats:

fully-qualified name	binary name	Class.getName	field descriptor
int	int	int	I
int[][]	int[][]	[[I	[[I
MyClass	MyClass	MyClass	LMyClass;
MyClass[]	MyClass[]	[LMyClass;	[LMyClass;
java.lang.Integer	java.lang.Integer	java.lang.Integer	Ljava/lang/Integer;
java.lang.Integer[]	java.lang.Integer[]	[Ljava.lang.Integer;	[Ljava/lang/Integer;
package.Outer.Inner	package.Outer\$Inner	package.Outer\$Inner	Lpackage/Outer\$Inner;
package.Outer.Inner[]	package.Outer\$Inner[]	[Lpackage.Outer\$Inner;	[Lpackage/Outer\$Inner;

Java defines one format for the string representation of a method signature:

@MethodDescriptor A *method descriptor* (JVMS §4.3.3) identifies a method's signature (its parameter and return types), just as a field descriptor identifies a type. The method descriptor for the method

```
Object mymethod(int i, double d, Thread t)
```

is

```
(IDLjava/lang/Thread;)Ljava/lang/Object;
```

13.2 What the Signature Checker checks

Certain methods in the JDK, such as `Class.forName`, are annotated indicating the type they require. The Signature Checker ensures that clients call them with the proper arguments. The Signature Checker does not reason about string operations such as concatenation, substring, parsing, etc.

To run the Signature Checker, supply the `-processor checkers.signature.SignatureChecker` command-line option to `javac`.

Chapter 14

Units checker

For many applications, it is important to use the correct units of measurement for primitive types. For example, NASA's Mars Climate Orbiter (cost: \$327 million) was lost because of a discrepancy between use of the metric unit Newtons and the imperial measure Pound-force.

The *Units Checker* ensures consistent usage of units. For example, consider the following code:

```
@m int meters = 5 * UnitsTools.m;  
@s int secs = 2 * UnitsTools.s;  
@mPERs int speed = meters / secs;
```

Due to the annotations `@m` and `@s`, the variables `meters` and `secs` are guaranteed to contain only values with meters and seconds as units of measurement. Utility class `UnitsTools` provides constants with which unqualified integer are multiplied to get values of the corresponding unit. The assignment of an unqualified value to `meters`, as in `meters = 99`, will be flagged as an error by the Units Checker.

The division `meters/secs` takes the types of the two operands into account and determines that the result is of type meters per second, signified by the `@mPERs` qualifier. We provide an extensible framework to define the result of operations on units.

14.1 Units annotations

The checker currently supports two varieties of units annotations: kind annotations (`@Length`, `@Mass`, ...) and the SI units (`@m`, `@kg`, ...).

Kind annotations can be used to declare what the expected unit of measurement is, without fixing the particular unit used. For example, one could write a method taking a `@Length` value, without specifying whether it will take meters or kilometers. The following kind annotations are defined:

```
@Area  
@Current  
@Length  
@Luminance  
@Mass  
@Speed  
@Substance  
@Temperature  
@Time
```

For each kind of unit, the corresponding SI unit of measurement is defined:

1. For `@Area`: the derived units square millimeters `@mm2`, square meters `@m2`, and square kilometers `@km2`

2. For `@Current: Ampere @A`
3. For `@Length: Meters @m` and the derived units `millimeters @mm` and `kilometers @km`
4. For `@Luminance: Candela @cd`
5. For `@Mass: kilograms @kg` and the derived unit `grams @g`
6. For `@Speed: meters per second @mPERs` and `kilometers per hour @kmPERh`
7. For `@Substance: Mole @mol`
8. For `@Temperature: Kelvin @K` and the derived unit `Celsius @C`
9. For `@Time: seconds @s` and the derived units `minutes @min` and `hours @h`

You may specify SI unit prefixes, using enumeration `Prefix`. The basic SI units (`@s`, `@m`, `@g`, `@A`, `@K`, `@mol`, `@cd`) take an optional `Prefix` enum as argument. For example, to use nanoseconds as unit, you could use `@s(Prefix.nano)` as a unit type. Furthermore, `@mm` is equivalent to `@m(Prefix.milli)`.

Class `UnitsTools` contains a constant for each SI unit. To create a value of the particular unit, multiply an unqualified value with one of these constants. By using static imports, this allows very natural notation; for example, after statically importing `UnitsTools.m`, the expression `5 * m` represents five meters. As all these unit constants are public, static, and final with value one, the compiler will optimize away these multiplications.

14.2 Extending the Units Checker

You can create new kind annotations and unit annotations that are specific to the particular needs of your project. An easy way to do this is by copying and adapting an existing annotation. (In addition, search for all uses of the annotation's name throughout the Units Checker implementation, to find other code to adapt; read on for details.)

Here is an example of a new unit annotation.

```
@Documented
@Retention(RetentionPolicy.RUNTIME)
@TypeQualifier
@SubtypeOf( { Time.class } )
@UnitsMultiple(quantity=s.class, prefix=Prefix.nano)
public @interface ns {}
```

The `@SubtypeOf` meta-annotation specifies that this annotation introduces an additional unit of time. The `@UnitsMultiple` meta-annotation specifies that this annotation should be a nano multiple of the basic unit `@s: @ns` and `@s(Prefix.nano)` behave equivalently and interchangeably. Most annotation definitions do not have a `@UnitsMultiple` meta-annotation.

To take full advantage of the additional unit qualifier, you need to do two additional steps. (1) Provide constants that convert from unqualified types to types that use the new unit. See class `UnitsTools` for examples (you will need to suppress a checker warning in just those few locations). (2) Put the new unit in relation to existing units. Provide an implementation of the `UnitsRelations` interface as a meta-annotation to one of the units.

See demonstration `demos/units-extension/` for an example extension that defines Hertz (hz) as scalar per second, and defines an implementation of `UnitsRelations` to enforce it.

14.3 What the Units Checker checks

The Units Checker ensures that unrelated types are not mixed.

All types with a particular unit annotation are disjoint from all unannotated types, from all types with a different unit annotation, and from all types with the same unit annotation but a different prefix.

Subtyping between the units and the unit kinds is taken into account, as is the `@UnitsMultiple` meta-annotation.

Multiplying a scalar with a unit type results in the same unit type.

The division of a unit type by the same unit type results in the unqualified type.

Multiplying or dividing different unit types, for which no unit relation is known to the system, will result in a `MixedUnits` type, which is separate from all other units. If you encounter a `MixedUnits` annotation in an error message, ensure that your operations are performed on correct units or refine your `UnitsRelations` implementation.

14.4 Running the Units Checker

The Units Checker can be invoked by running the following commands.

- If your code uses only the SI units that are provided by the framework, simply invoke the checker:

```
javac -processor checkers.units.UnitsChecker MyFile.java ...
```

- If you define your own units, provide the name of the annotations using the `-Aunits` option:

```
javac -processor checkers.units.UnitsChecker -Aunits=myprojectquals.MyUnit,myprojectquals
```

14.5 Suppressing warnings

One example of when you need to suppress warnings is when you initialize a variable with a unit type by a literal value. To remove this warning message, it is best to introduce a constant that represents the unit and to add a `@SuppressWarnings` annotation to that constant. For examples, see class `UnitsTools`.

14.6 References

- The GNU Units tool provides a comprehensive list of units:
<http://www.gnu.org/software/units/>
- The F# units of measurement system inspired some of our syntax:
http://en.wikibooks.org/wiki/F_Sharp_Programming/Units_of_Measure

Chapter 15

Basic checker

The Basic checker enforces only subtyping rules. It operates over annotations specified by a user on the command line. Thus, users can create a simple type checker without writing any code beyond definitions of the type qualifier annotations.

The Basic checker can accommodate all of the type system enhancements that can be declaratively specified (see Chapter 22). This includes type introduction rules (implicit annotations, e.g., literals are implicitly considered `@NonNull`) via the `@ImplicitFor` meta-annotation, and other features such as flow-sensitive type qualifier inference (Section 19.1.2) and qualifier polymorphism (Section 18.2).

The Basic checker is also useful to type system designers who wish to experiment with a checker before writing code; the Basic checker demonstrates the functionality that a checker inherits from the Checker Framework.

If you need typestate analysis, then you can extend a typestate checker, much as you would extend the Basic Checker if you do not need typestate analysis. For more details (including a definition of “typestate”), see Chapter 16.

For type systems that require special checks (e.g., warning about dereferences of possibly-null values), you will need to write code and extend the framework as discussed in Chapter 22.

15.1 Using the Basic checker

The Basic checker is used in the same way as other checkers (using the `-processor checkers.basic.BasicChecker` option; see Chapter 2), except that it requires an additional annotation processor argument via the standard “-A” switch:

- `-Aquals`: this option specifies a comma-no-space-separated list of the fully-qualified class names of the annotations used as qualifiers in the custom type system. For example,

```
javac -processor checkers.fenum.BasicChecker  
      -Aquals=myproject.quals.MyQual,myproject.quals.OtherQual MyFile.java ...
```

It serves the same purpose as the `@TypeQualifiers` annotation used by other checkers (see section 22.6).

The annotations listed in `-Aquals` must be accessible to the compiler during compilation in the classpath. In other words, they must already be compiled before you run the Basic checker with `javac`; it is not sufficient to supply their source files on the command line.

To suppress a warning issued by the basic checker, use a `@SuppressWarnings` annotation, with the argument being the unqualified, uncapitalized name of any of the annotations passed to `-Aquals`. This will suppress all warnings, regardless of which of the annotations is involved in the warning. (As a matter of style, you should choose one of the annotations as your `@SuppressWarnings` key and stick with it for that entire type hierarchy.)

15.2 Basic checker example

Consider a hypothetical `Encrypted` type qualifier, which denotes that the representation of an object (such as a `String`, `CharSequence`, or `byte[]`) is encrypted. To use the Basic checker for the `Encrypted` type system, follow three steps.

1. Define an annotation for the Encrypted qualifier:

```
package myquals;

import checkersquals.*;

/**
 * Denotes that the representation of an object is encrypted.
 * ...
 */
@TypeQualifier
@SubtypeOf(Unqualified.class)
@Target({ElementType.PARAMETER, ElementType.USE})
public @interface Encrypted {}
```

Don't forget to compile this class:

```
$ javac myquals/Encrypted.java
```

The resulting .class file should either be on your classpath, or on the processor path (set via the `-processorpath` command-line option to javac).

2. Write @Encrypted annotations in your program (YourProgram.java):

```
import myquals.Encrypted;

...

public @Encrypted String encrypt(String text) {
    // ...
}

// Only send encrypted data!
public void sendOverInternet(@Encrypted String msg) {
    // ...
}

void sendText() {
    // ...
    @Encrypted String ciphertext = encrypt(plaintext);
    sendOverInternet(ciphertext);
    // ...
}

void sendPassword() {
    String password = getUserPassword();
    sendOverInternet(password);
}
```

You may also need to add @SuppressWarnings annotations to the encrypt and decrypt methods. Analyzing them is beyond the capability of any realistic type system.

3. Invoke the compiler with the Basic checker, specifying the @Encrypted annotation using the `-Aquals` option. You should add the Encrypted classfile to the processor classpath:

```
$ javac -processorpath myqualspath -processor checkers.basic.BasicChecker \
    -Aquals=myquals.Encrypted YourProgram.java
```

```
YourProgram.java:42: incompatible types.  
found   : java.lang.String  
required: @myquals.Encrypted java.lang.String  
    sendOverInternet(password);  
                    ^
```

Chapter 16

Typestate checker

In a regular type system, a variable has the same type throughout its scope. In a typestate system, a variable's type can change as operations are performed on it.

The most common example of typestate is for a `File` object. Assume a file can be in two states, `@Open` and `@Closed`. Calling the `close()` method changes the file's state. Any subsequent attempt to read, write, or close the file will lead to a run-time error. It would be better for the type system to warn about such problems, or guarantee their absence, at compile time.

Just as you can extend the Basic Checker to create a type checker, you can extend a typestate checker to create a type checker that supports typestate analysis. Two extensible typestate analyses that build on the Checker Framework are available. One is by Adam Warski: <http://www.warski.org/typestate.html>. The other is by Daniel Wand: <http://typestate.ewand.de/>.

16.1 Comparison to flow-sensitive type refinement

The Checker Framework's flow-sensitive type refinement (Section 19.1.2) implements a form of typestate analysis. For example, after code that tests a variable against null, the Nullness Checker (Chapter 3) treats the variable's type as `@NonNull T`, for some `T`.

For many type systems, flow-sensitive type refinement is sufficient. But sometimes, you need full typestate analysis. This section compares the two. (Dependent types and unused variables (Section 19.3) also have similarities with typestate analysis and can occasionally substitute for it. For brevity, this discussion omits them.)

A typestate analysis is easier for a user to create or extend. Flow-sensitive type refinement is built into the Checker Framework and is optionally extended by each checker. Modifying the rules requires writing Java code in your checker. By contrast, it is possible to write a simple typestate checker declaratively, by writing annotations on the methods (such as `close()`) that change a reference's typestate.

A typestate analysis can change a reference's type to something that is not consistent with its original definition. For example, suppose that a programmer decides that the `@Open` and `@Closed` qualifiers are incomparable — neither is a subtype of the other. A typestate analysis can specify that the `close()` operation converts an `@Open File` into a `@Closed File`. By contrast, flow-sensitive type refinement can only give a new type that is a subtype of the declared type — for flow-sensitive type refinement to be effective, `@Closed` would need to be a child of `@Open` in the qualifier hierarchy (and `close()` would need to be treated specially by the checker).

Chapter 17

Third-party checkers

The checker framework has been used to build other checkers that are not distributed together with the framework. This chapter mentions just a few of them. If you want a reference to your checker included in this chapter, send us a link and a short description.

17.1 Units and dimensions checker

A checker for units and dimensions is available at <http://www.lexspoon.org/expannots/>.

Unlike the Units Checker that is distributed with the Checker Framework (see Section 14), this checker includes dynamic checks and permits annotation arguments that are Java expressions. This added flexibility, however, requires that you use a special version both of the Checker Framework and of the Type Annotations compiler.

17.2 Thread locality checker

Loci, a checker for thread locality, is available at <http://www.it.uu.se/research/upmarc/loci/>. Developer resources are available at the project page <http://java.net/projects/loci/>.

17.3 Safety-Critical Java checker

A checker for Safety-Critical Java (SCJ, JSR 302) is available at <http://sss.cs.purdue.edu/projects/oscj/checker/checker.html>. Developer resources are available at the project page <http://code.google.com/p/scj-jsr302/>.

17.4 Generic Universe Types checker

A checker for Generic Universe Types, a lightweight ownership type system, is available from <http://www.cs.washington.edu/homes/wmdietl/ownership/>.

17.5 EnerJ checker

A checker for EnerJ, an extension to Java that exposes hardware faults in a safe, principled manner to save energy with only slight sacrifices to the quality of service, is available from <http://sampa.cs.washington.edu/sampa/EnerJ>.

Chapter 18

Generics and polymorphism

This chapter describes support for Java generics and for the analogous capability over type qualifiers.

18.1 Generics (parametric polymorphism or type polymorphism)

The Checker Framework fully supports type-qualified Java generic types (also known in the research literature as “parametric polymorphism”). When instantiating a generic type, clients supply the qualifier along with the type argument, as in `List<@NonNull String>`.

Before running any checker, we recommend that you eliminate raw types from your code (e.g., your code should use `List<...>` as opposed to `List`). Using generics helps prevent type errors just as using a pluggable type-checker does, and the Checker Framework has not been extensively tested on code that uses raw types.

18.1.1 Restricting instantiation of a generic class

When you define a generic class in Java, the `extends` or `super` clause of the generic type parameter restricts how the class may be instantiated. For example, given the definition `class G<T extends Number> { ... }`, a client can instantiate it as `G<Integer>` but not `G<Date>`. Similarly, type qualifiers on the generic type parameters can restrict on how the class may be instantiated. For example, a generic list class might indicate that it can hold only non-null values. Similarly, a generic map class might indicate it requires an immutable key type, but that it supports both nullable and non-null value types.

There are two ways to restrict the type qualifiers that may be used on the actual type argument when instantiating a generic class.

The first technique is the standard Java approach of using the `extends` or `super` clause to supply an upper or lower bound. For example:

```
MyClass<T extends @NonNull Object> { ... }

MyClass<@NonNull String> m1;      // OK
MyClass<@Nullable String> m2;    // error
```

The second technique is to write a type annotation on the declaration of a generic type parameter, which specifies the exact annotation that is required on the actual type argument, rather than just a bound. For example:

```
class MyClassNN<@NonNull T> { ... }
class MyClassNble<@Nullable T> { ... }

MyClassNN<@NonNull Number> v1;    // OK
MyClassNN<@Nullable Number> v2;  // error
```

```
MyClassNble<@NonNull Number> v4; // error
MyClassNble<@Nullable Number> v3; // OK
```

18.1.2 A qualifier on a type parameter is like two bounds

A way to view a type annotation on a generic type parameter declaration is as syntactic sugar for the annotation on both the `extends` and the `super` clauses of the declaration. For example, these two declarations have the same effect:

```
class MyClassNN<@NonNull T> { ... }
class MyClassNN<T extends @NonNull Object super @NonNull void> { ... }
```

except that the latter is not legal Java syntax. The syntactic sugar is necessary because of two limitations in Java syntax: it is illegal to specify both the upper and the lower bound, and it is impossible to specify a type annotation for a lower bound without also specifying a type (use of `void` is illegal).

Suppose that a type parameter declaration is annotated with `@A`, and a bound is also given. Then the annotation `@A` applies to all bounds that have no explicit annotation. For example, the following pairs of declarations are identical.

```
class MyClass<@A T> { ... }
class MyClass<T extends @A Object super @A void> { ... }

class MyClass<@A T extends Number> { ... }
class MyClass<T extends @A Number super @A void> { ... }

class MyClass<@A T extends @B Number> { ... }
class MyClass<T extends @B Number super @A void> { ... }

class MyClass<@A T super Number> { ... }
class MyClass<T extends @A Object super @A Number> { ... }

class MyClass<@A T super @B Number> { ... }
class MyClass<T extends @A Object super @B Number> { ... }
```

18.1.3 Examples of qualifiers on a type parameter

Recall that `@Nullable X` is a supertype of `@NonNull X`, for any `X`. We can see from Section 18.1.2 that almost all of the following types mean different things:

```
class MyList1<@Nullable T> { ... }
class MyList2<@NonNull T> { ... }
class MyList3<T extends @Nullable Object> { ... }
class MyList4<T extends @NonNull Object> { ... } // same as MyList2
```

`MyList1` must be instantiated with a nullable type. The implementation must be able to consume (store) a null value and produce (retrieve) a null value.

`MyList2` and `MyList4` must be instantiated with non-null type. The implementation has to account for only non-null values — it does not have to account for consuming or producing null.

`MyList3` may be instantiated either way: with a nullable type or a non-null type. The implementation must consider that it may be instantiated either way — flexible enough to support either instantiation, yet rigorous enough to impose the correct constraints of the specific instantiation. It must also itself comply with the constraints of the potential instantiations.

One way to express the difference among `MyList1`, `MyList2`, `MyList3`, and `MyList4` is by comparing what expressions are legal in the implementation of the list — that is, what expressions may appear in the ellipsis, such as inside a method's body. Suppose each class has, in the ellipsis, these declarations:

```

T t;
@Nullable T nble;      // Section "Type annotations on a use of a generic type variable", below,
@NonNull T nn;        // further explains the meaning of "@Nullable T" and "@NonNull T".
void add(T arg) { ... }
T get(int i) { ... }

```

Then the following expressions would be legal, inside a given implementation. (Compilable source code appears as file `checker-framework/checkers/tests/nullness/GenericsExample.java`.)

	MyList1	MyList2	MyList3	MyList4
<code>t = null;</code>	OK	error	error	error
<code>t = nble;</code>	OK	error	error	error
<code>nble = null;</code>	OK	OK	OK	OK
<code>nn = null;</code>	error	error	error	error
<code>t = this.get(0);</code>	OK	OK	OK	OK
<code>nble = this.get(0);</code>	OK	OK	OK	OK
<code>nn = this.get(0);</code>	error	OK	error	OK
<code>this.add(t);</code>	OK	OK	OK	OK
<code>this.add(nble);</code>	OK	error	error	error
<code>this.add(nn);</code>	OK	OK	OK	OK

The differences are more significant when the qualifier hierarchy is more complicated than just `@Nullable` and `@NonNull`.

18.1.4 Defaults for bounds

Ordinarily, a type parameter declaration with no extends clause means the type parameter can be instantiated with any type argument at all. For example:

```

class C<T> { ... }
class C<T extends Object> { ... } // identical to previous line

```

However, instantiation may be restricted if a default qualifier is in effect (see Section 19.1.1). For example, the Nullness Checker (Chapter 3) uses a (configurable) default of `@NonNull` (see Section 3.3.2). That means that either declaration above is interpreted as

```

class C<T extends @NonNull Object> { ... }

```

and an instantiation such as `C<@Nullable Number>` is illegal. In such a case, to permit all type arguments, the programmer would write

```

class C<T extends @Nullable Object> { ... }

```

It is possible to set the default qualifier for upper bounds separately from other default qualifiers, by writing an annotation such as `@DefaultQualifier(value=Nullable.class, locations=DefaultLocation.UPPER_BOUNDS)`.

18.1.5 Type annotations on a use of a generic type variable

A type annotation on a use of a generic type variable overrides/ignores any type qualifier (in the same type hierarchy) on the corresponding actual type argument. For example, suppose that `T` is a formal type parameter. Then using `@Nullable T` within the scope of `T` applies the type qualifier `@Nullable` to the (unqualified) Java type of `T`. This feature is only rarely used.

Here is an example of applying a type annotation to a generic type variable:

```

class MyClass2<T> {
    ...
    @Nullable T myField = null;
    ...
}

```

The type annotation does not restrict how `MyClass2` may be instantiated. In other words, both `MyClass2<@NonNull String>` and `MyClass2<@Nullable String>` are legal, and in both cases `@Nullable T` means `@Nullable String`. In `MyClass2<@Interned String>`, `@Nullable T` means `@Nullable @Interned String`.

18.1.6 Covariant type parameters

Java types are *invariant* in their type parameter. This means that `A<X>` is a subtype of `B<Y>` only if `X` is identical to `Y`. For example, `ArrayList<Number>` is a subtype of `List<Number>`, but neither `ArrayList<Integer>` nor `List<Integer>` is a subtype of `List<Number>`. (If they were, there would be a type hole in the Java type system.) For the same reason, type parameter annotations are treated invariantly. For example, `List<@Nullable String>` is not a subtype of `List<String>`.

When a type parameter is used in a read-only way — that is, when values of that type are read but are never assigned — then it is safe for the type to be *covariant* in the type parameter. Use the `@Covariant` annotation to indicate this. When a type parameter is covariant, two instantiations of the class with different type arguments have the same subtyping relationship as the type arguments do.

For example, consider `Iterator`. Its elements can be read but not written, so `Iterator<@Nullable String>` can be a subtype of `Iterator<String>` without introducing a hole in the type system. Therefore, its type parameter is annotated with `@Covariant`. The first type parameter of `Map.Entry` is also covariant. Another example would be the type parameter of a hypothetical class `ImmutableList`.

The `@Covariant` annotation is trusted but not checked. If you incorrectly specify as covariant a type parameter that that can be written (say, the class performs a `set` operation or some other mutation on an object of that type), then you have created an unsoundness in the type system. For example, it would be incorrect to annotate the type parameter of `ListIterator` as covariant, because `ListIterator` supports a `set` operation.

18.2 Qualifier polymorphism

The Checker Framework also supports type *qualifier* polymorphism for methods, which permits a single method to have multiple different qualified type signatures. This is similar to Java’s generics, but is used in situations where you cannot use Java generics.

To *use* a polymorphic qualifier, just write it on a type. For example, you can write `@PolyNull` anywhere in a method that you would write `@NonNull` or `@Nullable`. A polymorphic qualifiers can be used on a method signature or body. It may not be used on a class or field.

A method written using a polymorphic qualifier conceptually has multiple versions, somewhat like a template in C++ or the generics feature of Java. In each version, each instance of the polymorphic qualifier has been replaced by the same other qualifier from the hierarchy. See the examples below in Section 18.2.1.

The method body must type-check with all signatures. A method call is type-correct if it type-checks under any one of the signatures. If a call matches multiple signatures, then the compiler uses the most specific matching signature for the purpose of type-checking. This is the same as Java’s rule for resolving overloaded methods.

To *define* a polymorphic qualifier, mark the definition with `@PolymorphicQualifier`. For example, `@PolyNull` is a polymorphic type qualifier for the Nullness type system:

```

@PolymorphicQualifier
@Target (ElementType.TYPE_USE)
public @interface PolyNull { }

```

See Section 18.2.5 for a way you can sometimes avoid defining a new polymorphic qualifier.

18.2.1 Examples of using polymorphic qualifiers

As an example of the use of `@PolyNull`, method `Class.cast` returns null if and only if its argument is null:

```
@PolyNull T cast(@PolyNull Object obj) { ... }
```

This is like writing:

```
@NonNull T cast( @NonNull Object obj) { ... }  
@Nullable T cast(@Nullable Object obj) { ... }
```

except that the latter is not legal Java, since it defines two methods with the same Java signature.

As another example, consider

```
// Returns null if either argument is null.  
@PolyNull T max(@PolyNull T x, @PolyNull T y);
```

which is like writing

```
@NonNull T max( @NonNull T x, @NonNull T y);  
@Nullable T max(@Nullable T x, @Nullable T y);
```

At a call site, the most specific applicable signature is selected.

Another way of thinking about which one of the two `max` variants is selected is that the nullness annotations of (the declared types of) both arguments are *unified* to a type that is a supertype of both, also known as the *least upper bound* or lub. If both arguments are `@NonNull`, their unification (lub) is `@NonNull`, and the method return type is `@NonNull`. But if even one of the arguments is `@Nullable`, then the unification (lub) is `@Nullable`, and so is the return type.

18.2.2 Relationship to subtyping and generics

Qualifier polymorphism has the same purpose and plays the same role as Java's generics. If a method is written using generics, it usually does not need qualifier polymorphism. If you have legacy code that is not written generically, and you cannot change it to use generics, then you can use qualifier polymorphism to achieve a similar effect, with respect to type qualifiers only. The base Java types are still treated non-generically.

Why not use ordinary subtyping to handle qualifier polymorphism? Ordinarily, when you want a method to work on multiple types, you can just use Java's subtyping. For example, the `equals` method is declared to take an `Object` as its first formal parameter, but it can be called on a `String` or a `Date` because those are subtypes of `Object`.

In most cases, the same subtyping mechanism works with type qualifiers. `String` is a supertype of `@InternedString`, so a method `toUpperCase` that is declared to take a `String` parameter can also be called on a `@InternedString` argument.

You use qualifier polymorphism in the same cases when you would use Java's generics. (If you can use Java's generics, then that is often better and you don't also need to use qualifier polymorphism.) One example is when you want a method to operate on collections with different types of elements. Another example is when you want two different formal parameters to be of the same type, without constraining them to be one specific type.

18.2.3 Using multiple polymorphic qualifiers in a method signature

Usually, it does not make sense to write only a single instance of a polymorphic qualifier in a method definition: if you write one instance of (say) `@PolyNull`, then you should use at least two. (An exception is a polymorphic qualifier on an array element type; this section ignores that case, but see below for further details.)

For example, there is no point to writing

```
void m(@PolyNull Object obj)
```

which expands to

```
void m(@NonNull Object obj)
void m(@Nullable Object obj)
```

This is no different (in terms of which calls to the method will type-check) than writing just

```
void m(@Nullable Object obj)
```

The benefit of polymorphic qualifiers comes when one is used multiple times in a method, since then each instance turns into the same type qualifier. Most frequently, the polymorphic qualifier appears on at least one formal parameter and also on the return type. It can also be useful to have polymorphic qualifiers on (only) multiple formal parameters, especially if the method side-effects one of its arguments. For example, consider

```
void moveBetweenStacks(Stack<@PolyNull Object> s1, Stack<@PolyNull Object> s2) {
    s1.push(s2.pop());
}
```

In this example, if it is acceptable to rewrite your code to use Java generics, the code can be even cleaner:

```
<T> void moveBetweenStacks(Stack<T> s1, Stack<T> s2) {
    s1.push(s2.pop());
}
```

18.2.4 Using a single polymorphic qualifier on an element type

There is an exception to the general rule that a polymorphic qualifier should be used multiple times in a signature. It can make sense to use a polymorphic qualifier just once, if it is on an array or generic element type.

For example, consider a routine that returns the index, in an array, of a given element:

```
public static int indexOf(@PolyNull Object[] a, @Nullable Object elt) { ... }
```

If `@PolyNull` were replaced with either `@Nullable` or `@NonNull`, then one of these safe client calls would be rejected:

```
@Nullable Object[] a1;
@NonNull Object[] a2;

indexOf(a1, someObject);
indexOf(a2, someObject);
```

Of course, it would be better style to use a generic method, as in either of these signatures:

```
public static <T extends @Nullable Object> int indexOf(T[] a, @Nullable Object elt) { ... }
public static <T extends @Nullable Object> int indexOf(T[] a, T elt) { ... }
```

The examples in this section use arrays, but analogous collection examples exist.

These examples show that use of a single polymorphic qualifier may be necessary in legacy code, but can often be avoided by use of better code style.

18.2.5 The `@PolyAll` qualifier applies to every type system

Ordinarily, you have to create a new polymorphic type qualifier for each type system you write. This can be tedious. More seriously, it can lead to an explosion in the number of type annotations, if some method is qualifier-polymorphic over multiple type qualifier hierarchies.

For example, a method that only performs `==` on array elements will work no matter what the array's element types are:

```

/** Searches for the first occurrence of the given element in the array,
 * testing for equality using == (not the equals method). */
public static int indexOfEq(@PolyAll Object[] a, @Nullable Object elt) {
    for (int i=0; i<a.length; i++)
        if (elt == a[i])
            return i;
    return -1;
}

```

The `@PolyAll` qualifier takes an optional argument so that you can specify multiple, independent polymorphic type qualifiers. For example, the method also works no matter what the type argument on the second argument is. This signature is overly restrictive:

```

/** Returns true if the arrays are elementwise equal,
 * testing for equality using == (not the equals method). */
public static int eltwiseEqualUsingEq(@PolyAll Object[] a, @PolyAll Object elt) {
    for (int i=0; i<a.length; i++)
        if (elt != a[i])
            return false;
    return true;
}

```

That signature requires the element type annotation to be identical for the two arguments. For example, it forbids this invocation:

```

@Mutable Object[] x;
@Immutable Object y;
... indexOf(x, y) ...

```

A better signature lets the two arrays' element types vary independently:

```

public static int eltwiseEqualUsingEq(@PolyAll(1) Object[] a, @PolyAll(2) Object elt)

```

Note that in this case, the `@Nullable` annotation on `elt`'s type is no longer necessary, since it is subsumed by `@PolyAll`.

The `@PolyAll` annotation applies to every type qualifier hierarchy for which no explicit qualifier is written. For example, a declaration like `@PolyAll @NonNull Object elt` is polymorphic over every type system *except* the nullness type system, for which the type is fixed at `@NonNull`. That would be the proper declaration for `elt` if the body had used `elt.equals(a[i])` instead of `elt == a[i]`.

Chapter 19

Advanced type system features

This chapter describes features that are automatically supported by every checker written with the Checker Framework. You may wish to skim or skip this chapter on first reading. After you have used a checker for a little while and want to be able to express more sophisticated and useful types, or to understand more about how the Checker Framework works, you can return to it.

19.1 The effective qualifier on a type (defaults and inference)

A checker sometimes treats a type as having a slightly different qualifier than what is written on the type — especially if the programmer wrote no qualifier at all. Most readers can skip this section on first reading, because you will probably find the system simply “does what you mean”, without forcing you to write too many qualifiers in your program. In particular, qualifiers in method bodies are extremely rare.

Most of this section is applicable only to source code that is being checked by a checker.

The following steps determine the effective qualifier on a type — the qualifier that the checkers treat as being present.

1. The type system adds implicit qualifiers. Implicit qualifiers can be built into a type system (Section 22.4), in which case the type system’s documentation should explain all of the type system’s implicit qualifiers. Or, a programmer may introduce an implicit annotation on each use of class *C* by writing a qualifier on the declaration of class *C*.
 - Example 1 (built-in): In the Nullness type system, `enum` values are never null, nor is a method receiver.
 - Example 2 (built-in): In the Interning type system, string literals and `enum` values are always interned.
2. If a type qualifier is present in the source code, that qualifier is used.

If the type has an implicit qualifier, then it is an error to write an explicit qualifier that is equal to (redundant with) or a supertype of (weaker than) the implicit qualifier. A programmer may strengthen (write a subtype of) an implicit qualifier, however.
3. If there is no implicit or explicit qualifier on a type, then a default qualifier may be applied; see Section 19.1.1.

At this point (after step 3), every type has a qualifier.
4. The type system may refine a qualified type on a local variable — that is, treat it as a subtype of how it was declared or defaulted. This refinement is always sound and has the effect of eliminating false positive error messages. See Section 19.1.2.

19.1.1 Default qualifier for unannotated types

A type system designer, or an end-user programmer, can cause unannotated references to be treated as if they had a default annotation.

There are several defaulting mechanisms, for convenience and flexibility. When determining the default qualifier for a use of a type, the following rules are used in order, until one applies.

- Use the innermost user-written `@DefaultQualifier`, as explained in this section.
- Use the default specified by the type system designer (Section 22.3.3).
- Use `@Unqualified`, which the framework inserts to avoid ambiguity and simplify the programming interface for type system designers. Users do not have to worry about this detail, but type system implementers can rely on the fact that some qualifier is present.

The end-user programmer specifies a default qualifier by writing the `@DefaultQualifier` annotation on a package, class, method, or variable declaration. The argument to `@DefaultQualifier` is the `String` name of an annotation. It may be a short name like `"NonNull"`, if an appropriate import statement exists. Otherwise, it should be fully-qualified, like `"checkers.nullnessquals.NonNull"`. The optional second argument indicates where the default applies. If the second argument is omitted, the specified annotation is the default in all locations. See the Javadoc of `DefaultQualifier` for details.

For example, using the Nullness type system (Chapter 3):

```
import checkers.quals.*;          // for DefaultQualifier[s]
import checkers.nullnessquals.NonNull;

@DefaultQualifier(NonNull.class)
class MyClass {

    public boolean compile(File myFile) { // myFile has type "@NonNull File"
        if (!myFile.exists())           // no warning: myFile is non-null
            return false;
        @Nullable File srcPath = ...; // must annotate to specify "@Nullable File"
        ...
        if (srcPath.exists())           // warning: srcPath might be null
            ...
    }

    @DefaultQualifier(Mutable.class)
    public boolean isJavaFile(File myFile) { // myFile has type "@Mutable File"
        ...
    }
}
```

If you wish to write multiple `@DefaultQualifier` annotations at a single location, use `@DefaultQualifiers` instead. For example:

```
@DefaultQualifiers({
    @DefaultQualifier(NonNull.class),
    @DefaultQualifier(Mutable.class)
})
```

If `@DefaultQualifier[s]` is placed on a package (via the `package-info.java` file), then it applies to the given package *and* all subpackages.

Recall that an annotation on a class definition indicates an implicit qualifier (Section 19.1) that can only be strengthened, not weakened. This can lead to unexpected results if the default qualifier applies to a class definition. Thus, you may want to put explicit qualifiers on class declarations (which prevents the default from taking effect), or exclude class declarations from defaulting.

When a programmer omits an `extends` clause at a declaration of a type parameter, the default still applies to the implicit upper bound. For example, consider these two declarations:

```
class C<T> { ... }
class C<T extends Object> { ... } // identical to previous line
```

The two declarations are treated identically by Java, and the default qualifier applies to the `Object` upper bound whether it is implicit or explicit. (The `@NonNull` default annotation applies only to the upper bound in the `extends` clause, not to the lower bound in the inexpressible implicit `super void` clause.)

19.1.2 Automatic type refinement (flow-sensitive type qualifier inference)

In order to reduce your burden of annotating types in your program, the checkers soundly treat certain variables and expressions as having a subtype of their declared or defaulted (Section 19.1.1) type. This functionality never introduces unsoundness nor causes an error to be missed.

By default, all checkers, including new checkers that you write, can take advantage of this functionality. Most of the time, users don't have to think about, and may not even notice, this feature of the framework. The checkers simply do the right thing even when a programmer omits an annotation on a local variable, or when a programmer writes an unnecessarily general type in a declaration.

The functionality has a variety of names: automatic type refinement, flow-sensitive type qualifier inference, local type inference, and sometimes just “flow”.

If you are curious or want more details about this feature, then read on.

As an example, the Nullness checker (Chapter 3) can automatically determine that certain variables are non-null, even if they were explicitly or by default annotated as nullable. The checker treats a variable or expression as `@NonNull`

- starting at the time that it is either assigned a non-null value or checked against null (e.g., via an assertion, `if` statement, or being dereferenced)
- until it might be re-assigned (e.g., via an assignment that might affect this variable, or via a method call that might affect this variable).

As with explicit annotations, the implicitly non-null types permit dereferences and assignments to non-null types, without compiler warnings.

Consider this code, along with comments indicating whether the Nullness checker (Chapter 3) issues a warning. Note that the same expression may yield a warning or not depending on its context.

```
// Requires an argument of type @NonNull String
void parse(@NonNull String toParse) { ... }

// Argument does NOT have a @NonNull type
void lex(@Nullable String toLex) {
    parse(toLex); // warning: toLex might be null
    if (toLex != null) {
        parse(toLex); // no warning: toLex is known to be non-null
    }
    parse(toLex); // warning: toLex might be null
    toLex = new String(...);
    parse(toLex); // no warning: toLex is known to be non-null
}
```

If you find examples where you think a value should be inferred to have (or not have) a given annotation, but the checker does not do so, please submit a bug report (see Section 25.2) that includes a small piece of Java code that reproduces the problem.

The inference indicates when a variable can be treated as having a subtype of its declared type — for instance, when an otherwise nullable type can be treated as a `@NonNull` one. The inference never treats a variable as a supertype of its declared type (e.g., an expression of `@NonNull` type is never inferred to be treated as possibly-null).

Type inference is never performed for method parameters of non-private methods, nor for non-private fields. More generally, the inferred information is never written to the `.class` file as user-written annotations are. If the checker did inference in externally-visible locations and wrote it to the `.class` file, then the resulting `.class` file would be different depending on whether an annotation processor had been run or not. It is a design goal that the same annotations appear in the `.class` file regardless of whether the class is compiled with or without the checker, and this requires that any public signature be fully annotated by the user rather than inferred.

19.1.3 Fields and flow sensitivity analysis

Flow sensitivity analysis infers the type of fields in some restricted cases:

- A final initialized field: Type inference is performed for final fields that are initialized to a compile-time constant at the declaration site; so the type of `protocol` is `@NonNull String` in the following declaration:

```
public final String protocol = "https";
```

Please note that such inferred type may leak to the public interface of the class. To override such behavior, you can explicitly insert the desired annotation, e.g.,

```
public final @Nullable String protocol = "https";
```

- Within method bodies: Type inference is performed for fields in the context of method bodies, like local variables, but method invocations invalidate any inferred information. Consider the following example, where `updatedAt` is a nullable field:

```
class DBObject {
    @Nullable Date updatedAt;

    void persistData() {
        ... // write to disk or other non-volatile memory
        updatedAt = null;
    }

    void update() {
        if (updatedAt == null)
            updatedAt = new Date();
        // updatedAt is nonnull
        log("Updating object at " + updatedAt.getTime());

        persistData();
        // updatedAt is nullable again
        log.debug("Saved object updated at " + updatedAt.getTime()); // invalid!
    }
}
```

Here the call to `persistData()` invalidates the inferred non-null type of `updatedAt`.

When methods do not modify any object state or have any identity side-effects (e.g., `log()` method here), you can annotate these methods as `Pure`. When a method is annotated as `Pure`, the flow analyzer carries the inferred types across the method invocation boundary.

19.1.4 Inherited defaults

In certain situations, it would be convenient for an annotation on a superclass member to be automatically inherited by subclasses that override it. This feature would reduce both annotation effort and program comprehensibility. In general, a program is read more often than it is edited/annotated, so the Checker Framework does not currently support this feature. Here are more detailed justifications:

- Currently, a user can determine the annotation on a parameter or return value by looking at a single file. If annotations could be inherited from supertypes, then a user would have to examine all supertypes to understand the meaning of an unannotated type in a given file.
- Different annotations might be inherited from a supertype and an interface, or from two interfaces. Presumably, the subtype's annotations would be stronger than either (the greatest lower bound in the type system), or an error would be thrown if no such annotations existed.

If these issues can be resolved, then the feature may be added in the future. Or, it may be added optionally, and each type-checker implementation can enable it if desired.

19.2 Writing Java expressions as annotation arguments

Sometimes, it is necessary to write a Java expression as the argument to an annotation. As of this writing, the annotations that take a Java expression as an argument are:

- `@KeyFor`
- `@NonNullOnEntry`
- `@AssertNonNullAfter`
- `@AssertNonNullIfTrue`
- `@AssertNonNullIfFalse`

The expression is a subset of legal Java expressions:

- the receiver object, `this`.
- a formal parameter. Write # followed by the one-based parameter index. For example: #1, #3. It is not permitted to write #0; use `this` instead.
- a local variable. This is not applicable for method annotations, but is applicable to type annotations such as `@KeyFor`. Write the variable name. For example: `myLocalVar`.
- a static variable. Write the class name and the variable, as in `System.out`. Within the class itself, just write the field name — do not write the class name.
- a field of any expression. For example: `next`, `this.next`, `#0.next`, `myLocalVar.next`.
- a method invocation on any expression. The method must be pure and have no formal parameters. For example: `myClass.getPackage()`, `myClass.getSuperclass()`, `myClass.getComponentType()`.
Warning: currently, annotations that use method calls are *not* checked. The annotation is trusted, and other code will rely on it, but it is not verified that other code establishes or maintains the validity of the annotation. Such expressions are still useful if a human verifies their correctness. They are used in the JDK annotations, for example.

You may optionally omit a leading “`this.`”, just as in Java. Thus, `this.next` and `next` are equivalent, assuming that there is no shadowing definition of `next`.

Limitation: Currently, only one level of field access can be checked; the checker cannot handle, for example, `"this.field1.field2"`.

(A side note: The formal parameter syntax #0 may seem less convenient than writing the formal parameter name. This syntax is necessary because in the `.class` file, no parameter name information is available. Running the compiler without a checker should create legal annotations in the `.class` file, so we cannot rely on the checker to translate names to indices.)

19.3 Unused fields and dependent types

In an inheritance hierarchy, subclasses often introduce new methods and fields. For example, a `Marsupial` (and its subclasses such as `Kangaroo`) might have a variable indicating the size of the animal's pouch. Because such variables would not exist in superclasses such as `Mammal` and `Animal`, any attempt to use them would be a compile-time error.

If you cannot use subtypes but wish to enforce similar requirements using type qualifiers, you can do so. To restrict which methods may be invoked, you can write an annotation on a method receiver; then the method may only be invoked on an expression whose type has the given annotation (or one of its subtypes). Section 19.3.1 describes how to restrict which fields may be accessed: in other words, a given field may only be accessed from an expression whose type has a given qualifier. Then, Section 19.3.2 describes an even more powerful mechanism, by which the qualifier of a field depends on the qualifier of the expression from which the field was accessed. (Also see the discussion of typestate checkers, in Chapter 16.)

19.3.1 Unused fields

A Java subtype can have more fields than its supertype. For example:

```
class Mammal extends Animal { ... }
class Marsupial extends Mammal {
    ...
    int pouchSize; // pouch capacity, in cubic centimeters
    ...
}
```

You can simulate the same effect for type qualifiers: a given field may not be accessed via a reference with a supertype qualifier, but can be accessed via a reference with a subtype qualifier. For example:

```
@interface Mammal { }
@interface Marsupial { }
class Animal {
    @Unused(when=Mammal.class)
    int pouchSize; // pouch capacity, in cubic centimeters
    ...
}

@Marsupial Animal joey = ...;
... joey.pouchSize ... // OK
@Mammal Animal mae = ...;
... mae.pouchSize ... // compile-time error
```

The `@Unused` annotation on a field declares that the field may *not* be accessed via a receiver of the given qualified type (or any *supertype*).

(It would probably be clearer to replace `@Unused` by an annotation that indicates when the field *may* be used.)

19.3.2 Dependent types

A variable has a *dependent type* if its type depends on some other value or type.

The Checker Framework supports a form of dependent types, via the `@Dependent` annotation. This annotation changes the type of a field or variable, based on the qualified type of the receiver (`this`). This can be viewed as a more expressive form of polymorphism (see Section 18). It can also be seen as a way of linking the meanings of two type qualifier hierarchies.

Here is a restatement of the example of Section 19.3.1, using `@Dependent`:

```
@interface Mammal { }
@interface Marsupial { }
class Animal { ...
    // pouch capacity, in cubic centimeters
    // (non-null if this animal is a marsupial)
```

```
    @Nullable @Dependent(result=NonNull.class, when=Marsupial.class) Integer pouchSize;
    ...
}

@Marsupial Animal joey = ...;
... joey.pouchSize.intValue() ... // OK
@mammal Animal mae = ...;
... mae.pouchSize.intValue() ... // compile-time error:
// dereference of possibly-null mae.pouchSize
```

However, when the `@Unused` annotation is sufficient, you should use it instead of `@Dependent`.

Chapter 20

Handling warnings and legacy code

Section 2.4.1 describes a methodology for applying annotations to legacy code. This chapter tells you what to do if, for some reason, you cannot change your code in such a way as to eliminate a checker warning.

Also recall that you can convert checker errors into warnings via the `-Awarns` command-line option; see Section 2.2.1.

20.1 Checking partially-annotated programs: handling unannotated code

Sometimes, you wish to type-check only part of your program. You might focus on the most mission-critical or error-prone part of your code. When you start to use a checker, you may not wish to annotate your entire program right away. You may not have enough knowledge to annotate poorly-documented libraries that your program uses.

If annotated code uses unannotated code, then the checker may issue warnings. For example, the Nullness checker (Chapter 3) will warn whenever an unannotated method result is used in a non-null context:

```
@NonNull myvar = unannotated_method(); // WARNING: unannotated_method may return null
```

If the call *can* return null, you should fix the bug in your program by removing the `@NonNull` annotation in your own program.

If the library call *never* returns null, there are several ways to eliminate the compiler warnings.

1. Annotate `unannotated_method` in full. This approach provides the strongest guarantees, but may require you to annotate additional methods that `unannotated_method` calls. See Chapter 21 for a discussion of how to annotate libraries for which you have no source code.
2. Annotate only the signature of `unannotated_method`, and suppress warnings in its body. Two ways to suppress the warnings are via a `@SuppressWarnings` annotation or by not running the checker on that file (see Section 20.2).
3. Suppress all warnings related to uses of `unannotated_method` via the `skipUses` processor option (see Section 20.2). Since this can suppress more warnings than you may expect, it is usually better to annotate at least the method's signature. If you choose the boundary between the annotated and unannotated code wisely, then you only have to annotate the signatures of a limited number of classes/methods (e.g., the public interface to a library or package).

Chapter 21 discusses adding annotations to signatures when you do not have source code available. Section 20.2 discusses suppressing warnings.

If you annotate a third-party library, please share it with us so that we can distribute the annotations with the Checker Framework; see Section 25.2.

20.2 Suppressing warnings

You may wish to suppress checker warnings because of unannotated libraries or un-annotated portions of your own code, because of application invariants that are beyond the capabilities of the type system, because of checker limitations, because you are interested in only some of the guarantees provided by a checker, or for other reasons. You can suppress warnings via

1. the `@SuppressWarnings` annotation,
2. the `-AskipUses` command-line option,
3. the `-AskipDefs` command-line option,
4. the `-Alint` command-line option,
5. not using the `-processor` command-line option, or
6. checker-specific mechanisms.

We now explain these mechanisms in turn.

20.2.1 `@SuppressWarnings` annotation

You can suppress specific errors and warnings by use of the `@SuppressWarnings("checkername")` annotation, for example `@SuppressWarnings("interning")` or `@SuppressWarnings("nullness")`. The argument *checkername* is in lower case and is derived from the way you invoke the checker; for example, if you invoke a checker as `javac -processor MyNiftyChecker . . .`, then you would suppress its error messages with `@SuppressWarnings("mynifty")`. (An exception is the Basic Checker, for which you use the annotation name; see Section 15.1).

A `@SuppressWarnings` annotation may be placed on program elements such as a local variable declaration, a method, or a class. It suppresses all warnings related to the given checker, for that program element.

For instance, one common use is to suppress warnings at a cast that you know is safe. Here is an example that uses the Tainting Checker (Section 9):

```
@SuppressWarnings("tainting")
String myvar = (@Untainted String) expr; // expr has type: @Tainted String
```

It is good practice to suppress warnings in the smallest possible scope. For example, if a particular expression causes a false positive warning, you should extract that expression into a local variable and place a `@SuppressWarnings` annotation on the variable declaration. As another example, if you have annotated the signatures but not the bodies of the methods in a class or package, put a `@SuppressWarnings` annotation on the class declaration or on the package's `package-info.java` file.

Another good practice is to use the most specific possible argument to `@SuppressWarnings`. The string can be of the form *checkername* or or *checkername:messagekey*. The *checkername* part is as described above. The *messagekey* part suppresses only errors/warnings relating to the given message key. For example, `cast.unsafe` is the key for warnings about an unsafe cast, and `cast.redundant` to the key for warnings about a redundant cast.

Thus, the above example could have been written as:

```
@SuppressWarnings("tainting")           // suppresses all tainting-related warnings
@SuppressWarnings("tainting:cast.unsafe") // suppresses tainting warnings about unsafe casts
@SuppressWarnings("tainting:cast")       // suppresses tainting warnings about casts
```

For a list of the message keys, see the `messages.properties` files in `checker-framework/checkers/src/checkers/checkername/messages.properties`. Each checker is built on the basetype checker and inherits its properties. Thus, to find all the error keys for a checker, you usually need to examine its own `messages.properties` file and that of basetype.

If a checker produces a warning/error and you want to determine its message key, you can re-run the checker, passing the the `-Anomsgtext` command-line option (Section 22.8).

20.2.2 `-AskipUses` command-line option

You can suppress all errors and warnings at all *uses* of a given class (but the class itself is still type-checked, unless you also use the `-AskipDefs` command-line option, see 20.2.3). Set the `-AskipUses` command-line option to a regular expression that matches class names (not file names) for which warnings and errors should be suppressed. For example, suppose that you use “`-AskipUses=^java\.`” on the command line (with appropriate quoting) when invoking `javac`. Then the checkers will suppress all warnings related to classes whose fully-qualified name starts with `java.`, such as all warnings relating to invalid arguments and all warnings relating to incorrect use of the return value.

To suppress all errors and warnings related to multiple classes, you can use the regular expression alternative operator “`|`”, as in “`-AskipUses="java\.lang\.|java\.util\."`” to suppress all warnings related to uses of classes belong to the `java.lang` or `java.util` packages.

20.2.3 `-AskipDefs` command-line option

You can suppress all errors and warnings in the *definition* of a given class (but uses of the class are still type-checked, unless you also use the `-AskipUses` command-line option, see 20.2.2). Set the `-AskipDefs` command-line option to a regular expression that matches class names (not file names) in whose definition warnings and errors should be suppressed. For example, if you use “`-AskipDefs=^mypackage\.`” on the command line (with appropriate quoting) when invoking `javac`, then the definitions of classes whose fully-qualified name starts with `mypackage.` will not be checked.

Another way not to type-check a file is not to pass it on the compiler command-line: the Checker Framework type-checks only files that are passed to the compiler on the command line, and does not type-check any file that is not passed to the compiler. The `-AskipDefs` command-line option is intended for situations in which the build system is hard to understand or change. In such a situation, a programmer may find it easier to supply an extra command-line argument, than to change the set of files that is compiled.

20.2.4 `-Alint` command-line option

The `-Alint` option enables or disables optional checks, analogously to `javac`’s `-Xlint` option. Each of the distributed checkers supports at least the following lint options:

- `cast:unsafe` (default: on) warn about unsafe casts that are not checked at run time, as in `((@NonNull String) myref)`. Such casts are generally not necessary when flow-sensitive local type refinement is enabled.
- `cast:redundant` (default: on) warn about redundant casts that are guaranteed to succeed at run time, as in `((@NonNull String) "m")`. Such casts are not necessary, because the target expression of the cast already has the given type qualifier.
- `cast` Enable or disable all cast-related warnings.
- `all` Enable or disable all lint warnings, including checker-specific ones if any. Examples include `nulltest` for the Nullness Checker (see Section 1) and `dotequals` for the Interning Checker (see Section 4.3). This option does not enable/disable the checker’s standard checks, just its optional ones.
- `none` The inverse of `all`: disable or enable all lint warnings, including checker-specific ones if any.

To activate a lint option, write `-Alint=` followed by a comma-delimited list of check names. If the option is preceded by a hyphen (`-`), the warning is disabled. For example, to disable all lint options except redundant casts, you can pass `-Alint=-all,cast:redundant` on the command line.

20.2.5 No `-processor` command-line option

You can also compile parts of your code without use of the `-processor` switch to `javac`. No checking is done during such compilations.

20.2.6 Checker-specific mechanisms

Finally, some checkers have special rules. For example, the Nullness checker (Chapter 3) uses `assert` statements that contain null checks, and the special `castNonNull` method, to suppress warnings (Section 3.4.1). This manual also explains special mechanisms for suppressing warnings issued by the Fenum checker (Section 8.4) and the Units checker (Section 14.5).

20.3 Backward compatibility with earlier versions of Java

Sometimes, your code needs to be compiled by people who are not using a compiler that supports type annotations. Sections 20.3.1–20.3.3 discuss this situation, which you can handle by writing annotations in comments.

In other cases, your code needs to be run by people who are not using a Java 8 JVM. Section 20.3.4 discusses this situation, which you can handle by passing the `-target 5` command-line argument.

(**Note:** These are features of the Type Annotations compiler that is distributed along with the Checker Framework. They are *not* supported by the mainline OpenJDK compiler. These features are the key difference between the Type Annotations compiler and the OpenJDK compiler on which it is built.)

20.3.1 Annotations in comments

A Java 4 compiler does not permit use of annotations, and a Java 5/6/7 compiler only permits annotations on declarations (but not on generic arguments, casts, `extends` clauses, method receiver, etc.).

So that your code can be compiled by any Java compiler (for any version of the Java language), you may write any single annotation inside a `/*...*/` Java comment, as in `List</*@NonNull*/ String>`. The Type Annotations compiler treats the code exactly as if you had not written the `/*` and `*/`. In other words, the Type Annotations compiler will recognize the annotation, but your code will still compile with any other Java compiler.

This feature only works if you provide no `-source` command-line argument to `javac`, or if the `-source` argument is `1.8` or `8`.

In a single program, you may write some annotations in comments, and others outside of comments.

By default, the compiler ignores any comment that contains spaces at the beginning or end, or between the `@` and the annotation name. In other words, it reads `/*@NonNull*/` as an annotation but ignores `/* @NonNull*/` or `/*@ NonNull*/` or `/*@NonNull */`. This feature enables backward compatibility with code that contains comments that start with `@` but are not annotations. (The ESC/Java [FLL⁺02], JML [LBR06], and Splint [Eva96] tools all use `“/*@”` or `“/* @”` as a comment marker.) Compiler flag `-XDTA:spacesincomments` causes the compiler to parse annotation comments even when they contain spaces. You may need to use `-XDTA:spacesincomments` if you use Eclipse’s “Source > Correct Indentation” command, since it inserts space in comments. But the annotation comments are less readable with spaces, so you may wish to disable inserting spaces: in the Formatter preferences, in the Comments tab, unselect the “enable block comment formatting” checkbox.

There is no way to turn off the annotations in comments feature. If you don’t want this feature, you can use a standard Java 8 compiler that supports type annotations but not annotations in comments. If your code already contains comments of the form `/*@...*/` that look like type annotations, and you want the Type Annotations compiler not to try to interpret them, then you can add spaces to the comments.

There is a more powerful mechanism that permits arbitrary code to be written in a comment that is formatted as `“/*>>>...*/”`, with the first three characters of the comment being greater-than signs. As with annotations in comments, the commented code is ignored by ordinary compilers but is treated like ordinary code by the UW Type Annotations compiler. This mechanism is intended only to support writing import statements and the new receiver (“this”) syntax, as in

```
public boolean getResult(/*>>> @ReadOnly MyClass this*/) { ... }
public boolean getResult(/*>>> @ReadOnly MyClass this, */ String argument) { ... }
```

for a method that does not modify its receiver. It can also be used for adding multiple annotations inside one comment.

It would be possible to abuse this mechanism to inject code only when using the Type Annotations compiler. Doing so is not a sanctioned use of the mechanism.

20.3.2 Implicit import statements

When writing source code with annotations, it is more convenient to write a short form such as `@NonNull` instead of `@checkers.nullness.quals.NonNull`. Here are ways to achieve this goal.

- The traditional way to do this is to write an import statement like “`import checkers.nullness.quals.*;`”. This works, but everyone who compiles the code (no matter what compiler they use, and even if the annotations are in comments) must have the annotation definitions (e.g., the `checkers.jar` or `checkers-quals.jar` file) on their classpath. The reason is that a Java compiler issues an error if an imported package is not on the classpath. See Section 2.1.1.
- Write an import statement in a comment, just as for annotations in comments:

```
/*>>> import checkers.nullness.quals.*; */
```

- An alternative is to set the shell environment variable `jsr308_imports` when you compile the code. The Type Annotations compiler treats this as if the given packages/classes were imported, but other compilers ignore the `jsr308_imports` environment variable — they do not need it, since they do not support annotations in comments. Thus, your code can compile whether or not the Type Annotations compiler is being used. You can specify multiple packages/classes separated by the classpath separator (same as the file path separator: `;` for Windows, and `:` for Unix and Mac). For example, to implicitly import the Nullness and Interning qualifiers, set `jsr308_imports` to `checkers.nullness.quals.*:checkers.interning.quals.*`.

Three ways to set an environment variable are:

- Set the environment variable in your shell. For example, in `bash`, you could do `export jsr308_imports='checkers.nullness.quals.*:checkers.interning.quals.*'`. This takes effect for all subsequent commands in that shell. To take effect for all shells that you run, put the command in your `~/.bashrc` file.
- Set the environment variable for a single command. For example, in `bash` prefix the `javac` command by `jsr308_imports='checkers.nullness.quals.*'`.
- Set the environment variable for a single command, via a `javac` argument. Use the `javac` command-line argument `-J-Djsr308_imports='checkers.nullness.quals.*'`.

If you issue the `javac` command from the command line or in a Makefile, you may need to add quotes (as shown above), to prevent your shell from expanding the `*` character. If you supply the `-J-Djsr308_imports` argument via an Ant buildfile, you do not need the extra quoting.

- If it is not possible to set the environment variable (for example, Maven bug PLXCOMP-190 makes it impossible to use the `-J-Djsr308_imports` command-line argument), you can instead use a different `javac` command-line argument: `-jsr308_imports ...` or `-Djsr308_imports=...` (they have the same effect). The same syntax for the packages/classes, and the same warnings about quoting from the command line, apply as for the `jsr308_imports` environment variable.

20.3.3 Migrating away from annotations in comments

Suppose that your codebase currently uses annotations in comments, but you wish to remove the comment characters around your annotations, because in the future you will use only compilers that support type annotations. This Unix command removes the comment characters, for all Java files in the current working directory or any subdirectory.

// TODO: This doesn't handle the `///` comments. I should adapt it to do so.

```
find . -type f -name '*.java' -print \  
  | xargs grep -l -P '/\*\s*@[^\s*/]+\s*\*/' \  
  | xargs perl -pi.bak -e 's|/\*\s*@[^\s*/]+\s*\*/|@|1|g'
```

You can customize this command:

- To process comments with embedded spaces and asterisks, change two instances of “`[^\s*/]`” to “`[^/]`”.
- To ignore comments with leading or trailing spaces, remove the four instances of “`\s*`”.
- To not make backups, remove “`.bak`”.

If your code used implicit import statements (Section 20.3.2), then after uncommenting the annotations, you may also need to introduce explicit import statements into your code.

20.3.4 Annotations in Java 5 `.class` files

If you supply the `-target 5` command-line argument along with no `-source` argument (or `-source 8`, which is equivalent), then the Type Annotations compiler creates a `.class` file that can be run on a Java 5 JVM, but that contains the type annotations. (It does not matter whether the type annotations were written in comments or not.) The fact that the `.class` file contains the type annotations is useful when type-checking client code. If you try to type-check client code against a library that lacks type annotations, then spurious warnings can result. So, use of `-target 5` gives backward compatibility with earlier JVMs while still permitting pluggable type-checking.

Ordinary Java compilers do not let you use a `-target` command-line argument with a value less than the `-source` argument.

Use of the `-source 5` command-line argument produces a `.class` file that does not contain type annotations. One reason you might want to periodically compile with the `-source 5` argument is to ensure that your code does not use any Java 8 features other than type annotations in comments.

Chapter 21

Annotating libraries

When annotated code uses an unannotated library, a checker may issue warnings. As described in Section 20.1, the best way to correct this problem is to add annotations to the library. (Alternately, you can instead suppress all warnings related to an unannotated library by use of the `-SkipUses` command-line option; see Section 20.2.) If you have source code for the library, you can easily add the annotations. This chapter tells you how to add annotations to a library for which you have no source code, because the library is distributed only in binary form (as `.class` files, possibly packaged in a `.jar` file). This chapter is also useful if you do not wish to edit the library's source code.

You can make the annotations known to the checkers in two ways.

- You can write annotations in a “stub file” containing classes with no method bodies. Section 21.2 describes how to create and use stub files.
- You can insert annotations in the compiled `.class` files of the library. You would express the annotations textually, typically as an annotation index file, and then insert them in the library by using the Annotation File Utilities (<http://types.cs.washington.edu/annotation-file-utilities/>). See the Annotation File Utilities documentation for full details.

The Checker Framework distribution contains annotations for popular libraries, such as the JDK. It uses both of the above mechanisms. The Nullness, Javari, IGJ, and Interning Checkers use an annotated JDK (Section 21.3), and all other checkers use stub files (Section 21.2).

If you annotate additional libraries, please share them with us so that we can distribute the annotations with the Checker Framework; see Section 25.2. You can determine the correct annotations for a library either automatically by running an inference tool, or manually by reading the documentation. Presently, type inference tools are available for the Nullness (Section 3.3.5) and Javari (Section 6.2.2) type systems.

21.1 Choosing between stub files and annotated `.class` files

A checker can read annotations either from a stub file or from a library's `.class` files. This section helps you choose between the two alternatives.

Once created, a stub file can be used directly; this makes it an easy way to get started with library annotations. When provided by the author of the checker, a stub file is used automatically, with no need for the user to supply a command-line option.

Inserting annotations in a library's `.class` files takes an extra step using an external tool, the Annotation File Utilities (<http://types.cs.washington.edu/annotation-file-utilities/>). However, this process does not suffer the limitations of stub files (Section 21.2.4).

21.2 Using stub classes

A stub file contains “stub classes” that contain annotated signatures, but no method bodies. A checker uses the annotated signatures at compile time, instead of or in addition to annotations that appear in the library.

Section 21.2.1 describes how to create stub classes. Section 21.2.2 describes how to use stub classes. These sections illustrate stub classes via the example of creating a `@Interned`-annotated version of `java.lang.String`. You don’t need to repeat these steps to handle `java.lang.String` for the Interning Checker, but you might do something similar for a different class and/or checker.

21.2.1 Creating a stub file

If you have access to the Java source code

Every Java file is a stub file. If you have access to the Java file, then you can use the Java file as the stub file, without removing any of the parts that the stub file format permits you to. Just add annotations to the signatures, leaving the method bodies unchanged. Optionally (but highly recommended!), run the type-checker to verify that your annotations are correct. When you run the type-checker on your annotations, there should not be any stub file that also contains annotations for the class. In particular, if you are type-checking the JDK itself, then you should use the `-Aignorejdkastub` command-line option.

This approach retains the original documentation and source code, making it easier for a programmer to double-check the annotations. It also enables creation of diffs, easing the process of upgrading when a library adds new methods. And, the annotations are in a format that the library maintainers can even incorporate.

The downside of this approach is that the stub files are larger. This can slow down parsing. Furthermore, a programmer must search the stub file for a given method rather than just skimming a few pages of method signatures.

If you do not have access to the Java source code

If you do not have access to the library source code, then you can create a stub file from the class file (Section 21.2.1), and then annotate it. The rest of this section describes this approach.

1. Create a stub file by running the stub class generator. (`checkers.jar` must be on your classpath.)

```
cd nullness-stub
java checkers.util.stub.StubGenerator java.lang.String > String.astub
```

Supply it with the fully-qualified name of the class for which you wish to generate a stub class. The stub class generator prints the stub class to standard out, so you may wish to redirect its output to a file.

2. Add import statements for the annotations. So you would need to add the following import statement at the beginning of the file:

```
import checkers.interningquals.*;
```

The stub file parser silently ignores any annotations that it cannot resolve to a type, so don’t forget the import statement.

3. Add annotations to the stub class. For example, you might annotate the `String.intern()` method as follows:

```
@Interned String intern();
```

You may also remove irrelevant parts of the stub file; see Section 21.2.3.

21.2.2 Using a stub file

The `-Astubs` argument causes the Checker Framework to read annotations from annotated stub classes in preference to the unannotated original library classes. For example:

```
javac -processor checkers.interning.InterningChecker -Astubs=String.astub:stubs MyFile.java MyOtherFile.java ...
```

Each stub path entry is a file or a directory; specifying a directory is equivalent to specifying every file in it whose name ends with `.astub`. The stub path entries are delimited by `File.pathSeparator` (':' for Linux and Mac, ';' for Windows).

A checker automatically reads the stub file `jdk.astub`. (The checker author should place it in the same directory as the Checker class, i.e., the subclass of `BaseTypeVisitor`.) Programmers should only use the `-Astubs` argument for additional stub files they create themselves.

21.2.3 Stub file format

Every Java file is a valid stub file. However, you can omit information that is not relevant to pluggable type-checking; this makes the stub file smaller and easier for people to read and write.

As an illustration, a stub file for the Interning type system (Chapter 4) could be:

```
import checkers.interning.qual.Interned;
package java.lang;
@Interned class Class<T> { }
class String {
    @Interned String intern();
}
```

The stub file format is allowed to differ from Java source code in the following ways:

Method bodies: The stub class does not require method bodies for classes; any method body may be replaced by a semicolon (;), as in an interface or abstract method declaration.

Method declarations: You only have to specify the methods that you need to annotate. Any method declaration may be omitted, in which case the checker reads its annotations from library's `.class` files. (If you are using a stub class, then typically the library is unannotated.)

Declaration specifiers: Declaration specifiers (e.g., `public`, `final`, `volatile`) may be omitted.

Import statements: The only required import statements are the ones to import type annotations. Such imports must be at the beginning of the file. Other import statements are optional.

Multiple classes and packages: The stub file format permits having multiple classes and packages. The packages are separated by a package statement: `package my.package;`. Each package declaration may occur only once; in other words, all classes from a package must appear together.

21.2.4 Limitations

The stub file reader has several limitations:

- It does not handle `enums`.
- It does not handle nested classes. To work around this, it permits a top-level class to be written with a `$` in its name, and applies the annotations to the appropriate nested class.

If these limitations are a problem, then you should insert annotations in the library's `.class` files instead.

21.3 Using distributed annotated JDKs

The Checker Framework distribution contains annotated JDKs at the path `checkers/jdk/jdk.jar`. The `javac` that is distributed with the Checker Framework uses the annotated JDKs by default.

If you use a different `javac`, then you must add a `-Xbootclasspath/p:` argument, which causes the compiler to read annotations from annotated JDK classes in preference to the unannotated original library classes. Supply `-Xbootclasspath/p:` in addition to whatever other arguments you usually use, including `-classpath`. For example:

```
javac -processor checkers.nullness.NullnessChecker -Xbootclasspath/p:${CHECKERS}/jdk/jdk.jar my_source_files
```

If you do not supply the `-Xbootclasspath/p:` option, the checker will print a message warning you to do so. In the unlikely event that you want to suppress this warning, use `-Anocheckjdk`.

The annotated JDK should *not* be in your classpath at run time, only at compile time.

The supplied annotated JDK is a version of JDK 6. If you wish to have an annotated version of JDK 7, you will need to create it yourself. Running `ant jdk.jar` from the `checkers/` directory will perform this process.

21.4 Troubleshooting/debugging annotated libraries

Sometimes, it may seem that a checker is treating a library as unannotated even though the library has annotations. The compiler has two flags that may help you in determining whether library files are read, and if they are read whether the library's annotations are parsed.

- `-verbose` Outputs info about compile phases — when the compiler reads/parses/attributes/writes any file. Also outputs the classpath and sourcepath paths.
- `-XDTA:parser` (which is equivalent to `-XDTA:reader` plus `-XDTA:writer`) Sets the internal `debugJSR308` flag, which output information about reading and writing.

Chapter 22

How to create a new checker

This chapter describes how to create a checker — a type-checking compiler plugin that detects bugs or verifies their absence. After a programmer annotates a program, the checker plugin verifies that the code is consistent with the annotations. If you only want to *use* a checker, you do not need to read this chapter.

Writing a simple checker is easy! For example, here is a complete, useful type checker:

```
@TypeQualifier
@SubtypeOf(Unqualified.class)
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface Encrypted {}
```

This checker is so short because it builds on the Basic Checker (Chapter 15). See Section 15.2 for more details about this particular checker. When you wish to create a new checker, it is often easiest to begin by building it declaratively on top of the Basic Checker, and then return to this chapter when you need more expressiveness or power than the Basic Checker affords.

You can also create your own checker by customizing a different existing checker. Specific checkers that are designed for extension (besides the Basic Checker) include the Fake Enumeration Checker (Chapter 8, page 45), the Units Checker (Chapter 14, page 61), and the tystate checkers (Chapter 16, page 67). Or, you can copy and then modify a different existing checker — whether one distributed with the Checker Framework or a third-party one.

You can place your checker's source files wherever you like. When you compile your checker, `$CHECKERS/binary/jsr308-all.jar` should be on your classpath. (If you wish to modify an existing checker in place, or to place the source code for your new checker in your own private copy of the Checker Framework source code, then you need to be able to re-compile the Checker Framework, as described in Section 25.3.)

The rest of this chapter contains many details for people who want to write more powerful checkers. You do not need all of the details, at least at first. In addition to reading this chapter of the manual, you may find it helpful to examine the implementations of the checkers that are distributed with the Checker Framework. You can even create your checker by modifying one of those. The Javadoc documentation of the framework and the checkers is in the distribution and is also available online at <http://types.cs.washington.edu/checker-framework/current/doc/>.

If you write a new checker and wish to advertise it to the world, let us know so we can mention it in the Checker Framework manual, link to it from the webpages, or include it in the Checker Framework distribution. For examples, see Chapters 16 and 17.

22.1 Relationship of the Checker Framework to other tools

This table shows the relationship among various tools. All of the tools use the Type Annotations (JSR 308) syntax. You use the Checker Framework to build pluggable type systems, and the Annotation File Utilities to manipulate `.java` and `.class` files.

Basic Checker	Nullness Checker	Mutation Checker	Tainting Checker	...	Your Checker		
Base Checker (enforces subtyping rules)						Type inference	Other tools
Checker Framework (enables creation of pluggable type-checkers)						Annotation File Utilities (.java ↔ .class files)	
Type Annotations syntax and classfile format (“JSR 308”) (no built-in semantics)							

The Base Checker enforces the standard subtyping rules on extended types. The Basic Checker is a simple use of the Base Checker that supports providing type qualifiers on the command line. You usually want to build your checker on the Base Checker.

22.2 The parts of a checker

The Checker Framework provides abstract base classes (default implementations), and a specific checker overrides as little or as much of the default implementations as necessary. Sections 22.3–22.6 describe the components of a type system as written using the Checker Framework:

22.3 Type qualifiers and hierarchy. You define the annotations for the type system and the subtyping relationships among qualified types (for instance, that `@NonNull Object` is a subtype of `@Nullable Object`).

22.4 Type introduction rules. For some types and expressions, a qualifier should be treated as implicitly present even if a programmer did not explicitly write it. For example, in the Nullness type system every literal other than `null` has a `@NonNull` type; examples of literals include `"some string"` and `java.util.Date.class`.

22.5 Type rules. You specify the type system semantics (type rules), violation of which yields a type error. There are two types of rules.

- Subtyping rules related to the type hierarchy, such as that every assignment and pseudo-assignment satisfies a subtyping relationship. Your checker automatically inherits these subtyping rules from the Base Checker (Chapter 15).
- Additional rules that are specific to your particular checker. For example, in the Nullness type system, only references with a `@NonNull` type may be dereferenced. You write these additional rules yourself.

22.6 Interface to the compiler. The compiler interface indicates which annotations are part of the type system, which command-line options and `@SuppressWarnings` annotations the checker recognizes, etc.

22.3 Annotations: Type qualifiers and hierarchy

A type system designer specifies the qualifiers in the type system and the type hierarchy that relates them.

Type qualifiers are defined as Java annotations [Dar06]. In Java, an annotation is defined using the Java `@interface` keyword. For example:

```
// Define an annotation for the @NonNull type qualifier.
@interface NonNull
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface NonNull { }
```

Write the `@TypeQualifier` meta-annotation on the annotation definition to indicate that the annotation represents a type qualifier and should be processed by the checker. Also write a `@Target` meta-annotation to indicate where the annotation may be written. (An annotation that is written on an annotation definition, such as `@TypeQualifier`, is called a *meta-annotation*.)

The type hierarchy induced by the qualifiers can be defined either declaratively via meta-annotations (Section 22.3.1), or procedurally through subclassing `QualifierHierarchy` or `TypeHierarchy` (Section 22.3.2).

22.3.1 Declaratively defining the qualifier and type hierarchy

Declaratively, the type system designer uses two meta-annotations (written on the declaration of qualifier annotations) to specify the qualifier hierarchy.

- `@SubtypeOf` denotes that a qualifier is a subtype of another qualifier or qualifiers, specified as an array of class literals. For example, for any type T , `@NonNull T` is a subtype of `@Nullable T`:

```
@TypeQualifier
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
@SubtypeOf( { Nullable.class } )
public @interface NonNull { }
```

`@SubtypeOf` accepts multiple annotation classes as an argument, permitting the type hierarchy to be an arbitrary DAG. For example, in the IGJ type system (Section 5.2), `@Mutable` and `@Immutable` induce two mutually exclusive subtypes of the `@ReadOnly` qualifier.

All type qualifiers, except for polymorphic qualifiers (see below and also Section 18.2), need to be properly annotated with `SubtypeOf`.

The top qualifier is annotated with `@SubtypeOf({ })`. The top qualifier is the qualifier that is a supertype of all other qualifiers. For example, `@Nullable` is the top qualifier of the Nullness type system, hence is defined as:

```
@TypeQualifier
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
@SubtypeOf( { } )
public @interface Nullable { }
```

If the top qualifier of the hierarchy is the unqualified type, then its children will use `@SubtypeOf(Unqualified.class)`, but no `@SubtypeOf({ })` annotation on the top qualifier is necessary. For an example, see the Encrypted type system of Section 15.2.

- `@PolymorphicQualifier` denotes that a qualifier is a polymorphic qualifier. For example:

```
@TypeQualifier
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
@PolymorphicQualifier
public @interface PolyNull { }
```

For a description of polymorphic qualifiers, see Section 18.2. A polymorphic qualifier needs no `@SubtypeOf` meta-annotation and need not be mentioned in any other `@SubtypeOf` meta-annotation.

The declarative and procedural mechanisms for specifying the hierarchy can be used together. In particular, when using the `@SubtypeOf` meta-annotation, further customizations may be performed procedurally (Section 22.3.2) by overriding the `isSubtype` method in the checker class (Section 22.6). However, the declarative mechanism is sufficient for most type systems.

22.3.2 Procedurally defining the qualifier and type hierarchy

While the declarative syntax suffices for many cases, more complex type hierarchies can be expressed by overriding, in `BaseTypeChecker`, either `createQualifierHierarchy` or `createTypeHierarchy` (typically only one of these needs to be overridden). For more details, see the Javadoc of those methods and of the classes `QualifierHierarchy` and `TypeHierarchy`.

The `QualifierHierarchy` class represents the qualifier hierarchy (not the type hierarchy), e.g., `Mutable` is a subtype of `ReadOnly`. A type-system designer may subclass `QualifierHierarchy` to express customized qualifier relationships (e.g., relationships based on annotation arguments).

The `TypeHierarchy` class represents the type hierarchy — that is, relationships between annotated types, rather than merely type qualifiers, e.g., `@Mutable Date` is a subtype of `@ReadOnly Date`. The default `TypeHierarchy` uses `QualifierHierarchy` to determine all subtyping relationships. The default `TypeHierarchy` handles generic

type arguments, array components, type variables, and wildcards in a similar manner to the Java standard subtype relationship but with taking qualifiers into consideration. Some type systems may need to override that behavior. For instance, the Java Language Specification specifies that two generic types are subtypes only if their type arguments are identical: for example, `List<Date>` is not a subtype of `List<Object>`, or of any other generic `List`. (In the technical jargon, the generic arguments are “invariant” or “novariant”.) The Javari type system overrides this behavior to allow some type arguments to change covariantly in a type-safe manner (e.g., `List<@Mutable Date>` is a subtype of `List<@QReadOnly Date>`).

22.3.3 Defining a default annotation

A type system designer may set a default annotation. A user may override the default; see Section 19.1.1.

The type system designer may specify a default annotation declaratively, using the `@DefaultQualifierInHierarchy` meta-annotation. Note that the default will apply to any source code that the checker reads, including stub libraries, but will not apply to compiled `.class` files that the checker reads.

Alternately, the type system designer may specify a default procedurally, by calling the `QualifierDefaults.addAbsoluteDefault` method. You may do this even if you have declaratively defined the qualifier hierarchy; see the Nullness checker’s implementation for an example.

Recall that defaults are distinct from implicit annotations; see Sections 19.1 and 22.4.

22.3.4 Completeness of the type hierarchy

Bottom qualifier It is usually a good idea to have a bottom qualifier in your type hierarchy — a qualifier that is a (direct or indirect) subtype of every other qualifier. For instance, the hierarchy of Figure 5.1 lacks a bottom qualifier, because there is no qualifier that is a subtype of both `@Immutable` and `@Mutable`. The bottom qualifier is the natural type for the null value, which can be viewed as having any type at all. Without a bottom qualifier, type-checking becomes less precise. Users should never write the bottom qualifier explicitly; it is merely used for the null value.

The actual IGJ hierarchy does contain a (non-user-visible) bottom qualifier, defined like this:

```
@TypeQualifier
@SubtypeOf({Mutable.class, Immutable.class, I.class})
@Target({}) // forbids a programmer from writing it in a program
@ImplicitFor(trees = { Kind.NULL_LITERAL, Kind.CLASS, Kind.NEW_ARRAY },
             typeClasses = { AnnotatedPrimitiveType.class })
@interface IGJBottom { }
```

Top qualifier Similarly, it is usually a good idea to have a top qualifier in your type hierarchy — a qualifier that is a (direct or indirect) supertype of every other qualifier. When a type system lacks a top qualifier (or any other qualifier), then users lose flexibility in expressing defaults.

For instance, the `@Encrypted` type system of Section 15.2 lacks an explicit top qualifier:

```
@TypeQualifier
@SubtypeOf(Unqualified.class)
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface Encrypted { }
```

The default is `@Unqualified`, and there is no way for a user to change that. (A user could not specify `@DefaultQualifier(Unqualified.class)` because many type systems share `@Unqualified` and it is not clear which type system the annotation is intended to refer to.)

By contrast, the Nullness type system has an explicit qualifier for every possible meaning: both `@Nullable` and `@NonNull`. Because it has no built-in meaning for unannotated types; a user may specify a default qualifier. This greater flexibility for the user is usually preferable.

There are reasons to omit the top qualifier. The ability to omit the top qualifier is a convenience when writing a type system, because it reduces the number of qualifiers that must be defined; this is especially convenient when using the Basic Checker (Section 15). More importantly, omitting the top qualifier restricts the user in ways that the type system designer may have intended.

22.4 Type factory: Implicit annotations

For some types and expressions, a qualifier should be treated as present even if a programmer did not explicitly write it. For example, every literal (other than `null`) has a `@NonNull` type.

The implicit annotations may be specified declaratively and/or procedurally.

22.4.1 Declaratively specifying implicit annotations

The `@ImplicitFor` meta-annotation indicates implicit annotations. When written on a qualifier, `ImplicitFor` specifies the trees (AST nodes) and types for which the framework should automatically add that qualifier.

In short, the types and trees can be specified via any combination of five fields in `ImplicitFor`:

- `trees`: an array of `com.sun.source.tree.Tree.Kind`, e.g., `NEW_ARRAY` or `METHOD_INVOCATION`
- `types`: an array of `TypeKind`, e.g., `ARRAY` or `BOOLEAN`
- `treeClasses`: an array of class literals for classes implementing `Tree`, e.g., `LiteralTree.class` or `ExpressionTree.class`
- `typeClasses`: an array of class literals for classes implementing `javac.lang.model.type.TypeMirror`, e.g., `javac.lang.model.type.PrimitiveType`. Often you should use a subclass of `AnnotatedTypeMirror`.
- `stringPatterns`: an array of regular expressions that will be matched against string literals, e.g., `"[01]+"` for a binary number. Useful for annotations that indicate the format of a string.

For example, consider the definitions of the `@NonNull` and `@Nullable` type qualifiers:

```
@TypeQualifier
@SubtypeOf( { Nullable.class } )
@ImplicitFor(
    types={TypeKind.PACKAGE},
    typeClasses={AnnotatedPrimitiveType.class},
    trees={
        Tree.Kind.NEW_CLASS,
        Tree.Kind.NEW_ARRAY,
        Tree.Kind.PLUS,
        // All literals except NULL_LITERAL:
        Tree.Kind.BOOLEAN_LITERAL, Tree.Kind.CHAR_LITERAL, Tree.Kind.DOUBLE_LITERAL, Tree.Kind.FLOAT_LITERAL,
        Tree.Kind.INT_LITERAL, Tree.Kind.LONG_LITERAL, Tree.Kind.STRING_LITERAL
    })
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface NonNull { }

@TypeQualifier
@SubtypeOf({})
@ImplicitFor(trees={Tree.Kind.NULL_LITERAL})
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface Nullable { }
```

For more details, see the Javadoc for the `ImplicitFor` annotation, and the Javadoc for the javac classes that are linked from it. You only need to understand a small amount about the javac AST, such as the `Tree.Kind` and `TypeKind` enums. All the information you need is in the Javadoc, and Section 22.9 can help you get started.

22.4.2 Procedurally specifying implicit annotations

The Checker Framework provides a representation of annotated types, `AnnotatedTypeMirror`, that extends the standard `TypeMirror` interface but integrates a representation of the annotations into a type representation. A checker's

type factory class, given an AST node, returns the annotated type of that expression. The Checker Framework's abstract *base type factory* class, `AnnotatedTypeFactory`, supplies a uniform, Tree-API-based interface for querying the annotations on a program element, regardless of whether that element is declared in a source file or in a class file. It also handles default annotations, and it optionally performs flow-sensitive local type inference.

`AnnotatedTypeFactory` inserts the qualifiers that the programmer explicitly inserted in the code. Yet, certain constructs should be treated as having a type qualifier even when the programmer has not written one. The type system designer may subclass `AnnotatedTypeFactory` and override `annotateImplicit(Tree, AnnotatedTypeMirror)` and `annotateImplicit(Element, AnnotatedTypeMirror)` to account for such constructs.

22.4.3 Flow-sensitive type qualifier inference

The Checker Framework provides automatic type refinement as described in Section 19.1.2.

Class `BasicAnnotatedTypeFactory` provides a 3 parameter constructor that allows subclasses to disable flow inference. By default the 2 parameter constructor performs flow inference. To disable flow inference, call `super(checker, root, false)`; in your subtype of `BasicAnnotatedTypeFactory`.

22.5 Visitor: Type rules

A type system's rules define which operations on values of a particular type are forbidden. These rules must be defined procedurally, not declaratively.

The Checker Framework provides a *base visitor class*, `BaseTypeVisitor`, that performs type-checking at each node of a source file's AST. It uses the visitor design pattern to traverse Java syntax trees as provided by Oracle's Tree API, and it issues a warning whenever the type system is violated.

A checker's visitor overrides one method in the base visitor for each special rule in the type qualifier system. Most type-checkers override only a few methods in `BaseTypeVisitor`. For example, the visitor for the Nullness type system of Chapter 3 contains a single 4-line method that warns if an expression of nullable type is dereferenced, as in:

```
myObject.hashCode(); // invalid dereference
```

By default, `BaseTypeVisitor` performs subtyping checks that are similar to Java subtype rules, but taking the type qualifiers into account. `BaseTypeVisitor` issues these errors:

- invalid assignment (`type.incompatible`) for an assignment from an expression type to an incompatible type. The assignment may be a simple assignment, or pseudo-assignment like return expressions or argument passing in a method invocation

In particular, in every assignment and pseudo-assignment, the left-hand side of the assignment is a supertype of (or the same type as) the right-hand side. For example, this assignment is not permitted:

```
@Nullable Object myObject;
@NonNull Object myNonNullObject;
...
myNonNullObject = myObject; // invalid assignment
```

- invalid generic argument (`generic.argument.invalid`) when a type is bound to an incompatible generic type variable
- invalid method invocation (`method.invocation.invalid`) when a method is invoked on an object whose type is incompatible with the method receiver type
- invalid overriding parameter type (`override.parameter.invalid`) when a parameter in a method declaration is incompatible with that parameter in the overridden method's declaration
- invalid overriding return type (`override.return.invalid`) when a parameter in a method declaration is incompatible with that parameter in the overridden method's declaration
- invalid overriding receiver type (`override.receiver.invalid`) when a receiver in a method declaration is incompatible with that receiver in the overridden method's declaration

22.5.1 AST traversal

The Checker Framework needs to do its own traversal of the AST even though it operates as an ordinary annotation processor [Dar06]. Annotation processors can utilize a visitor for Java code, but that visitor only visits the public elements of Java code, such as classes, fields, methods, and method arguments — it does not visit code bodies or various other locations. The Checker Framework hardly uses the built-in visitor — as soon as the built-in visitor starts to visit a class, then the Checker Framework’s visitor takes over and visits all of the class’s source code.

Because there is no standard API for the AST of Java code¹, the Checker Framework uses the `javac` implementation. This is why the Checker Framework is not deeply integrated with Eclipse, but runs as an external tool (see Section 23.6).

22.5.2 Avoid hardcoding

It may be tempting to write a type-checking rule for method invocation, where your rule checks the name of the method being called and then treats the method in a special way. This is usually the wrong approach. It is better to write annotations, in a stub file (Chapter 21), and leave the work to the standard type-checking rules.

22.6 The checker class: Compiler interface

A checker’s entry point is a subclass of `BaseTypeChecker`. This entry point, which we call the checker class, serves two roles: an interface to the compiler and a factory for constructing type-system classes.

Because the Checker Framework provides reasonable defaults, oftentimes the checker class has no work to do. Here are the complete definitions of the checker classes for the Interning and Nullness checkers:

```
@TypeQualifiers({ Interned.class, PolyInterned.class })
@SupportedLintOptions({"dotequals"})
public final class InterningChecker extends BaseTypeChecker { }

@TypeQualifiers({ Nullable.class, Raw.class, NonNull.class, PolyNull.class })
@SupportedLintOptions({"flow", "cast", "cast:redundant"})
public class NullnessChecker extends BaseTypeChecker { }
```

The checker class must be annotated by `@TypeQualifiers`, which lists the annotations that make up the type hierarchy for this checker (including polymorphic qualifiers), provided as an array of class literals. Each one is a type qualifier whose definition bears the `@TypeQualifier` meta-annotation (or is returned by the `BaseTypeChecker.getSupportedTypeQualifiers` method).

The checker class bridges between the compiler and the rest of the checker. It invokes the type-rule check visitor on every Java source file being compiled, and provides a simple API, `report`, to issue errors using the compiler error reporting mechanism.

Also, the checker class follows the factory method pattern to construct the concrete classes (e.g., visitor, factory) and annotation hierarchy representation. It is a convention that, for a type system named `Foo`, the compiler interface (checker), the visitor, and the annotated type factory are named as `FooChecker`, `FooVisitor`, and `FooAnnotatedTypeFactory`. `BaseTypeChecker` uses the convention to reflectively construct the components. Otherwise, the checker writer must specify the component classes for construction.

A checker can customize the default error messages through a `Properties`-loadable text file named `messages.properties` that appears in the same directory as the checker class. The property file keys are the strings passed to `report` (like `type.incompatible`) and the values are the strings to be printed ("`cannot assign ...`"). The `messages.properties` file only need to mention the new messages that the checker defines. It is also allowed to override messages defined in superclasses, but this is rarely needed. For more details about message keys, see Section 20.2.1 (page 84).

¹Actually, there is a standard API for Java ASTs — JSR 198 (Extension API for Integrated Development Environments) [Cro06]. If tools were to implement it (which would just require writing wrappers or adapters), then the Checker Framework and similar tools could be portable among different compilers and IDEs.

22.6.1 Bundling multiple checkers

To run a checker, a user supplies the `-processor` command-line option. There are two ways to run multiple related checkers as a unit.

1. A user can pass multiple `-processor` command-line options, like:

```
javac -processor DistanceUnitChecker -processor SpeedUnitChecker ... files ...
```

This is verbose, and it is also error-prone, since a user might omit one of several related checkers that are designed to be run together.

2. You can define an aggregate checker class that combines multiple checkers. Extend `AggregateChecker` and override the `getSupportedTypeCheckers` method, like the following:

```
public class UnitCheckers extends AggregateChecker {
    protected Collection<Class<? extends SourceChecker>> getSupportedCheckers() {
        return Arrays.asList(DistanceUnitChecker.class, SpeedUnitChecker.class);
    }
}
```

Now, a user can pass a single `-processor` argument on the command line:

```
javac -processor UnitCheckers ... files ...
```

22.6.2 Providing command-line options

A checker can provide two kinds of command-line options: boolean flags and named string values (the standard annotation processor options).

Boolean flags

To specify a simple boolean flag, add:

```
@SupportedLintOptions({"flag"})
```

to your checker subclass. The value of the flag can be queried using

```
checker.getLintOption("flag", false)
```

The second argument sets the default value that should be returned.

To pass a flag on the command line, call `javac` as follows:

```
javac -processor Mine -Aflag=flag
```

Named string values

For more complicated options, one can use the standard annotation processing `SupportedOptions` annotation on the checker, as in:

```
@SupportedOptions({"info"})
```

The value of the option can be queried using

```
env.getOptions().get("info")
```

where `env` is the current `ProcessingEnvironment`.

To pass an option on the command line, call `javac` as follows:

```
javac -processor Mine -Ainfo=p1,p2
```

The value is returned as a single string and you have to perform the required parsing of the option.

22.7 Testing framework

The Checker Framework comes with a testing framework that is used for testing the distributed checkers. It is easy to use this testing framework to ensure correctness of your checker!

You first need to provide a subclass of `ParameterizedCheckerTest` that determines the checker to use and all command-line options that should be provided. This class can be run as a JUnit test runner. Note that you always need to use the `-Anomsgtext` option to suppress the substitution of message keys by human-readable values. See the test setup classes in directory `tests/src/tests/` for examples.

Locate all your test cases in a subdirectory of the `tests` directory. The individual test cases are normal Java files that use stylized comments to indicate expected error messages. For example, consider this test case from the Nullness checker:

```
//:: error: (dereference.of.nullable)
s.toString();
```

An expected error message is introduced by the `//::` comment. The next token is either `error:` or `warning:`, distinguishing what kind of message is expected. Finally, the message key for the expected message is given.

Multiple expected messages can be given using the `"//:: A :: B :: C"` syntax.

As an alternative, expected errors can be specified in a separate file using the `.out` file extension. These files are of the following format:

```
:19: error: (dereference.of.nullable)
```

The number between the colons is the line number of the expected error message. This format is a lot harder to maintain and we suggest using the in-line comment format.

22.8 Debugging options

The Checker Framework provides debugging options that can be helpful when writing a checker. These are provided via the standard `javac` `-A` switch, which is used to pass options to an annotation processor.

- `-Anomsgtext`: use message keys (such as `"type.invalid"`) rather than full message text when reporting errors or warnings
- `-Ashowchecks`: print debugging information for each pseudo-assignment check (as performed by `BaseTypeVisitor`; see Section 22.5 above)
- `-Afilenames`: print the name of each file before type-checking it
- `-AprintErrorStack`: print a stack trace together with internal Checker Framework error messages
- `-AprintAllQualifiers`: print all type qualifiers, including qualifiers like `@Unqualified` which are usually not shown. (Use the `@InvisibleQualifier` meta-annotation on a qualifier to hide it.)
- `-Aignorejdkastub`: ignore the `jdk.astub` file in the checker directory. Files passed through the `-Astubs` option are still processed. This is useful when compiling the source code that is described by the stub file and experimenting with an alternative stub file.

The following example demonstrates how these options are used:

```
$ javac -processor checkers.interning.InterningChecker \
  examples/InternedExampleWithWarnings.java -Ashowchecks -Anomsgtext -Afilenames
```

```
[InterningChecker] InterningExampleWithWarnings.java
success (line 18): STRING_LITERAL "foo"
  actual: DECLARED @checkers.interningquals.Interned java.lang.String
  expected: DECLARED @checkers.interningquals.Interned java.lang.String
success (line 19): NEW_CLASS new String("bar")
  actual: DECLARED java.lang.String
  expected: DECLARED java.lang.String
```

```

examples/InterningExampleWithWarnings.java:21: (not.interned)
    if (foo == bar)
        ^
success (line 22): STRING_LITERAL "foo == bar"
    actual: DECLARED @checkers.interningquals.Interned java.lang.String
    expected: DECLARED java.lang.String
1 error

```

You can use any standard debugger to observe the execution of your checker. Set the execution main class to `com.sun.tools.javac.Main`, and insert the JSR 308 `javac.jar` (resides in `.../jsr308-langtools/dist/lib/javac.jar`). If using an IDE, it is recommended that you add `.../jsr308-langtools` as a project, so you can step into its source code if needed.

22.9 javac implementation survival guide

A checker built using the Checker Framework makes use of a few interfaces from the underlying compiler (Oracle's OpenJDK `javac`). This section describes those interfaces.

22.9.1 Checker access to compiler information

The compiler uses and exposes three hierarchies to model the Java source code and classfiles.

Types - Java Language Model API

A `TypeMirror` represents a Java type.

There is a `TypeMirror` interface to represent each type kind, e.g., `PrimitiveType` for primitive types, `ExecutableType` for method types, and `NullType` for the type of the `null` literal.

`TypeMirror` does not represent annotated types though. Checkers should use the Checker Framework types API, `AnnotatedTypeMirror`, instead. `AnnotatedTypeMirror` parallels the `TypeMirror` API, but also present the type annotations associated with the type.

The Checker Framework and the checkers use the types API extensively.

Elements - Java Language Model API

An `Element` represents a potentially-public declaration that can be accessed from elsewhere: classes, interfaces, methods, constructors, and fields. `Element` represents elements found in both source code and bytecode.

There is an `Element` interface to represent each construct, e.g., `TypeElement` for class/interfaces, `ExecutableElement` for methods/constructors, `VariableElement` for local variables and method parameters.

If you need to operate on the declaration level, always use elements rather than trees (see below). This allows the code to work on both source and bytecode elements.

Example: retrieve declaration annotations, check variable modifiers (e.g., `strictfp`, `synchronized`)

Trees - Compiler Tree API

A `Tree` represents a syntactic unit in the source code, like a method declaration, statement, block, `for` loop, etc. Trees only represent source code to be compiled (or found in `-sourcepath`); no tree is available for classes read from bytecode.

There is a `Tree` interface for each Java source structure, e.g., `ClassTree` for class declaration, `MethodInvocationTree` for a method invocation, and `ForEachTree` for an enhanced-for-loop statement.

You should limit your use of trees. Checkers use Trees mainly to traverse the source code and retrieve the types/elements corresponding to them. Then, the checker performs any needed checks on the types/elements instead.

Using the APIs

The three APIs use some common idioms and conventions; knowing them will help you to create your checker.

Type-checking: Do not use `instanceof` to determine the class of the object, because you cannot necessarily predict the run-time type of the object that implements an interface. Instead, use the `getKind()` method. The method returns `TypeKind`, `ElementKind`, and `Tree.Kind` for the three interfaces, respectively.

Visitors and Scanners: The compiler and the Checker Framework use the visitor pattern extensively. For example, visitors are used to traverse the source tree (`BaseTypeVisitor` extends `TreePathScanner`) and for type checking (`TreeAnnotator` implements `TreeVisitor`).

Utility classes: Some useful methods appear in a utility class. The Oracle convention is that the utility class for a `Foo` hierarchy is `Foos` (e.g., `Types`, `Elements`, and `Trees`). The Checker Framework uses a common `Utils` suffix instead (e.g., `TypesUtils`, `TreeUtils`, `ElementUtils`), with one notable exception: `AnnotatedTypes`.

22.9.2 How a checker fits in the compiler as an annotation processor

The Checker Framework builds on the Annotation Processing API introduced in Java 6. A type annotation processor is one that extends `AbstractTypeProcessor`; these get run on each class source file after the compiler confirms that the class is valid Java code.

The most important methods of `AbstractTypeProcessor` are `typeProcess` and `getSupportedSourceVersion`. The former class is where you would insert any sort of method call to walk the AST, and the latter just returns a constant indicating that we are targeting version 8 of the compiler. Implementing these two methods should be enough for a basic plugin; see the Javadoc for the class for other methods that you may find useful later on.

The Checker Framework uses Oracle's Tree API to access a program's AST. The Tree API is specific to the Oracle OpenJDK, so the Checker Framework only works with the OpenJDK `javac`, not with Eclipse's compiler `ecj` or with `gcj`. This also limits the tightness of the integration of the Checker Framework into other IDEs such as IntelliJ IDEA. An implementation-neutral API would be preferable. In the future, the Checker Framework can be migrated to use the Java Model AST of JSR 198 (Extension API for Integrated Development Environments) [Cro06], which gives access to the source code of a method. But, at present no tools implement JSR 198. Also see Section 22.5.1.

Learning more about javac

Sun's `javac` compiler interfaces can be daunting to a newcomer, and its documentation is a bit sparse. The Checker Framework aims to abstract a lot of these complexities. You do not have to understand the implementation of `javac` to build powerful and useful checkers. Beyond this document, other useful resources include the Java Infrastructure Developer's guide at http://wiki.netbeans.org/Java_DevelopersGuide and the compiler mailing list archives at <http://news.gmane.org/gmane.comp.java.openjdk.compiler.devel> (subscribe at <http://mail.openjdk.java.net/mailman/listinfo/compiler-dev>).

Chapter 23

Integration with external tools

This chapter discusses how to run a checker from your favorite IDE.

Or, if your favorite isn't here, you should customize how it runs the `javac` command on your behalf. See the IDE documentation to learn how to customize it, adapting the instructions for `javac` in Section 2.2. If you make another tool support running a checker, please inform us via the mailing list or issue tracker so we can add it to this manual.

This chapter also discusses type inference tools (see Section 23.8).

23.1 Javac Compiler

If you use `javac` compiler from the command line, then you can use the Type annotations compiler (a variant of the OpenJDK `javac`) that is bundled with the Checker Framework. The bundled `javac` recognizes type annotations, and annotations in comments (see Section 20.3.1). (Eventually, you will be able to use any Java compiler, such as the OpenJDK compiler, but Oracle has been slow to incorporate all the patches, so the bundled `javac` is superior, for the purpose of pluggable type-checking, and is equivalent in all other respects.)

This section describes how you can install and use the bundled `javac`, using either Unix/Linux/macOS (see Section 23.1.1) or Windows (see Section 23.1.2). The instructions are identical to those of Section 1.2, but are given as commands that you can cut and paste into your command shell.

23.1.1 Unix/Linux/macOS installation

These instructions assume that you use the `bash` or `sh` shell. If you use a different shell, you may need to slightly adjust the commands.

1. Download the latest Checker Framework distribution and unzip it. You can put it anywhere you like by changing the definition of environment variable `JSR308` below; a standard place is in a new directory named `jsr308`.

```
export JSR308=$HOME/jsr308
mkdir -p ${JSR308}
cd ${JSR308}
# or: wget http://types.cs.washington.edu/checker-framework/current/checkers.zip
curl -O http://types.cs.washington.edu/checker-framework/current/checkers.zip
unzip checkers.zip
chmod +x checker-framework/checkers/binary/javac
checker-framework/checkers/binary/javac -version
```

The output of the last command should be:

```
javac 1.7.0-jsr308-1.4.0-b1
```

2. Place the following commands in your `.bashrc` file:


```
export JSR308=$HOME/jsr308
export CHECKERS=$JSR308/checker-framework/checkers
export PATH=$CHECKERS/binary:${PATH}
```

Also execute them on the command line, or log out and back in. Then, verify that the installation works. From the command line, run:

```
javac -version
```

The output should be:

```
javac 1.7.0-jsr308-1.4.0-b1
```

That's all there is to it! Now you are ready to start using the checkers with the new javac compiler.

23.1.2 Windows installation

1. Download the latest Checker Framework distribution and unzip it to create a checkers directory. You can put it anywhere you like; a standard place is in a new directory under C:\Program Files.

- (a) Save the file <http://types.cs.washington.edu/checker-framework/current/checkers.zip> to your Desktop.
- (b) Double-click the checkers.zip file on your computer. Click on the checkers directory, then Select Extract all files, and use C:\Program Files as the destination. You will obtain a new C:\Program Files\checker-framework folder.
- (c) Verify that the installation works. From a Windows command prompt, run:

```
set CHECKERS = C:\Program Files\checker-framework\checkers
java -Xbootclasspath/p:%CHECKERS%\binary\jsr308-all.jar -jar C:%CHECKERS%\binary\jsr308-all.jar
```

The output should be:

```
javac 1.7.0-jsr308-1.4.0-b1
```

2. In order to use the updated compiler when you type javac, add the directory C:\Program Files\checker-framework\checkers to the beginning of your path variable. Also set a CHECKERS variable.

To set an environment variable, you have two options: make the change temporarily or permanently.

- To make the change **temporarily**, type at the command shell prompt:

```
path = newdir;%PATH%
```

For example:

```
set CHECKERS = C:\Program Files\checker-framework\checkers
path = %CHECKERS%\binary;%PATH%
```

This is a temporary change that endures until the window is closed, and you must re-do it every time you start a new command shell.

- To make the change **permanently**, Right-click the My Computer icon and select Properties. Select the Advanced tab and click the Environment Variables button. You can set the variable as a “System Variable” (visible to all users) or as a “User Variable” (visible to just this user). Both work; the instructions below show how to set as a “System Variable”. In the System Variables pane, select Path from the list and click Edit. In the Edit System Variable dialog box, move the cursor to the beginning of the string in the Variable Value field and type the full directory name (not using the %CHECKERS% environment variable) followed by a semicolon (;).

Similarly, set the CHECKERS variable.

This is a permanent change that only needs to be done once ever.

Now, verify that the installation works. From the command line, run:

```
javac -version
```

The output should be:

```
javac 1.7.0-jsr308-1.4.0-b1
```

23.2 Ant task

If you use the Ant build tool to compile your software, then you can add an Ant task that runs a checker. We assume that your Ant file already contains a compilation target that uses the `javac` task.

1. Set the `jsr308javac` property:

```
<property environment="env"/>

<dirname property="checkers" file="${env.CHECKERS}" />

<presetdef name="jsr308.javac">
  <javac fork="yes">
    <!-- JSR308 related compiler arguments -->
    <compilerarg value="-version"/>
    <!-- optional, so .class files work with older JVMs: <compilerarg line="-target 5"/> -->
    <compilerarg value="-implicit:class"/>
    <!-- optional, to issue warnings instead of errors: <compilerarg line="-Awarns -Xmaxwarns 10000"/> -->
    <compilerarg value="-J-Xbootclasspath/p:${checkers}/binary/jsr308-all.jar"/>

    <classpath>
      <pathelement location="${checkers}/checkers.jar"/>
    </classpath>
  </javac>
</presetdef>
```

2. Duplicate the compilation target, then modify it slightly as indicated in this example:

```
<target name="check-nullness"
  description="Check for null pointer dereferences"
  depends="clean,...">
  <!-- use jsr308.javac instead of javac -->
  <jsr308.javac ... >
    <compilerarg line="-processor checkers.nullness.NullnessChecker"/>
    <compilerarg value="-Xbootclasspath/p:${checkers}/jdk/jdk.jar"/>
    <!-- optional, for implicit imports: <compilerarg value="-J-Djsr308_imports=checkers.nullnessquals.*"/> -->
    <!-- optional, to not check uses of library methods: <compilerarg value="-AskipUses^(java\.awt\.|javax\.swing\.)"/> -->
    ...
  </jsr308.javac>
</target>
```

Fill in each ellipsis (...) from the original compilation target. But, don't include any `-source` argument with value other than 1.8 or 8. Doing so will disable the annotations in comments feature (see Section 20.3.1, page 86).

In the example, the target is named `check-nullness`, but you can name it whatever you like.

23.2.1 Explanation

This section explains each part of the Ant task.

1. Definition of `jsr308.javac`:

The `fork` field of the `javac` task ensures that an external `javac` program is called. Otherwise, Ant will run `javac` via a Java method call, and there is no guarantee that it will get the JSR 308 version that is distributed with the Checker Framework.

The `-version` compiler argument is just for debugging; you may omit it.

The `-target 5` compiler argument is optional, if you use Java 5 in ordinary compilation when not performing pluggable type-checking (see Section 20.3.4, page 88).

The `-implicit:class` compiler argument causes annotation processing to be performed on implicitly compiled files. (An implicitly compiled file is one that was not specified on the command line, but for which the source code is newer than the `.class` file.) This is the default, but supplying the argument explicitly suppresses a compiler warning.

The `-Awarns ...` compiler argument is optional, and causes the checker to treat errors as warnings so that compilation does not fail even if pluggable type-checking fails; see Section 2.2.1.

2. The `check-nullness` target:

The target assumes the existence of a `clean` target that removes all `.class` files. That is necessary because Ant's `javac` target doesn't re-compile `.java` files for which a `.class` file already exists.

The `-processor ...` compiler argument indicates which checker to run. You can supply additional arguments to the checker as well.

23.3 Maven plugin

If you use the Maven project tool, then you can specify the distributed checkers as part of your build process.

1. First, you need to add the repositories in your `pom.xml` file:

```
<repositories>
  <repository>
    <id>checker-framework-repo</id>
    <url>http://types.cs.washington.edu/m2-repo</url>
  </repository>
</repositories>
<pluginRepositories>
  <pluginRepository>
    <id>checker-framework-repo</id>
    <url>http://types.cs.washington.edu/m2-repo</url>
  </pluginRepository>
</pluginRepositories>
```

2. Then, to use the annotations used by the distributed checkers, you'll have to declared as a dependency:

```
<dependencies>
  <!-- annotations for the standard checkers: nullness, interning, mutability -->
  <dependency>
    <groupId>types.checkers</groupId>
    <artifactId>checkers-quals</artifactId>
    <version>1.1.1</version>
  </dependency>

  <!-- other dependencies -->
</dependencies>
```

3. And finally, you need to attach the plugin to your build lifecycle:

```
<build>
  <plugins>
    <plugin>
      <groupId>types.checkers</groupId>
      <artifactId>checkersplugin</artifactId>
      <version>0.1</version>
      <executions>
        <execution>
          <!-- run the checkers after compilation; this can also be any later phase -->
          <phase>process-classes</phase>
          <goals>
            <goal>check</goal>
          </goals>
        </execution>
      </executions>
    </plugin>
  </plugins>
</build>
```

```

        </execution>
    </executions>
    <configuration>
        <!-- required configuration options -->
        <!-- a list of processors to run -->
        <processors>
            <processor>checkers.nullness.NullnessChecker</processor>
            <processor>checkers.interning.InterningChecker</processor>
        </processors>

        <!-- other optional configuration -->
        <!-- full path to a java executable that should be used to create the forked JVM -->
        <executable>/opt/javal.6/bin/java</executable>
        <!-- should an error reported by a checker cause a build failure, or only be logged -->
        <failOnError>true|false</failOnError>
        <!-- a list of patterns to include, in the standard maven syntax; defaults to **/* -->
        <includes>
            <include>org/company/important/**/*.*.java</include>
        </includes>
        <!-- a list of patterns to exclude, in the standard maven syntax; defaults to an empty list -->
        <excludes>
            <exclude>org/company/notimportant/**/*.*.java</exclude>
        </excludes>
        <!-- additional parameters passed to the JSR308 java compiler -->
        <javacParams>-Alint</javacParams>
        <!-- additional parameters to pass to the forked JVM -->
        <javaParams>-Xdebug</javaParams>
        <!-- versions of checkers to use; defaults to the current newest version: 1.0.6 -->
        <checkersVersion>0.8.8</checkersVersion>
    </configuration>
</plugin>
</plugins>
</build>

```

The plugin was contributed by Adam Warski.

23.4 Gradle

Gradle lets you add command-line arguments to a javac invocation by setting the compilerArgs property of the compiler options. This is adequate for running the Checker Framework, because it is run by specifying command-line arguments. See the instructions elsewhere in this manual for a list of command-line arguments.

To specify command-line arguments, set `compile.options.compilerArgs`. Here is a possible example:

```

allprojects {
    tasks.withType(Compile).allTasks { Compile compile ->
        compile.options.debug = true
        compile.options.compilerArgs = [
            '-version',
            '-implicit:class',
            '-J-Xbootclasspath/p:${env.CHECKERS}/binary/jsr308-all.jar',

```

```

        '-processor', 'checkers.nullness.NullnessChecker',
        '-Xbootclasspath/p:${env.CHECKERS}/jdk/jdk.jar' ]
    }
}

```

You don't need to use a special version of `javac`; you only need to give the `-J-Xbootclasspath/p:...` argument. If you choose to use a special version of `javac` instead of supplying the command-line argument, then you can do so in the following way:

```

tasks.withType(Compile).allObjects { compile ->
    compile.options.fork.executable = "$CHECKERS/binary/javac"
}

```

23.5 IntelliJ IDEA

IntelliJ IDEA (Maia release) supports the Type Annotations (JSR-308) syntax. See <http://blogs.jetbrains.com/idea/2009/07/type-annotations-jsr-308-support/>.

23.6 Eclipse

There are two ways to run a checker from within the Eclipse IDE: via Ant or using an Eclipse plug-in.

Using an Ant task Add an Ant target as described in Section 23.2. You can run the Ant target by executing the following steps (instructions copied from http://www.eclipse.org/documentation/?topic=/org.eclipse.platform.doc.user/gettingStarted/qs-84_run_ant.htm):

1. Select `build.xml` in one of the navigation views and choose **Run As > Ant Build...** from its context menu.
2. A launch configuration dialog is opened on a launch configuration for this Ant buildfile.
3. In the **Targets** tab, select the new ant task (e.g., `check-interning`).
4. Click **Run**.
5. The Ant buildfile is run, and the output is sent to the Console view.

Eclipse plug-in for the Checker Framework The Checker Plugin is an Eclipse plugin that enables the use of the Checker Framework. Its website (<http://types.cs.washington.edu/checker-framework/eclipse/>). The website contains instructions for installing and using the plugin.

Eclipse plug-in for Type Annotations A prototype version of Type Annotations support for Eclipse is available from the Eclipse project. The goal is to enable full support for writing type annotations outside of comments. You do not need this to run the Checker Framework, whether or not you write your type annotations in comments.

(Update: this apparently needs a username and password, so it may not be publicly available.) Use the following information to check out the CVS repository:

Host: dev.eclipse.org
Repository path: /cvsroot/eclipse
Module name: org.eclipse.jdt.core
Branch: JSR_308

23.7 tIDE

tIDE, an open-source Java IDE, supports the Checker Framework. See its documentation at <http://tide.olympo-network.com/>.

23.8 Type inference tools

23.8.1 Varieties of type inference

There are two different tasks that are commonly called “type inference”.

1. Type inference during type checking (Section 19.1.2): During type checking, if certain variables have no type qualifier, the type-checker determines whether there is some type qualifier that would permit the program to type check. If so, the type checker uses that type qualifier, but never tells the programmer what it was. Each time the type checker runs, it re-infers the type qualifier for that variable. If no type qualifier exists that permits the program to type-check, the type-checker issues a type warning.

This variety of type inference is built into the Checker Framework. Every checker can take advantage of it at no extra effort. However, it only works within a method, not across method boundaries.

Advantages of this variety of type inference include:

- If the type qualifier is obvious to the programmer, then omitting it can reduce annotation clutter in the program.
 - The type inference can take advantage of only the code currently being compiled, rather than having to be correct for all possible calls. Additionally, if the code changes, then there is no old annotation to update.
2. Type inference to annotate a program (Section 23.8.2): As a separate step before type checking, a type inference tool takes the program as input, and outputs a set of type qualifiers that would type-check. These qualifiers are inserted into the source code or the class file. They can be viewed and adjusted by the programmer, and can be used by tools such as the type checker.

This variety of type inference must be provided by a separate tool. It is not built into the Checker Framework.

Advantages of this variety of type inference include:

- The program contains documentation in the form of type qualifiers, which can aid programmer understanding.
- Error messages may be more comprehensible. With type inference during type checking, error messages can be obscure, because the compiler has already inferred (possibly incorrect) types for a number of variables.
- A minor advantage is speed: type-checking can be modular, which can be faster than re-doing type inference every time the program is type-checked.

Advantages of both varieties of inference include:

- Less work for the programmer.
- The tool chooses the most general type, whereas a programmer might accidentally write a more specific, less generally-useful annotation.

Each variety of type inference has its place. When using the Checker Framework, type inference during type checking is performed only *within* a method (Section 19.1.2). Every method signature (arguments and return values) and field must be explicitly annotated, either by the programmer or by a separate type checking tool (Section 23.8.2). This choice reduces programmer effort (typically, a programmer does not have to write any qualifiers inside the body of a method) while still retaining modular checking and documentation benefits.

23.8.2 Type inference to annotate a program

This section lists tools that take a program and output a set of annotations for it.

Section 3.3.5 lists several tools that infer annotations for the Nullness Checker.

Section 6.2.2 lists a tool that infers annotations for the Javari Checker, which detects mutation errors.

Chapter 24

Frequently Asked Questions (FAQs)

These are some common questions about the Checker Framework and about pluggable type-checking in general. Feel free to suggest improvements to the answers, or other questions to include here.

There is a separate FAQ for the type annotations syntax (<http://types.cs.washington.edu/jsr308/jsr308-faq.html>).

Contents:

- 24.1: Are type annotations easy to read and write?
- 24.2: Will my code become cluttered with type annotations?
- 24.3: I don't make type errors, so would pluggable type checking help me?
- 24.4: What should I do if a checker issues a warning about my code?
- 24.5: Can a pluggable type-checker give an absolute guarantee of correctness?
- 24.6: How do I make compilation succeed even if a checker issues errors?
- 24.7: Will using the Checker Framework slow down my program? Will it slow down the compiler?
- 24.8: How can I do run-time monitoring of properties that were not statically checked?
- 24.9: How do I get started annotating an existing program?
- 24.10: How do I shorten the command line when invoking a checker?
- 24.11: When should I use type qualifiers, and when should I use subclasses?
- 24.12: How do I create a new checker?
- 24.13: Why is there no declarative syntax for writing type rules?
- 24.14: Why not just use a bug detector (like FindBugs)?
- 24.15: How does pluggable type-checking compare with JML?
- 24.16: Why shouldn't a qualifier apply to both types and declarations?
- 24.17: What is the meaning of array annotations such as `@NonNull Object @Nullable []`?
- 24.18: Why are the type parameters to `List` and `Map` annotated as `@NonNull`?
- 24.19: Is the Checker Framework an official part of Java?

24.1 Are type annotations easy to read and write?

The paper “Practical pluggable types for Java” [PAC⁺08] discusses case studies in which programmers found type annotations to be natural to read and write. The code continued to feel like Java, and the type-checking errors were easy to comprehend and often led to real bugs.

You don't have to take our word for it, though. You can try the Checker Framework for yourself.

The difficulty of adding and verifying annotations depends on your program. If your program is well-designed and -documented, then skimming the existing documentation and writing type annotations is extremely easy. Otherwise, you may find yourself spending a lot of time trying to understand, reverse-engineer, or fix bugs in your program, and

then just a moment writing a type annotation that describes what you discovered. This process inevitably improves your code. You must decide whether it is a good use of your time. For code that is not causing trouble now and is unlikely to do so in the future (the code is bug-free, and you do not anticipate changing it or using it in new contexts), then the effort of writing type annotations for it may not be justified.

24.2 Will my code become cluttered with type annotations?

As with any language feature, it is possible to write ugly code that over-uses annotations. However, in normal use, very few annotations need to be written. Figure 1 of the paper Practical pluggable types for Java [PAC⁺08] reports data for over 350,000 lines of type-annotated code:

- 1 annotation per 62 lines for nullness annotations (@NonNull, @Nullable, etc.)
- 1 annotation per 1736 lines for interning annotations (@Interned)
- 1 annotation per 27 lines for immutability annotations (IGJ type system)

These numbers are for annotating existing code. New code that is written with the type annotation system in mind is cleaner and more correct, so it requires even fewer annotations.

Each annotation that a programmer writes replaces a sentence or phrase of English descriptive text that would otherwise have been written in the Javadoc. So, use of annotations actually reduces the overall size of the documentation, at the same time as making it less ambiguous and machine-processable.

In summary: annotations do not clutter code; they are used much less frequently than generic types, which Java programmers find acceptable; and they reduce the overall volume of documentation that a codebase needs.

24.3 I don't make type errors, so would pluggable type checking help me?

Occasionally, a developer says that he makes no errors that typechecking could catch, or that any such errors are unimportant because they have low impact and are easy to fix. When I investigate the claim, I invariably find that the developer is mistaken.

Very frequently, the developer has underestimated what typechecking can discover. Not every type error leads to an exception being thrown; and even if an exception is thrown, it may not seem related to classical types. Remember that a type system can discover null pointer dereferences, incorrect side effects, security errors such as information leakage or SQL injection, partially-initialized data, wrong units of measurement, and many other errors. Even where type-checking does not discover a problem directly, it can indicate code with bad smells, thus revealing problems, improving documentation, and making future maintenance easier.

There are other ways to discover errors, including extensive testing and debugging. You should continue to use these. But type-checking is a good complement to these. Type-checking is more effective for some problems, and less effective for other problems. It can reduce (but not eliminate) the time and effort that you spend on other approaches. There are many important errors that type checking and other automated approaches cannot find; pluggable typechecking gives you more time to focus on those.

24.4 What should I do if a checker issues a warning about my code?

For a discussion of this issue, see Section 2.4.7.

24.5 Can a pluggable type-checker give an absolute guarantee of correctness?

Each checker looks for certain errors. You can use multiple checkers, but even then your program might still contain other kinds of errors.

If you run a pluggable checker on only part of the code of a program, then you do not get a guarantee that all parts of the program satisfy the type system (that is, are error-free). An example is a framework that clients are intended to extend. In this case, you should recommend that clients run the pluggable checker. There is no way to force users to do so, so you may want to retain dynamic checks or use other mechanisms to detect errors.

There are other circumstances in which a static type-checker may fail to detect a possible type error. In Java, these include arrays, casts, raw types, reflection, separate compilation (bytecodes from unverified sources), native code, etc. (For details, see section 2.3.) Java uses dynamic checks for most of these, so that the type error cannot cause a security vulnerability or a crash. The pluggable type-checkers inherit many (not all) of these weaknesses of Java type-checking, but do not currently have built-in dynamic checkers. Writing dynamic checkers would be an interesting and valuable project.

Even if a tool such as a pluggable checker cannot give an ironclad guarantee of correctness, it is still useful. It can find errors, excluding certain types of possible problems (e.g., restricting the possible class of problems), and increasing confidence in a piece of software.

24.6 How do I make compilation succeed even if a checker issues errors?

Section 2.2 describes the `-Awarns` command-line option that turns checker errors into warnings, so type-checking errors will not cause `javac` to exit with a failure status.

24.7 Will using the Checker Framework slow down my program? Will it slow down the compiler?

Using the Checker Framework has no impact on the execution of your program: the Type Annotations compiler emits the identical bytecodes as the Java 8 compiler and so there is no run-time effect. Because there is no run-time representation of type qualifiers, there is no way to use reflection to query the qualifier on a given object, though you can use reflection to examine a class/method/field declaration.

Using the Checker Framework does increase compilation time. In theory it should only add a few percent overhead, but our current implementation can double the compilation time — or more, if you run many pluggable type-checkers at once. This is especially true if you run pluggable type-checking on every file (as we recommend) instead of just on the ones that have recently changed. Nonetheless, compilation with pluggable type-checking still feels like compilation, and you can do it as part of your normal development process.

24.8 How can I do run-time monitoring of properties that were not statically checked?

Some properties are not checked statically (see Chapter 20 for reasons that code might not be statically checked). In such cases, it would be desirable to check the property dynamically, at run time. Currently, the Checker Framework has no support for adding code to perform run-time checking.

Adding such support would be an interesting and valuable project. An example would be an option that causes the Checker Framework to automatically insert a run-time check anywhere that static checking is suppressed. If you are able to add run-time verification functionality, we would gladly welcome it as a contribution to the Checker Framework.

Some checkers have library methods that you can explicitly insert in your source code. Examples include the Nullness Checker's `NullnessUtils.castNonNull` method (see Section 3.4.1) and the Regex Checker's `RegexUtil` class (see Section 11.2.4). But, it would be better to have more general support that does not require the user to explicitly insert method calls.

24.9 How do I get started annotating an existing program?

See Section 2.4.1.

24.10 How do I shorten the command line when invoking a checker?

The compile options to `javac` can be a pain to type; for example, `javac -processor checkers.nullness.NullnessChecker . . .`. See Section 2.2.2 for a way to avoid the need for the `-processor` command-line option.

24.11 When should I use type qualifiers, and when should I use subclasses?

In brief, use subtypes when you can, and use type qualifiers when you cannot use subtypes. For more details, see section 2.4.6.

24.12 How do I create a new checker?

In addition to using the checkers that are distributed with the Checker Framework, you can write your own checker to check specific properties that you care about. Thus, you can find and prevent the bugs that are most important to you.

Chapter 22 gives complete details regarding how to write a checker. It also suggests places to look for more help, such as the Checker Framework API documentation (Javadoc) and the source code of the distributed checkers.

To whet your interest and demonstrate how easy it is to get started, here is an example of a complete, useful type checker.

```
@TypeQualifier
@SubtypeOf(Unqualified.class)
@Target({ElementType.TYPE_PARAMETER, ElementType.TYPE_USE})
public @interface Encrypted { }
```

Section 15.2 explains this checker and tells you how to run it.

24.13 Why is there no declarative syntax for writing type rules?

A type system implementer can declaratively specify the type qualifier hierarchy (Section 22.3.1) and the type introduction rules (Section 22.4.1). However, the Checker Framework uses a procedural syntax for specifying type-checking rules (Section 22.5). A declarative syntax might be more concise, more readable, and more verifiable than a procedural syntax.

We have not found the procedural syntax to be the most important impediment to writing a checker.

Previous attempts to devise a declarative syntax for realistic type systems have failed; see a technical paper [PAC⁺08] for a discussion. When an adequate syntax exists, then the Checker Framework can be extended to support it.

24.14 Why not just use a bug detector (like FindBugs)?

Pluggable type-checking finds more bugs than a bug detector does, for any given variety of bug.

A bug detector like FindBugs [HP04, HSP05], Jlint [Art01], or PMD [Cop05] aims to find *some* of the most obvious bugs in your program. It uses a lightweight analysis, then uses heuristics to discard some of its warnings. Thus, even if the tool prints no warnings, your code might still have errors — maybe the analysis was too weak to find them, or the tool's heuristics classified the warnings as likely false positives and discarded them.

A type checker aims to find *all* the bugs (of certain varieties). It requires you to write type qualifiers in your program, or to use a tool that infers types. Thus, it requires more work from the programmer, and in return it gives stronger guarantees.

Each tool is useful in different circumstances, depending on how important your code is and your desired level of confidence in your code. For more details on the comparison, see section 25.5. For a case study that compared the nullness analysis of FindBugs, Jlint, PMD, and the Checker Framework, see section 6 of the paper “Practical pluggable types for Java” [PAC⁺08].

24.15 How does pluggable type-checking compare with JML?

JML, the Java Modeling Language [LBR06], is a language for writing formal specifications. JML aims to be more expressive than pluggable type-checking. JML is not as practical as pluggable type-checking.

A programmer can write a JML specification that describes arbitrary facts about program behavior. Then, the programmer can use formal reasoning or a theorem-proving tool to verify that the code meets the specification. Run-time checking is also possible. By contrast, pluggable type-checking can express a more limited set of properties about your program.

The JML toolset is less mature. For instance, if your code uses generics or other features of Java 5, then you cannot use JML. However, JML has a run-time checker, which the Checker Framework currently lacks.

24.16 Why shouldn't a qualifier apply to both types and declarations?

It is bad style for an annotation to apply to both types and declarations. In other words, every annotation should have a `@Target` meta-annotation, and the `@Target` meta-annotation should list either only declaration locations or only type annotations. (It's OK for an annotation to target both `ElementType.TYPE_PARAMETER` and `ElementType.TYPE_USE`, but no other declaration location along with `ElementType.TYPE_USE`.)

Sometimes, it may seem tempting for an annotation to apply to both type uses and (say) method declarations. Here is a hypothetical example:

“Each `Widget` type may have a `@Version` annotation. I wish to prove that versions of widgets don't get assigned to incompatible variables, and that older code does not call newer code (to avoid problems when backporting).

A `@Version` annotation could be written like so:

```
@Version("2.0") Widget createWidget(String value) { ... }
```

`@Version("2.0")` on the method could mean that the `createWidget` method only appears in the 2.0 version. `@Version("2.0")` on the return type could mean that the returned `Widget` should only be used by code that uses the 2.0 API of `Widget`. It should be possible to specify these independently, such as a 2.0 method that returns a value that allows the 1.0 API method invocations.”

Both of these are type properties and should be specified with type annotations. No method annotation is necessary or desirable. The best way to require that the receiver has a certain property is to use a type annotation on the receiver of the method. (Slightly more formally, the property being checked is compatibility between the annotation on the type of the formal parameter receiver and the annotation on the type of the actual receiver.)

Another example of a type-and-declaration annotation that represents poor design is JCIP's `@GuardedBy` annotation [GPB⁺06]. As discussed in Section 7.1.3, it means two different things when applied to a field or a method. To reduce confusion and increase expressiveness, the Lock Checker (see Chapter 7) uses the `@Holding` annotation for one of these meanings, rather than overloading `@GuardedBy` with two distinct meanings.

24.17 What is the meaning of array annotations such as `@NonNull Object @Nullable []`?

You should parse this as: `(@NonNull Object) (@Nullable [])`. Each annotation precedes the component of the type that it qualifies.

Thus, `@NonNull Object @Nullable []` is a possibly-null array of non-null objects. Note that the first token in the type, “`@NonNull`”, applies to the element type `Object`, not to the array type as a whole. The annotation `@Nullable` applies to the array `[]`.

Similarly, `@Nullable Object @NonNull []` is a non-null array of possibly-null objects.

24.18 Why are the type parameters to `List` and `Map` annotated as `@NonNull`?

The annotation on `java.util.Collection` only allows non-null elements:

```
public interface Collection<E extends @NonNull Object> {
    ...
}
```

Thus, you will get a type error if you write code like `Collection<@Nullable Object>`. A nullable type parameter is also forbidden for certain other collections, including `AbstractCollection`, `List`, `Map`, and `Queue`.

The `extends @NonNull Object` bound is a direct consequence of the design of the collections classes; it merely formalizes the Javadoc specification. The Javadoc for `Collection` states:

Some list implementations have restrictions on the elements that they may contain. For example, some implementations prohibit null elements, ...

Here are some consequences of the requirement to detect all nullness errors at compile time. If even one subclass of a given collection class may prohibit null, then the collection class and all its subclasses must prohibit null. Conversely, if a collection class is specified to accept null, then all its subclasses must honor that specification.

The Checker Framework’s annotations make apparent a flaw in the JDK design, and helps you to avoid problems that might be caused by that flaw.

Justification from type theory Suppose `B` is a subtype of `A`. Then an overriding method in `B` must have a stronger (or equal) signature than the overridden method in `A`. In a stronger signature, the formal parameter types may be supertypes, and the return type may be a subtype. Here are examples:

```
class A          { @NonNull Object Number m1( @NonNull Object arg) { ... } }
class B extends A { @Nullable Object Number m1( @NonNull Object arg) { ... } } // error!
class C extends A { @NonNull Object Number m1( @Nullable Object arg) { ... } } // OK
class D          { @Nullable Object Number m2( @Nullable Object arg) { ... } }
class E extends D { @NonNull Object Number m2( @Nullable Object arg) { ... } } // OK
class F extends D { @Nullable Object Number m2( @NonNull Object arg) { ... } } // error!
```

According to these rules, since some subclasses of `Collection` do not permit nulls, then `Collection` cannot either:

```
// does not permit null elements
class PriorityQueue<E> implements Collection<E> {
    boolean add(E);
    ...
}
// must not permit null elements, or PriorityQueue would not be a subtype of Collection
```

```
interface Collection<E> {
    boolean add(E);
    ...
}
```

Justification from checker behavior Suppose that you changed the bound in the `Collection` declaration to extends `Nullable Object`. Then, the checker would issue no warning for this method:

```
static void addNull(Collection l) {
    l.add(null);
}
```

However, calling this method *can* result in a null pointer exception, for instance caused by the following code:

```
addNull(new PriorityQueue());
```

Therefore, the bound must remain as extends `NonNull Object`.

By contrast, this code is OK because `ArrayList` is documented to support null elements:

```
static void addNull(ArrayList l) {
    l.add(null);
}
```

Therefore, the upper bound in `ArrayList` is extends `Nullable Object`. Any subclass of `ArrayList` must also support null elements.

Suppressing warnings Suppose your program has a list variable, and you know that any list referenced by that variable will definitely support null. Then, you can suppress the warning:

```
@SuppressWarnings("nullness:generic.argument")
static void addNull(List l) {
    l.add(null);
}
```

You need to use `@SuppressWarnings("nullness:generic.argument")` whenever you use a collection that may contain null elements in contradiction to its documentation. Fortunately, such uses are relatively rare.

For more details on suppressing nullness warnings, see Section 3.4.

24.19 Is the Checker Framework an official part of Java?

The Checker Framework is not an official part of Java, though it relies on type annotations, which are part of Java 8. See the Type Annotations (JSR 308) FAQ for more details.

Chapter 25

Troubleshooting and getting help

Please read the entire manual, including this chapter and the FAQ (Chapter 24), because the manual might already answer your question. If not, you can use the mailing list, `checker-framework-discuss@googlegroups.com`, to ask other users for help. For archives and to subscribe, see <http://groups.google.com/group/checker-framework-discuss>. To report bugs, use the issue tracker at <http://code.google.com/p/checker-framework/issues/list>. If you want to help out, you can choose a bug and fix it, or select a project from the ideas list at <http://code.google.com/p/checker-framework/wiki/Ideas>.

25.1 Common problems and solutions

- To verify that you are using the compiler you think you are, you can add `-version` to the command line. For instance, instead of running `javac -g MyFile.java`, you can run `javac -version -g MyFile.java`. Then, `javac` will print out its version number in addition to doing its normal processing.

25.1.1 Unable to run the checker, or checker crashes

If you are unable to run the checker, or if the checker or the compiler crashes, then the problem may be a problem with your environment. This section describes some possible problems and solutions.

- If you get the error

```
com.sun.tools.javac.code.Symbol$CompletionFailure: class file for com.sun.source.tree.Tree not found
```

then you are using the source installation and file `tools.jar` is not on your classpath. See the installation instructions (Section 1.2).
- If you get an error such as

```
package checkers.nullnessquals does not exist
```

despite no apparent use of `import checkers.nullnessquals.*;` in the source code, then perhaps `jsr308_imports` is set as a Java system property, a shell environment variable, or a command-line option (see Section 20.3.2). You can solve this by unsetting the variable/option, or by ensuring that the `checkers.jar` file is on your classpath. If the error is

```
package 'checkers.nullnessquals does not exist
```

(note the extra apostrophe!), then you have probably misused quoting when supplying the `jsr308_imports` environment variable.
- If you get an error like the following when using the Ant task (Section 23.2),

```
...\build.xml:59: Error running ${env.CHECKERS}\binary\javac.bat compiler
```

then the problem may be that you have not set the `CHECKERS` environment variable, as described in Section 23.1.2. Or, maybe you made it a user variable instead of a system variable.

- If you get one of these errors:

The hierarchy of the type `ClassName` is inconsistent

The type `com.sun.source.util.AbstractTypeProcessor` cannot be resolved.

It is indirectly referenced from required ".class files",
then you are missing `jsr308-all.jar` from your classpath.

- If you get the error

```
java.lang.ArrayStoreException: sun.reflect.annotation.TypeNotPresentExceptionProxy
```

then an annotation is not present at run time that was present at compile time. For example, maybe when you compiled the code, the `@Nullable` annotation was available, but it was not available at run time. You can use JDK 8 at run time, or compile with a Java 7 compiler that will ignore the annotations in comments.

- A “class file not found” error may be due to a JDK version mismatch. For instance, you might be using JDK 7, but you get an error that refers to a class that was in a previous version of the JDK but has subsequently been removed, such as:

```
class file for java.io.File$LazyInitialization not found
class file for java.util.Hashtable$EmptyIterator not found
java.lang.NoClassDefFoundError: java/util/Hashtable$EmptyEnumerator
```

Or, you might be using JDK 6, but you get an error that refers to a class that has been introduced in a newer version of the JDK, such as:

```
class file for java.util.Vector$Itr not found
```

This problem occurs when your classpath contains code that was compiled with one version of the JDK and refers to its implementation details, but your classpath does not contain that version of the JDK itself.

You can solve the problem by re-generating `jdk/jdk.jar` and `binary/jdk.jar`. You can do this by running

```
cd checkers
ant jdk.jar bindist
```

- A `NoSuchFieldError` such as this:

```
java.lang.NoSuchFieldError: NATIVE_HEADER_OUTPUT
```

Field `NATIVE_HEADER_OUTPUT` was added in JDK 8. The error message suggests that you’re not executing with the right bootclasspath: some classes were compiled with the JDK 8 version and expect the field, but you’re executing the compiler on a JDK without the field.

One possibility is that you are not running the Type Annotations compiler — use `javac -version` to check this, then use the right one. (Maybe the Type Annotations `javac` is at the end rather than the beginning of your path.)

If you are using Ant, then one possibility is that the `javac` compiler is using the same JDK as Ant is using. You can correct this by being sure to use `fork="yes"` (see Section 23.2) and/or setting the `build.compiler` property to `extJavac`.

If you are building from source, you might need to rebuild the Annotation File Utilities before recompiling or using the Checker Framework.

- If you get an error that contains lines like these:

```
Caused by: java.util.zip.ZipException: error in opening zip file
    at java.util.zip.ZipFile.open(Native Method)
    at java.util.zip.ZipFile.<init>(ZipFile.java:131)
```

then one possibility is that you have installed the Checker Framework in a directory that contains special characters that Java’s `ZipFile` implementation cannot handle. For instance, if the directory name contains “+”, then Java 1.6 throws a `ZipException`, and Java 1.7 throws a `FileNotFoundException` and prints out the directory name with “+” replaced by blanks.

25.1.2 Unexpected type-checking results

This section describes possible problems that can lead the type-checker to give unexpected results.

- If the Checker Framework is unable to verify a property that you know is true, then it is helpful to formulate an argument about why the property is true. Recall that the Checker Framework does modular verification, one procedure at a time; it observes the specifications, but not the implementations, of other methods. If any aspects of your argument are not expressed as annotations, then you may need to write more annotations. If any aspects of your argument are not expressible as annotations, then you may need to extend the type-checker.
- If a checker seems to be ignoring the annotation on a method, then it is possible that the checker is reading the method's signature from its `.class` file, but the `.class` file was not created by the JSR 308 compiler. You can check whether the annotations actually appear in the `.class` file by using the `javap` tool. If the annotations do not appear in the `.class` file, here are two ways to solve the problem:
 - Re-compile the method's class with the Type Annotations compiler. This will ensure that the type annotations are written to the class file, even if no type-checking happens during that execution.
 - Pass the method's file explicitly on the command line when type-checking, so that the compiler reads its source code instead of its `.class` file.

- If the compiler reports that it cannot find a method from the JDK or another external library, then maybe the stub/skeleton file for that class is incomplete. You can edit it to add the missing method. The libraries appear, for example, at `checkers/jdk/nullness/src/` for the Nullness checker.

The error might take one of these forms:

```
method sleep in class Thread cannot be applied to given types
cannot find symbol: constructor StringBuffer(StringBuffer)
```

25.1.3 Unable to build the checker, or to run programs

An error like this

```
Unsupported major.minor version 51.0
```

means that you have compiled some files into the Java 7 format (version 51.0), but you are trying to run them with Java 6. Run `java -version` to determine the version of Java you are using and use a newer version, and/or use the `-target` command-line option to `javac` to specify the version of the class files that are created, such as `javac -target 6`

25.2 How to report problems (bug reporting)

If you have a problem with any checker, or with the Checker Framework, please file a bug at <http://code.google.com/p/checker-framework/issues/list>. (First, check whether there is an existing bug report for that issue.)

Alternately (especially if your communication is not a bug report), you can send mail to `checker-framework-dev@googlegroups.com`. We welcome suggestions, annotated libraries, bug fixes, new features, new checker plugins, and other improvements.

Please ensure that your bug report is clear and that it is complete. Otherwise, we may be unable to understand it or to reproduce it, either of which would prevent us from fixing the bug. Your bug report will be most helpful if you:

- Add `-version -verbose -AprintErrorStack -printAllQualifiers` to the `javac` options. This causes the compiler to output debugging information, including its version number.
- Indicate exactly what you did. Don't skip any steps, and don't merely describe your actions in words. Show the exact commands by attaching a file or using cut-and-paste from your command shell;
- Include all files that are necessary to reproduce the problem. This includes every file that is used by any of the commands you reported, and possibly other files as well.

- Indicate exactly what the result was by attaching a file or using cut-and-paste from your command shell (don't merely describe it in words). Also indicate what you expected the result to be — remember, a bug is a difference between desired and actual outcomes.

A particularly useful format for a test case is as a new file, or a diff to an existing file, for the existing Checker Framework test suite. For instance, for the Nullness Checker, see directory `checker-framework/checkers/tests/nullness/`. But, please report your bug even if you do not report it in this format.

25.3 Building from source

The Checker Framework release (Section 1.2) contains everything that most users need, both to use the distributed checkers and to write your own checkers. This section describes how to compile its binaries from source. You will be using the latest development version of the Checker Framework, rather than an official release.

25.3.1 Obtain the source

Obtain the latest source code from the version control repository:

```
export JSR308=$HOME/jsr308
mkdir -p $JSR308
cd $JSR308
hg clone https://code.google.com/p/jsr308-langtools/ jsr308-langtools
hg clone https://code.google.com/p/checker-framework/ checker-framework
hg clone https://code.google.com/p/annotation-tools/ annotation-tools
```

(Alternately, you could use the version of the source code that is packaged in the Checker Framework release.)

25.3.2 Build the Type Annotations compiler

1. Set the `JAVA_HOME` environment variable to the location of your JDK 7 installation (not the JRE installation, and not JDK 6). This needs to be an Oracle JDK. (The `JAVA_HOME` environment variable might already be set, because it is needed for Ant to work.)

In the bash shell, the following command *sometimes* works (it might not because `java` might be the version in the JDK or in the JRE):

```
export JAVA_HOME=${JAVA_HOME:-$(dirname $(dirname $(dirname $(readlink -f $(/usr/bin/which java))))))
```

2. Compile the Type Annotations javac compiler and the javap tool:

```
cd $JSR308/jsr308-langtools/make
ant clean build-javac build-javap
```

3. Add the `jsr308-langtools/dist/bin` directory to the front of your `PATH` environment variable. Example command:

```
export PATH=$JSR308/jsr308-langtools/dist/bin:${PATH}
```

25.3.3 Build the Annotation File Utilities

This is simply done by:

```
cd $JSR308/annotation-tools
ant
```

You do not need to add the Annotation File Utilities to the path, as the Checker Framework build finds it using relative paths.

25.3.4 Build the Checker Framework

1. Run `ant` to create `checkers.jar`:

```
cd $JSR308/checker-framework/checkers
ant
```

2. Add `tools.jar` and `checkers.jar` to your classpath. (If you do not do this, you will have to supply the `-cp` option whenever you run `javac` and use a checker plugin.) Example command:

```
export CLASSPATH=${CLASSPATH}:$JAVA_HOME/lib/tools.jar:$JSR308/checker-framework/checkers/checkers.jar
```

3. Test that everything works:

- Run `ant all-tests` in the `checkers` directory:

```
cd $JSR308/checker-framework/checkers
ant all-tests
```
- Run the Nullness checker examples (see Section 3.8, page 29).

25.3.5 Build the Checker Framework manual (this document)

1. To build the manual you will need `plume-bib` (<http://code.google.com/p/plume-bib/>) and `HEVEA` (<http://hevea.inria.fr/>) installed.
2. Run `make` in the `checkers/manual` directory to build both the PDF and HTML versions of the manual.

25.4 Learning more

The technical paper “Practical pluggable types for Java” [PAC⁺08] (<http://www.cs.washington.edu/homes/mernst/pubs/pluggable-checkers-issta2008.pdf>) gives more technical detail about many aspects of the Checker Framework and its implementation. The technical paper also describes case studies in which each of the checkers found previously-unknown errors in real software.

The paper “Building and using pluggable type-checkers” [DDE⁺11] (<http://www.cs.washington.edu/homes/mernst/pubs/pluggable-checkers-icse2011.pdf>) discusses further experience with the Checker Framework, increasing the number of lines of verified code to 3 million.

In addition to these papers that discuss use the Checker Framework directly, other academic papers use the Checker Framework in their implementation or evaluation. Most educational use of the Checker Framework is never published, and most commercial use of the Checker Framework is never discussed publicly.

25.5 Comparison to other tools

A pluggable type-checker, such as those created by the Checker Framework, aims to help you prevent or detect all errors of a given variety. An alternate approach is to use a bug detector such as FindBugs, JLint, or PMD.

A pluggable type-checker differs from a bug detector in several ways:

- A type-checker aims to find *all* errors. Thus, it can verify the *absence* of errors: if the type checker says there are no null pointer errors in your code, then there are none. (This guarantee only holds for the code it checks, of course; see Section 2.3.)
A bug detector aims to find *some* of the most obvious errors. Even if it reports no errors, then there may still be errors in your code.
Both types of tools may issue false positive warnings; see Section 20.2.
- A type-checker requires you to annotate your code with type qualifiers, or to run an inference tool that does so for you. A bug detector may not require annotations. This means that it may be easier to get started running a bug detector.

- A type-checker may use a more sophisticated and complete analysis. A bug detector typically does a more lightweight analysis, coupled with heuristics to suppress false positives. As one example, a type-checker can take advantage of annotations on generic type parameters, such as `List<@NonNull String>`, permitting it to be much more precise for code that uses generics.

A case study [PAC⁺08, §6] compared the Checker Framework's nullness checker with those of FindBugs, Jlint, and PMD. The case study was on a well-tested program in daily use. The Checker Framework tool found 8 nullness errors (that is, null pointer dereferences). None of the other tools found any errors.

Also see the JSR 308 [Ern08] documentation for a detailed discussion of related work.

25.6 Credits and changelog

The key developers of the Checker Framework are Mahmood Ali, Telmo Correa, Werner M. Dietl, Michael D. Ernst, and Matthew M. Papi. Many other developers have also contributed, for example by writing the checkers that are distributed with the Checker Framework. Many, many users to list have provided valuable feedback, for which we are grateful.

Differences from previous versions of the checkers and framework can be found in the `changelog-checkers.txt` file. This file is included in the Checker Framework distribution and is also available on the web at <http://types.cs.washington.edu/checker-framework/current/changelog-checkers.txt>.

Bibliography

- [Art01] Cyrille Artho. Finding faults in multi-threaded programs. Master's thesis, Swiss Federal Institute of Technology, March 15, 2001.
- [Cop05] Tom Copeland. *PMD Applied*. Centennial Books, November 2005.
- [Cro06] Jose Cronembold. JSR 198: A standard extension API for Integrated Development Environments. <http://jcp.org/en/jsr/detail?id=198>, May 8, 2006.
- [Dar06] Joe Darcy. JSR 269: Pluggable annotation processing API. <http://jcp.org/en/jsr/detail?id=269>, May 17, 2006. Public review version.
- [DDE⁺11] Werner Dietl, Stephanie Dietzel, Michael D. Ernst, Kıvanç Muşlu, and Todd Schiller. Building and using pluggable type-checkers. In *ICSE'11, Proceedings of the 33rd International Conference on Software Engineering*, pages 681–690, Waikiki, Hawaii, USA, May 25–27, 2011.
- [Ern08] Michael D. Ernst. Type Annotations specification (JSR 308). <http://types.cs.washington.edu/jsr308/>, September 12, 2008.
- [Eva96] David Evans. Static detection of dynamic memory errors. In *PLDI 1996, Proceedings of the SIGPLAN '96 Conference on Programming Language Design and Implementation*, pages 44–53, Philadelphia, PA, USA, May 21–24, 1996.
- [FL03] Manuel Fähndrich and K. Rustan M. Leino. Declaring and checking non-null types in an object-oriented language. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 2003)*, pages 302–312, Anaheim, CA, USA, November 6–8, 2003.
- [FLL⁺02] Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In *PLDI 2002, Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*, pages 234–245, Berlin, Germany, June 17–19, 2002.
- [Goe06] Brian Goetz. The pseudo-typedef antipattern: Extension is not type definition. <http://www.ibm.com/developerworks/java/library/j-jtp02216/>, February 21, 2006.
- [GPB⁺06] Brian Goetz, Tim Peierls, Joshua Bloch, Joseph Bowbeer, David Holmes, and Doug Lea. *Java Concurrency in Practice*. Addison-Wesley, 2006.
- [HP04] David Hovemeyer and William Pugh. Finding bugs is easy. In *Companion to Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 2004)*, pages 132–136, Vancouver, BC, Canada, October 26–28, 2004.
- [HSP05] David Hovemeyer, Jaime Spacco, and William Pugh. Evaluating and tuning a static analysis to find null pointer bugs. In *ACM SIGPLAN/SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE 2005)*, pages 13–19, Lisbon, Portugal, September 5–6, 2005.

- [LBR06] Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: A behavioral interface specification language for Java. *ACM SIGSOFT Software Engineering Notes*, 31(3), March 2006.
- [PAC⁺08] Matthew M. Papi, Mahmood Ali, Telmo Luis Correa Jr., Jeff H. Perkins, and Michael D. Ernst. Practical pluggable types for Java. In *ISSTA 2008, Proceedings of the 2008 International Symposium on Software Testing and Analysis*, pages 201–212, Seattle, WA, USA, July 22–24, 2008.
- [QTE08] Jaime Quinonez, Matthew S. Tschantz, and Michael D. Ernst. Inference of reference immutability. In *ECOOP 2008 — Object-Oriented Programming, 22nd European Conference*, pages 616–641, Paphos, Cyprus, July 9–11, 2008.
- [TE05] Matthew S. Tschantz and Michael D. Ernst. Javari: Adding reference immutability to Java. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 2005)*, pages 211–230, San Diego, CA, USA, October 18–20, 2005.
- [ZPA⁺07] Yoav Zibin, Alex Potanin, Mahmood Ali, Shay Artzi, Adam Kiezun, and Michael D. Ernst. Object and reference immutability using Java generics. In *ESEC/FSE 2007: Proceedings of the 11th European Software Engineering Conference and the 15th ACM SIGSOFT Symposium on the Foundations of Software Engineering*, pages 75–84, Dubrovnik, Croatia, September 5–7, 2007.