

Assignment 3 – Software Analysis

Extended Java Typechecking

Claudio Maggioni

1 Project selection

The assignment description requires to find a project with more than 1000 lines of code making significant use of arrays or strings.

Given these requirements, I decide to analyze the Apache Commons Text project in the GitHub repository `apache/commons-text`.

1.1 The Apache Commons Lang Project

The Apache Commons family of libraries is an Apache Software Foundation¹ sponsored collection of Java libraries designed to complement the standard libraries of Java. The Apache Commons Text project focuses on text manipulation, encoding and decoding of *Strings* and *CharSequence*-implementing classes in general.

All the source and test classes are contained within in the package `org.apache.commons.text` or in a sub-package of that package. For the sake of brevity, this prefix is omitted from now on when mentioning file paths and classes in the project.

I choose to analyze the project at the *git* commit `78fac0f157f74feb804140613e4ffec449070990` as it is the latest commit on the *master* branch at the time of writing.

To verify that the project satisfies the 1000 lines of code requirement, I run the *cloc* tool. Results are shown in table 1. Given the project has more than 29,000 lines of Java code, this requirement is satisfied.

Language	Files	Blank	Comment	Code
Java	194	5642	18704	26589
XML	16	205	425	1370
Text	6	194	0	667
Maven	1	23	24	536
YAML	6	39	110	160
Markdown	4	40	106	109
Velocity Template Language	1	21	31	87
CSV	1	0	0	5
Properties	2	2	28	5
Bourne Shell	1	0	2	2
Total	232	6166	19430	29530

Table 1: Output of the *cloc* tool for the Apache Commons Text project at tag `78fac0f1` (before refactoring is carried out).

¹<https://apache.org/>

2 Running the CheckerFramework Type Checker

The relevant source code to analyze has been copied to the directory *sources* in the assignment repository

usi-si-teaching/msde/2022-2023/software-analysis/maggioni/assignment-3

on *gitlab.com*. The Maven build specification for the project has been modified to run the CheckerFramework extended type checker (version 3.33.0) as an annotation processor to be ran on top of the Java compiler. Both source code and test code is checked with the tool for violations, which are reported with compilation warnings. To run the type checker simply run:

```
mvn clean compile
```

in a suitable environment (i.e. with JDK 1.8 or greater and Maven installed). To additionally run the Apache Commons Text test suite and enable `assert` assertions (later useful for CheckerFramework `@AssumeAssertion(index)` assertions) simply run:

```
env MAVEN_OPTS="-ea" mvn clean test
```

Apache Commons Text includes classes that have been deprecated. As changing the interface and behaviour of these classes would be useless, as alternatives to them exist in the library already, I choose to ignore them while refactoring by adding a `@SuppressWarnings` annotation in each of them. The state of the assignment repository after the deprecated classes are annotated and when the type checker was first ran successfully is pinned by the *git* tag *before-refactor*. A copy of the CheckerFramework relevant portion of the compilation output at that tag is stored in the file *before-refactor.txt*.

No CheckerFramework checkers other than the index checker is used in this analysis as the code in the project mainly manipulates strings and arrays and a significant number of warnings are generated even by using this checker only..

3 Refactoring

Warning type	Before refactoring	After refactoring
argument	254	241
array.access.unsafe.high	130	117
array.access.unsafe.high.constant	31	28
array.access.unsafe.high.range	22	22
array.access.unsafe.low	59	58
array.length.negative	3	3
cast.unsafe	2	2
override.return	12	12
Total	513	483

Table 2: Number of CheckerFramework Type Checker warnings by category before and after refactoring, ignoring deprecated classes.

Table 2 provides a summary on the extent of the refactoring performed in response to index checker warnings across the Apache Commons Text project. In total, 513 warnings are found before refactoring, with 30 of them later being extinguished by introducing annotations and assertions in the code in the following classes:

- AlphabetConverter
- StringSubstitutor
- similarity.LongestCommonSubsequence
- translate.AggregateTranslator
- translate.CharSequenceTranslator
- translate.CodePointTranslator
- translate.CsvTranslators
- translate.JavaUnicodeEscaper
- translate.SinglePassTranslator
- translate.UnicodeEscaper

The strategy I adopt to perform the refactoring is based on the compiler errors thrown on the original code. In every flagged statement I attempt to find the root cause of the warning and eliminate it with either extended type qualifier annotations or assertions when adding only the former fails.

Instead of using `@SuppressWarnings` annotations I choose to use `@AssumeAssertion`-annotated assertions as I aim to use the existing Commons Text test suite to aid in finding incorrectly-placed annotations. As mentioned before in the report, I run the test suite of the project by enabling assertions and I verify that all tests still pass and no *AssertionErrors* are thrown.

In total, the refactor consists in the placement of 16 extended type qualifiers and 14 assertions. A more detailed description of salient refactoring decisions taken to extinguish the warnings follows.

3.1 Class *AlphabetConverter*

```

387 for (int j = 0; j < encoded.length();) {
388     final int i = encoded.codePointAt(j);
389     final String s = codePointToString(i);
390
391     if (s.equals(originalToEncoded.get(i))) {
392         result.append(s);
393         j++; // because we do not encode in Unicode extended the
394             // length of each encoded char is 1
395     } else {
396         if (j + encodedLetterLength > encoded.length()) {
397             throw new UnsupportedOperationException("Unexpected end "
398                 + "of string while decoding " + encoded);
399         }
400         final String nextGroup = encoded.substring(j,
401             j + encodedLetterLength);

```

Here the `substring(...)` call at line 151 is flagged by CheckerFramework warning the start and end index may be negative and that the start index may be greater than the end index. As the attribute `encodedLetterLength` is positive according to the contract of the class constructor and `j` is only incremented in the for loop or by a factor of `encodedLetterLength`, the code is correct. After introducing a `@Positive` annotation on the declaration of `j` and an `assert encodedLength > 0` after line 395, CheckerFramework agrees with my judgement.

3.2 Class *StringSubstitutor*

```

910 /** [...]
911  * @throws IllegalArgumentException if variable is not found when its allowed to
912  *   ↪ throw exception
913  * @throws StringIndexOutOfBoundsException if {@code offset} is not in the
914  *   range {@code 0 <= offset <= source.length()}
915  * @throws StringIndexOutOfBoundsException if {@code length < 0}

```

```

915  * @throws StringIndexOutOfBoundsException if {@code offset + length >
    ↪ source.length()}
916  */
917  public String replace(final String source, final int offset, final int length) {
918      if (source == null) {
919          return null;
920      }
921      final TextStringBuilder buf = new TextStringBuilder(length).append(source,
    ↪ offset, length);
922      if (!substitute(buf, 0, length)) {
923          return source.substring(offset, offset + length);
924      }
925      return buf.toString();
926  }

```

The implementation of method `replace` is flagged by the extended type checker as the indices `offset` and `length` are not bound checked against the string `source`. As the unsafe behaviour of the method is documented in its *javadoc* with appropriate `@throws` clauses, I simply add this implied preconditions to the method's contract by using extended type qualifiers:

```

public String replace(final String source,
                    final @IndexOrHigh("#1") int offset,
                    final @NonNegative @LTLengthOf(value = "#1", offset = "#2 -
    ↪ 1") int length)

```

3.3 Class *translate.CharSequenceTranslator* and implementors

Apache Commons Text provides the aforementioned abstract class implementation as a template method of sorts for expressing text encoding and decoding algorithms. The class essentially provides facilities to scan UTF-16 code points sequentially, and delegating the translation of each code point to the implementation of the abstract method:

```

public abstract int translate(CharSequence input, int index, Writer writer) throws
    ↪ IOException;

```

CheckerFramework gives some warnings about some of the implementations of this method, highlighting that they assume the `input` *CharSequence* is non-empty and the `index` parameter is a valid index for the string.

Even if the method is public, I choose to interpret this hierarchy to mainly be a template method pattern, with high coupling between the algorithm in the abstract class and each abstract method implementation. Given this, I decide to restrict the method's precondition to highlight conditions already provided by the caller algorithm, namely the length and index constraints provided by CheckerFramework.

The new signature of the abstract method is this:

```

public abstract int translate(@MinLen(1) CharSequence input,
                            @NonNegative @LTLengthOf("#1") int index,
                            Writer writer) throws IOException;

```

As some methods have a more forgiving implementation, and a broader child method argument type from a more restrictive parent type does not break the rules of inheritance (thanks to contravariance), I choose to propagate the extended type annotations only when needed and avoid introducing additional preconditions to more tolerant implementations of the template method.

3.4 Class *translate.SinglePassTranslator* and implementors

SinglePassTranslator is one of the implementor classes of the aforementioned *CharSequenceTranslator* template method. However, the class is itself a template method pattern “for processing whole input in single pass”², i.e. essentially performing an abstraction inversion of the codepoint-by-codepoint algorithm in *CharSequenceTranslator* by implementing the encoding or decoding process in a single go.

The class immediately delegates the implementation of the translation algorithm to the abstract package-private method:

```
abstract void translateWhole(CharSequence input, Writer writer) throws IOException;
```

and requires callers of the public implementation of `translate` to call it with `index` equal to 0.

I simply propagate the non-empty extended type annotation on `input` (i.e. `@MinLen(1)`) on this new abstract method and implementors.

4 Conclusions

Did using the checker help you find any bugs or other questionable design and implementation choices?

∫ no bugs found, couple of design choices

How complex was it to apply the checker, and what benefits did you gain in return?

∫ not so complex, lots of false positives

Compare the checker’s trade-off between complexity of usage and analysis power to that of other software analysis techniques you’re familiar with (in particular, those used in previous assignments).

²According to the class *javadoc* description.