

UiO • **Department of Informatics**
University of Oslo

Automated Composition of Refactorings

Implementing and evaluating a search-based Extract and Move Method refactoring

Erlend Kristiansen
Master's Thesis Spring 2014



Abstract

Remove all todos (including list) before delivery/printing!!!
Can be done by removing “draft” from documentclass.

Write abstract

Contents

1	Introduction	15
1.1	Motivation and structure	15
1.2	What is refactoring?	15
1.2.1	Defining refactoring	16
1.2.2	The etymology of 'refactoring'	17
1.2.3	Reasons for refactoring	17
1.2.4	The magical number seven	19
1.2.5	Notable contributions to the refactoring literature	20
1.2.6	Tool support (for Java)	21
1.2.7	The relation to design patterns	22
1.2.8	The impact on software quality	23
1.2.9	Composite refactorings	24
1.2.10	Manual vs. automated refactorings	25
1.2.11	Correctness of refactorings	25
1.2.12	Refactoring and the importance of testing	27
1.3	The Project	28
1.3.1	Project description	28
1.3.2	The premises	28
1.3.3	The primitive refactorings	30
1.3.4	The Extract and Move Method refactoring	30
1.3.5	The Coupling Between Object Classes metric	31
1.3.6	Research questions	32
1.3.7	Methodology	34
1.3.8	Case study	35
1.3.9	Dogfooding	35
1.4	Related Work	35
1.4.1	Refactoring safety	35
1.4.2	Search-based refactoring	36
1.4.3	The compositional paradigm of refactoring	36
2	The search-based Extract and Move Method refactoring	37
2.1	The inputs to the refactoring	37
2.2	Defining a text selection	37
2.3	Where we look for text selections	38
2.3.1	Text selections are found in methods	38
2.3.2	The possible text selections of a method body	38
2.4	Disqualifying a selection	41

2.4.1	A call to a protected or package-private method . . .	41
2.4.2	A double class instance creation	42
2.4.3	Instantiation of non-static inner class	42
2.4.4	References to enclosing instances of the enclosing class	42
2.4.5	Inconsistent return statements	43
2.4.6	Ambiguous return values	44
2.4.7	Illegal statements	44
2.5	Disqualifying selections from the example	45
2.6	Finding a move target	45
2.7	Unfixes	46
2.8	Finding the example selections that have possible targets . .	47
2.9	Choosing the selection	49
2.10	Concluding the example	50
2.11	??	51
3	Refactorings in Eclipse JDT: Design and Shortcomings	53
3.1	Design	53
3.1.1	The Language Toolkit	53
3.2	Shortcomings	55
3.2.1	Absence of Generics in Eclipse Source Code	55
3.2.2	Composite Refactorings Will Not Appear as Atomic Actions	55
4	Composite Refactorings in Eclipse	57
4.1	A Simple Ad Hoc Model	57
4.1.1	A typical RefaktorChanger	57
4.2	The Extract and Move Method Refactoring	57
4.2.1	The Building Blocks	57
4.2.2	The ExtractAndMoveMethodChanger	58
4.2.3	The SearchBasedExtractAndMoveMethodChanger . .	61
4.2.4	The Prefix Class	62
4.2.5	The PrefixSet Class	62
4.2.6	Hacking the Refactoring Undo History	63
5	Analyzing Source Code in Eclipse	65
5.1	The Java model	65
5.2	The Abstract Syntax Tree	65
5.2.1	The AST in Eclipse	67
5.3	The ASTVisitor	69
5.4	Property collectors	71
5.4.1	The PrefixesCollector	71
5.4.2	The UnfixesCollector	72
5.4.3	The ContainsReturnStatementCollector	72
5.4.4	The LastStatementCollector	72
5.5	Checkers	72
5.5.1	The CallToProtectedOrPackagePrivateMethodChecker	72
5.5.2	The DoubleClassInstanceCreationChecker	73
5.5.3	The InstantiationOfNonStaticInnerClassChecker . . .	73

5.5.4	The EnclosingInstanceReferenceChecker	73
5.5.5	The ReturnStatementsChecker	73
5.5.6	The AmbiguousReturnValueChecker	74
5.5.7	The IllegalStatementsChecker	75
6	Technicalities	77
6.1	Source code organization	77
6.1.1	The no.uio.ifi.refaktor project	78
6.2	Continuous integration	80
6.2.1	Problems with AspectJ	80
7	Benchmarking	83
7.1	The benchmark setup	83
7.1.1	The ProjectImporter	83
7.2	Statistics	84
7.2.1	AspectJ	84
7.2.2	The Statistics class	84
7.2.3	Advices	85
7.3	Optimizations	85
7.3.1	Caching	85
7.3.2	Candidates stored as mementos	86
7.4	Handling failures	87
8	Case Studies	89
8.1	The tools	89
8.2	The SonarQube quality profile	90
8.3	The input	91
8.4	The experiment	91
8.5	Case 1: The Eclipse JDT UI project	92
8.5.1	Statistics	92
8.5.2	SonarQube analysis	93
8.5.3	Unit tests	95
8.5.4	Conclusions	96
8.6	Case 2: The <code>no.uio.ifi.refaktor</code> project	97
8.6.1	Statistics	97
8.6.2	SonarQube analysis	98
8.6.3	Unit tests	99
8.6.4	An informal experiment	99
8.6.5	Conclusions	99
8.7	Summary	100
9	Conclusions and Future Work	107
9.1	Conclusions	107
9.2	Future work	109

A Eclipse Bugs Submitted	111
A.1 Eclipse bug 420726: Code is broken when moving a method that is assigning to the parameter that is also the move destination	111
A.2 Eclipse bug 429416: IAE when moving method from anonymous class	111
A.3 Eclipse bug 429954: Extracting statement with reference to local type breaks code	112

List of Figures

1.1	The Extract Superclass refactoring, with united interfaces. . .	25
5.1	The Java model of Eclipse. “{ SomeElement }*” means “ SomeElement zero or more times“. For recursive structures, “...” is used.	67
5.2	Interrupted compilation process. (Full compilation process borrowed from <i>Compiler construction: principles and practice</i> by Kenneth C. Louden [Lou97].)	68
5.3	The abstract syntax tree for the expression (5 + 7) * 2 . . .	68
5.4	The format of the abstract syntax tree in Eclipse.	69
5.5	The Visitor Pattern.	70

List of Tables

5.1	The elements of the Java Model [Vog12].	66
8.1	The IFI Refaktor Case Study quality profile (version 6). . . .	91
8.2	Configuration for Case 1.	92
8.3	Statistics after batch refactoring the Eclipse JDT UI project with the <i>Extract and Move Method</i> refactoring.	101
8.4	Results for analyzing the Eclipse JDT UI project, before and after the refactoring, with SonarQube and the IFI Refaktor Case Study quality profile. (Bold numbers are better.)	102
8.5	Configuration for Case 2.	103
8.6	Statistics after batch refactoring the no.uio.ifi.refaktor project with the <i>Extract and Move Method</i> refactoring. . . .	104
8.7	Results for analyzing the no.uio.ifi.refaktor project, before and after the refactoring, with SonarQube and the IFI Refaktor Case Study quality profile. (Bold numbers are better.)	105

List of listings

1	Moving method f from C to X	21
2	The target and the destination for the composition of the Extract Method and <i>Move Method</i> refactorings.	26
3	The result of the composed refactoring.	27
4	An example of an <i>Extract Method</i> refactoring.	30
5	An example of a <i>Move Method</i> refactoring.	31
6	An example of the <i>Extract and Move Method</i> refactoring. . .	31
7	An example of improving CBO. Class C has a CBO value of 4 before refactoring it, and 3 after.	33
8	Classes A and B are both public. The methods foo and bar are public members of class A	39
9	Example of how the text selections generator would generate text selections based on a lists of statements. Each highlighted rectangle represents a text selection.	40
10	When the <i>Extract and Move Method</i> tries to use a variable with a local type as the move target, an intermediate step is performed that is not allowed. Here: LocalClass is not in the scope of fooBar in its intermediate location.	47
11	The result after refactoring the class C of listing 8 on page 39 with the <i>Extract and Move Method</i> refactoring with ((16, 17), b.a , <i>f</i> (2)) as input.	51
12	An ASTVisitor that visits all the names in a subtree and adds them to a collection, except those names that are children of any QualifiedName	71
13	An example of a pointcut named methodAnalyze , and an advice defined to be applied after it has occurred.	85
14	Showing the basic usage of soft references. Weak references is used the same way. (The references are part of the java.lang.ref package.)	86

Chapter 1

Introduction

1.1 Motivation and structure

For large software projects, complex program source code is an issue. It impacts the cost of maintenance in a negative way. It often stalls the implementation of new functionality and other program changes. The code may be difficult to understand, the changes may introduce new bugs that are hard to find and its complexity can simply keep people from doing code changes in fear of breaking some dependent piece of code. All these problems are related, and often lead to a vicious circle that slowly degrades the overall quality of a project.

More specifically, and in an object-oriented context, a class may depend on a number of other classes. Sometimes these intimate relationships are appropriate, and sometimes they are not. Inappropriate *coupling* between classes can make it difficult to know whether or not a change that is aimed at fixing a specific problem also alters the behavior of another part of a program.

One of the tools that are used to fight complexity and coupling in program source code is *refactoring*. The intention for this master's thesis is therefore to create an automated composite refactoring that reduces coupling between classes. The refactoring shall be able to operate automatically in all phases of a refactoring, from performing analysis to executing changes. It is also a requirement that it should be able to process large quantities of source code in a reasonable amount of time.

The current chapter proceeds in section 1.2 by describing what refactoring is. Then the project is presented in section 1.3, before the chapter is concluded with a brief discussion of related work in section 1.4.

Chapter 2 shows the workings of our refactoring together with a simple example illustrating this.

Structure. Write later. . .

1.2 What is refactoring?

This question is best answered by first defining the concept of a *refactoring*, what it is to *refactor*, and then discuss what aspects of programming make

people want to refactor their code.

1.2.1 Defining refactoring

Martin Fowler, in his classic book on refactoring [Fow99], defines a refactoring like this:

Refactoring (noun): a change made to the internal structure¹ of software to make it easier to understand and cheaper to modify without changing its observable behavior. [Fow99, p. 53]

This definition assigns additional meaning to the word *refactoring*, beyond the composition of the prefix *re-*, usually meaning something like “again” or “anew”, and the word *factoring*, that can mean to isolate the *factors* of something. Here a *factor* would be close to the mathematical definition of something that divides a quantity, without leaving a remainder. Fowler is mixing the *motivation* behind refactoring into his definition. Instead it could be more refined, formed to only consider the *mechanical* and *behavioral* aspects of refactoring. That is to factor the program again, putting it together in a different way than before, while preserving the behavior of the program. An alternative definition could then be:

Definition. A *refactoring* is a transformation done to a program without altering its external behavior.

From this we can conclude that a refactoring primarily changes how the *code* of a program is perceived by the *programmer*, and not the *behavior* experienced by any user of the program. Although the logical meaning is preserved, such changes could potentially alter the program’s behavior when it comes to performance gain or -penalties. So any logic depending on the performance of a program could make the program behave differently after a refactoring.

In the extreme case one could argue that software obfuscation is refactoring. It is often used to protect proprietary software. It restrains uninvited viewers, so they have a hard time analyzing code that they are not supposed to know how works. This could be a problem when using a language that is possible to decompile, such as Java.

Obfuscation could be done composing many, more or less randomly chosen, refactorings. Then the question arises whether it can be called a *composite refactoring* or not (see section 1.2.9 on page 24)? The answer is not obvious. First, there is no way to describe the mechanics of software obfuscation, because there are infinitely many ways to do that. Second, obfuscation can be thought of as *one operation*: Either the code is obfuscated, or it is not. Third, it makes no sense to call software obfuscation *a refactoring*, since it holds different meaning to different people.

This last point is important, since one of the motivations behind defining different refactorings, is to establish a *vocabulary* for software professionals to use when reasoning about and discussing programs, similar to the motivation behind design patterns [Gam+95].

¹The structure observable by the programmer.

1.2.2 The etymology of 'refactoring'

It is a little difficult to pinpoint the exact origin of the word “refactoring”, as it seems to have evolved as part of a colloquial terminology, more than a scientific term. There is no authoritative source for a formal definition of it.

According to Martin Fowler [Fow03], there may also be more than one origin of the word. The most well-known source, when it comes to the origin of *refactoring*, is the Smalltalk¹ community and their infamous Refactoring Browser² described in the article *A Refactoring Tool for Smalltalk* [RBJ97], published in 1997. Allegedly [Fow03], the metaphor of factoring programs was also present in the Forth¹ community, and the word “refactoring” is mentioned in a book by Leo Brodie, called *Thinking Forth* [Bro04], first published in 1984³. The exact word is only printed one place [Bro04, p. 232], but the term *factoring* is prominent in the book, that also contains a whole chapter dedicated to (re)factoring, and how to keep the (Forth) code clean and maintainable.

... good factoring technique is perhaps the most important skill for a Forth programmer. [Bro04, p. 172]

Brodie also express what *factoring* means to him:

Factoring means organizing code into useful fragments. To make a fragment useful, you often must separate reusable parts from non-reusable parts. The reusable parts become new definitions. The non-reusable parts become arguments or parameters to the definitions. [Bro04, p. 172]

Fowler claims that the usage of the word *refactoring* did not pass between the Forth and Smalltalk communities, but that it emerged independently in each of the communities.

1.2.3 Reasons for refactoring

There are many reasons why people want to refactor their programs. They can for instance do it to remove duplication, break up long methods or to introduce design patterns into their software systems. The shared trait for all these are that peoples' intentions are to make their programs *better*, in some sense. But what aspects of their programs are becoming improved?

As just mentioned, people often refactor to get rid of duplication. They are moving identical or similar code into methods, and are pushing methods up or down in their class hierarchies. They are making template methods for overlapping algorithms/functionality, and so on. It is all about gathering

¹Programming language

²<http://st-www.cs.illinois.edu/users/brant/Refactory/RefactoringBrowser.html>

³*Thinking Forth* was first published in 1984 by the Forth Interest Group. Then it was reprinted in 1994 with minor typographical corrections, before it was transcribed into an electronic edition typeset in L^AT_EX and published under a Creative Commons licence in 2004. The edition cited here is the 2004 edition, but the content should essentially be as in 1984.

what belongs together and putting it all in one place. The resulting code is then easier to maintain. When removing the implicit coupling¹ between code snippets, the location of a bug is limited to only one place, and new functionality need only to be added to this one place, instead of a number of places people might not even remember.

A problem you often encounter when programming, is that a program contains a lot of long and hard-to-grasp methods. It can then help to break the methods into smaller ones, using the *Extract Method* refactoring [Fow99]. Then you may discover something about a program that you were not aware of before; revealing bugs you did not know about or could not find due to the complex structure of your program. Making the methods smaller and giving good names to the new ones clarifies the algorithms and enhances the *understandability* of the program (see section 1.2.4 on the facing page). This makes refactoring an excellent method for exploring unknown program code, or code that you had forgotten that you wrote.

Most primitive refactorings are simple, and usually involves moving code around [Ker05]. The motivation behind them may first be revealed when they are combined into larger — higher level — refactorings, called *composite refactorings* (see section 1.2.9 on page 24). Often the goal of such a series of refactorings is a design pattern. Thus the design can *evolve* throughout the lifetime of a program, as opposed to designing up-front. It is all about being structured and taking small steps to improve a program's design.

Many software design pattern are aimed at lowering the coupling between different classes and different layers of logic. One of the most famous is perhaps the *Model-View-Controller* [Gam+95] pattern. It is aimed at lowering the coupling between the user interface, the business logic and the data representation of a program. This also has the added benefit that the business logic could much easier be the target of automated tests, thus increasing the productivity in the software development process.

Another effect of refactoring is that with the increased separation of concerns coming out of many refactorings, the *performance* can be improved. When profiling programs, the problematic parts are narrowed down to smaller parts of the code, which are easier to tune, and optimization can be performed only where needed and in a more effective way [Fow99].

Last, but not least, and this should probably be the best reason to refactor, is to refactor to *facilitate a program change*. If one has managed to keep one's code clean and tidy, and the code is not bloated with design patterns that are not ever going to be needed, then some refactoring might be needed to introduce a design pattern that is appropriate for the change that is going to happen.

Refactoring program code — with a goal in mind — can give the code itself more value. That is in the form of robustness to bugs, understandability and maintainability. Having robust code is an obvious advantage, but understandability and maintainability are both very

¹When duplicating code, the duplicate pieces of code might not be coupled, apart from representing the same functionality. So if this functionality is going to change, it might need to change in more than one place, thus creating an implicit coupling between multiple pieces of code.

important aspects of software development. By incorporating refactoring in the development process, bugs are found faster, new functionality is added more easily and code is easier to understand by the next person exposed to it, which might as well be the person who wrote it. The consequence of this, is that refactoring can increase the average productivity of the development process, and thus also add to the monetary value of a business in the long run. The perspective on productivity and money should also be able to open the eyes of the many nearsighted managers that seldom see beyond the next milestone.

1.2.4 The magical number seven

The article *The magical number seven, plus or minus two: some limits on our capacity for processing information* [Mil56] by George A. Miller, was published in the journal *Psychological Review* in 1956. It presents evidence that support that the capacity of the number of objects a human being can hold in its working memory is roughly seven, plus or minus two objects. This number varies a bit depending on the nature and complexity of the objects, but is according to Miller "...never changing so much as to be unrecognizable."

Miller's article culminates in the section called *Recoding*, a term he borrows from communication theory. The central result in this section is that by recoding information, the capacity of the amount of information that a human can process at a time is increased. By *recoding*, Miller means to group objects together in chunks, and give each chunk a new name that it can be remembered by.

...recoding is an extremely powerful weapon for increasing the amount of information that we can deal with. [Mil56, p. 95]

By organizing objects into patterns of ever growing depth, one can memorize and process a much larger amount of data than if it were to be represented as its basic pieces. This grouping and renaming is analogous to how many refactorings work, by grouping pieces of code and give them a new name. Examples are the fundamental *Extract Method* and *Extract Class* refactorings [Fow99].

An example from the article addresses the problem of memorizing a sequence of binary digits. The example presented here is a slightly modified version of the one presented in the original article [Mil56], but it preserves the essence of it. Let us say we have the following sequence of 16 binary digits: "1010001001110011". Most of us will have a hard time memorizing this sequence by only reading it once or twice. Imagine if we instead translate it to this sequence: "A273". If you have a background from computer science, it will be obvious that the latter sequence is the first sequence recoded to be represented by digits in base 16. Most people should be able to memorize this last sequence by only looking at it once.

Another result from the Miller article is that when the amount of information a human must interpret increases, it is crucial that the

translation from one code to another must be almost automatic for the subject to be able to remember the translation, before he is presented with new information to recompile. Thus learning and understanding how to best organize certain kinds of data is essential to efficiently handle that kind of data in the future. This is much like when humans learn to read. First they must learn how to recognize letters. Then they can learn distinct words, and later read sequences of words that form whole sentences. Eventually, most of them will be able to read whole books and briefly retell the important parts of its content. This suggests that the use of design patterns is a good idea when reasoning about computer programs. With extensive use of design patterns when creating complex program structures, one does not always have to read whole classes of code to comprehend how they function, it may be sufficient to only see the name of a class to almost fully understand its responsibilities.

Our language is tremendously useful for repackaging material into a few chunks rich in information. [Mil56, p. 95]

Without further evidence, these results at least indicate that refactoring source code into smaller units with higher cohesion and, when needed, introducing appropriate design patterns, should aid in the cause of creating computer programs that are easier to maintain and have code that is easier (and better) understood.

1.2.5 Notable contributions to the refactoring literature

1992 William F. Opdyke submits his doctoral dissertation called *Refactoring Object-Oriented Frameworks* [Opd92]. This work defines a set of refactorings, that are behavior preserving given that their preconditions are met. The dissertation is focused on the automation of refactorings.

1999 Martin Fowler et al.: *Refactoring: Improving the Design of Existing Code* [Fow99]. This is maybe the most influential text on refactoring. It bares similarities with Opdykes thesis [Opd92] in the way that it provides a catalog of refactorings. But Fowler’s book is more about the craft of refactoring, as he focuses on establishing a vocabulary for refactoring, together with the mechanics of different refactorings and when to perform them. His methodology is also founded on the principles of test-driven development.

2005 Joshua Kerievsky: *Refactoring to Patterns* [Ker05]. This book is heavily influenced by Fowler’s *Refactoring* [Fow99] and the “Gang of Four” *Design Patterns* [Gam+95]. It is building on the refactoring catalogue from Fowler’s book, but is trying to bridge the gap between *refactoring* and *design patterns* by providing a series of higher-level composite refactorings, that makes code evolve toward or away from certain design patterns. The book is trying to build up the reader’s intuition around *why* one would want to use a particular design

pattern, and not just *how*. The book is encouraging evolutionary design (see section 1.2.7 on the following page).

1.2.6 Tool support (for Java)

This section will briefly compare the refactoring support of the three IDEs Eclipse¹, IntelliJ IDEA² and NetBeans³. These are the most popular Java IDEs [11].

All three IDEs provide support for the most useful refactorings, like the different extract, move and rename refactorings. In fact, Java-targeted IDEs are known for their good refactoring support, so this did not appear as a big surprise.

The IDEs seem to have excellent support for the *Extract Method* refactoring, so at least they have all passed the first “refactoring rubicon” [Fow01; VJ12].

Regarding the *Move Method* refactoring, the Eclipse and IntelliJ IDEs do the job in very similar manners. In most situations they both do a satisfying job by producing the expected outcome. But they do nothing to check that the result does not break the semantics of the program (see section 1.2.11 on page 25). The NetBeans IDE implements this refactoring in a somewhat unsophisticated way. For starters, the refactoring’s default destination for the move, is the same class as the method already resides in, although it refuses to perform the refactoring if chosen. But the worst part is, that if moving the method **f** of the class **C** to the class **X**, it will break the code. The result is shown in listing 1 on this page.

```
public class C {
    private X x;
    ...
    public void f() {
        x.m();
        x.n();
    }
}

public class X {
    ...
    public void f(C c) {
        c.x.m();
        c.x.n();
    }
}
```

Listing 1: Moving method **f** from **C** to **X**.

NetBeans will try to create code that call the methods **m** and **n** of **X** by accessing them through **c.x**, where **c** is a parameter of type **C** that is added the method **f** when it is moved. (This is seldom the desired outcome of this refactoring, but ironically, this “feature” keeps NetBeans from breaking the code in the example from section 1.2.11 on page 25.) If **c.x** for some reason is inaccessible to **X**, as in this case, the refactoring breaks the code, and it will not compile. NetBeans presents a preview of the refactoring outcome, but the preview does not catch it if the IDE is about break the program.

¹<http://www.eclipse.org/>

²The IDE under comparison is the Community Edition, <http://www.jetbrains.com/idea/>

³<https://netbeans.org/>

The IDEs under investigation seem to have fairly good support for primitive refactorings, but what about more complex ones, such as *Extract Class* [Fow99]? IntelliJ handles this in a fairly good manner, although, in the case of private methods, it leaves unused methods behind. These are methods that delegate to a field with the type of the new class, but are not used anywhere. Eclipse has added its own quirk to the *Extract Class* refactoring, and only allows for *fields* to be moved to a new class, *not methods*. This makes it effectively only extracting a data structure, and calling it *Extract Class* is a little misleading. One would often be better off with textual extract and paste than using the *Extract Class* refactoring in Eclipse. When it comes to NetBeans, it does not even show an attempt on providing this refactoring.

1.2.7 The relation to design patterns

Refactoring and design patterns have at least one thing in common, they are both promoted by advocates of *clean code* [MC09] as fundamental tools on the road to more maintainable and extendable source code.

Design patterns help you determine how to reorganize a design, and they can reduce the amount of refactoring you need to do later. [Gam+95, p. 353]

Although sometimes associated with over-engineering [Ker05; Fow99], design patterns are in general assumed to be good for maintainability of source code. That may be because many of them are designed to support the *open/closed principle* of object-oriented programming. The principle was first formulated by Bertrand Meyer, the creator of the Eiffel programming language, like this: “Modules should be both open and closed.” [Mey88] It has been popularized, with this as a common version:

Software entities (classes, modules, functions, etc.) should be open for extension, but closed for modification.¹

Maintainability is often thought of as the ability to be able to introduce new functionality without having to change too much of the old code. When refactoring, the motivation is often to facilitate adding new functionality. It is about factoring the old code in a way that makes the new functionality being able to benefit from the functionality already residing in a software system, without having to copy old code into new. Then, next time someone shall add new functionality, it is less likely that the old code has to change. Assuming that a design pattern is the best way to get rid of duplication and assist in implementing new functionality, it is reasonable to conclude that a design pattern often is the target of a series of refactorings. Having a repertoire of design patterns can also help in knowing when and how to refactor a program to make it reflect certain desired characteristics.

¹See <http://c2.com/cgi/wiki?OpenClosedPrinciple> or https://en.wikipedia.org/wiki/Open/closed_principle

There is a natural relation between patterns and refactorings. Patterns are where you want to be; refactorings are ways to get there from somewhere else. [Fow99, p. 107]

This quote is wise in many contexts, but it is not always appropriate to say “Patterns are where you want to be...”. *Sometimes*, patterns are where you want to be, but only because it will benefit your design. It is not true that one should always try to incorporate as many design patterns as possible into a program. It is not like they have intrinsic value. They only add value to a system when they support its design. Otherwise, the use of design patterns may only lead to a program that is more complex than necessary.

The overuse of patterns tends to result from being patterns happy. We are *patterns happy* when we become so enamored of patterns that we simply must use them in our code. [Ker05, p. 24]

This can easily happen when relying largely on up-front design. Then it is natural, in the very beginning, to try to build in all the flexibility that one believes will be necessary throughout the lifetime of a software system. According to Joshua Kerievsky “That sounds reasonable — if you happen to be psychic.” [Ker05, p. 1] He is advocating what he believes is a better approach: To let software continually evolve. To start with a simple design that meets today’s needs, and tackle future needs by refactoring to satisfy them. He believes that this is a more economic approach than investing time and money into a design that inevitably is going to change. By relying on continuously refactoring a system, its design can be made simpler without sacrificing flexibility. To be able to fully rely on this approach, it is of utter importance to have a reliable suit of tests to lean on (see section 1.2.12 on page 27). This makes the design process more natural and less characterized by difficult decisions that has to be made before proceeding in the process, and that is going to define a project for all of its unforeseeable future.

1.2.8 The impact on software quality

What is software quality?

The term *software quality* has many meanings. It all depends on the context we put it in. If we look at it with the eyes of a software developer, it usually means that the software is easily maintainable and testable, or in other words, that it is *well designed*. This often correlates with the management scale, where *keeping the schedule* and *customer satisfaction* is at the center. From the customers point of view, in addition to good usability, *performance* and *lack of bugs* is always appreciated, measurements that are also shared by the software developer. (In addition, such things as good documentation could be measured, but this is out of the scope of this document.)

The impact on performance

Refactoring certainly will make software go more slowly¹, but it also makes the software more amenable to performance tuning. [Fow99, p. 69]

There is a common belief that refactoring compromises performance, due to increased degree of indirection and that polymorphism is slower than conditionals.

In a survey, Demeyer [Dem02] disproves this view in the case of polymorphism. He did an experiment on, what he calls, “Transform Self Type Checks” where you introduce a new polymorphic method and a new class hierarchy to get rid of a class’ type checking of a “type attribute“. He uses this kind of transformation to represent other ways of replacing conditionals with polymorphism as well. The experiment is performed on the C++ programming language and with three different compilers and platforms. Demeyer concludes that, with compiler optimization turned on, polymorphism beats middle to large sized if-statements and does as well as case-statements. (In accordance with his hypothesis, due to similarities between the way C++ handles polymorphism and case-statements.)

The interesting thing about performance is that if you analyze most programs, you find that they waste most of their time in a small fraction of the code. [Fow99, p. 70]

So, although an increased amount of method calls could potentially slow down programs, one should avoid premature optimization and sacrificing good design, leaving the performance tuning until after profiling the software and having isolated the actual problem areas.

1.2.9 Composite refactorings

Generally, when thinking about refactoring, at the mechanical level, there are essentially two kinds of refactorings. There are the *primitive* refactorings, and the *composite* refactorings.

Definition. A *primitive refactoring* is a refactoring that cannot be expressed in terms of other refactorings.

Examples are the *Pull Up Field* and *Pull Up Method* refactorings [Fow99], that move members up in their class hierarchies.

Definition. A *composite refactoring* is a refactoring that can be expressed in terms of two or more other refactorings.

An example of a composite refactoring is the *Extract Superclass* refactoring [Fow99]. In its simplest form, it is composed of the previously described primitive refactorings, in addition to the *Pull Up Constructor Body* refactoring [Fow99]. It works by creating an abstract superclass that the target

¹With today’s compiler optimization techniques and performance tuning of e.g. the Java virtual machine, the penalties of object creation and method calls are debatable.

class(es) inherits from, then by applying *Pull Up Field*, *Pull Up Method* and *Pull Up Constructor Body* on the members that are to be members of the new superclass. If there are multiple classes in play, their interfaces may need to be united with the help of some rename refactorings, before extracting the superclass. For an overview of the *Extract Superclass* refactoring, see fig. 1.1 on the current page.

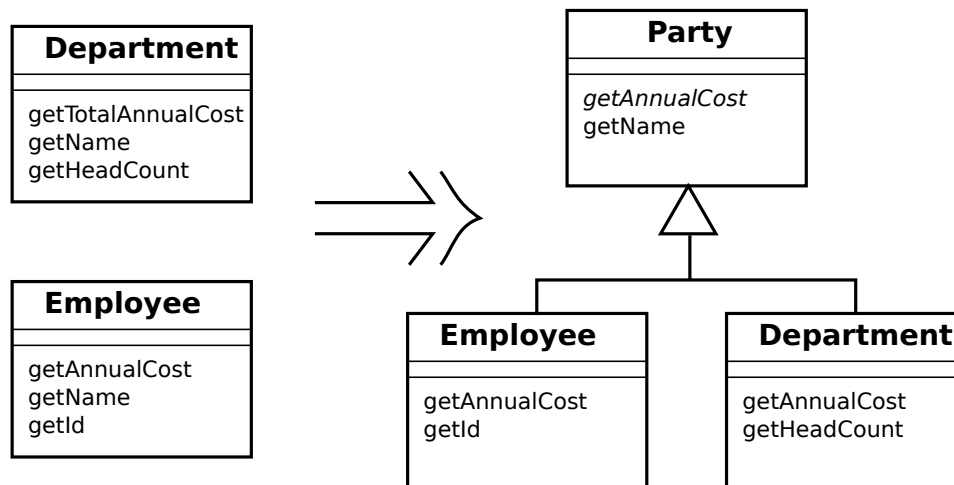


Figure 1.1: The Extract Superclass refactoring, with united interfaces.

1.2.10 Manual vs. automated refactorings

Refactoring is something every programmer does, even if she does not know the term *refactoring*. Every refinement of source code that does not alter the program's behavior is a refactoring. For small refactorings, such as *Extract Method*, executing it manually is a manageable task, but is still prone to errors. Getting it right the first time is not easy, considering the method signature and all the other aspects of the refactoring that has to be in place.

Consider the renaming of classes, methods and fields. For complex programs these refactorings are almost impossible to get right. Attacking them with textual search and replace, or even regular expressions, will fall short on these tasks. Then it is crucial to have proper tool support that can perform them automatically. Tools that can parse source code and thus have semantic knowledge about which occurrences of which names belong to what construct in the program. For even trying to perform one of these complex task manually, one would have to be very confident on the existing test suite (see section 1.2.12 on page 27).

1.2.11 Correctness of refactorings

For automated refactorings to be truly useful, they must show a high degree of behavior preservation. This last sentence might seem obvious, but there are examples of refactorings in existing tools that break programs. In an ideal world, every automated refactoring would be "complete", in the sense

that it would never break a program. In an ideal world, every program would also be free from bugs. In modern IDEs the implemented automated refactorings are working for *most* cases, which is enough for making them useful.

I will now present an example of a *corner case* where a program breaks when a refactoring is applied. The example shows an *Extract Method* refactoring followed by a *Move Method* refactoring that breaks a program in both the Eclipse and IntelliJ IDEs¹. The target and the destination for the composed refactoring is shown in listing 2 on this page. Note that the method `m(C c)` of class `X` assigns to the field `x` of the argument `c` that has type `C`.

———— Refactoring target ————	———— Method destination ————
<pre> 1 public class C { 2 public X x = new X(); 3 4 public void f() { 5 x.m(this); 6 // Not the same x 7 x.n(); 8 } 9 }</pre>	<pre> public class X { public void m(C c) { c.x = new X(); // If m is called from // c, then c.x no longer // equals 'this' } public void n() {} }</pre>

Listing 2: The target and the destination for the composition of the Extract Method and *Move Method* refactorings.

The refactoring sequence works by extracting line 6 through 8 from the original class `C` into a method `f` with the statements from those lines as its method body (but with the comment left out, since it will no longer hold any meaning). The method is then moved to the class `X`. The result is shown in listing 3 on the next page.

Before the refactoring, the methods `m` and `n` of class `X` are called on different object instances (see line 6 and 8 of the original class `C` in listing 2). After the refactoring, they are called on the same object, and the statement on line 3 of class `X` (in listing 3) no longer has the desired effect in our example. The method `f` of class `C` is now calling the method `f` of class `X` (see line 5 of class `C` in listing 3), and the program now behaves different than before.

The bug introduced in the previous example is of such a nature² that it is very difficult to spot if the refactored code is not covered by tests. It does not generate compilation errors, and will thus only result in a runtime error or corrupted data, which might be hard to detect.

¹The NetBeans IDE handles this particular situation without altering the program's behavior, mainly because its *Move Method* refactoring implementation is a bit flawed in other ways (see section 1.2.6 on page 21).

²Caused by aliasing. See [https://en.wikipedia.org/wiki/Aliasing_\(computing\)](https://en.wikipedia.org/wiki/Aliasing_(computing))

```

1 public class C {
2     public X x = new X();
3
4     public void f() {
5         x.f(this);
6     }
7 }

1 public class X {
2     public void m(C c) {
3         c.x = new X();
4     }
5     public void n() {}
6     // Extracted and
7     // moved method
8     public void f(C c) {
9         m(c);
10        n();
11    }
12 }

```

Listing 3: The result of the composed refactoring.

1.2.12 Refactoring and the importance of testing

If you want to refactor, the essential precondition is having solid tests. [Fow99]

When refactoring, there are roughly three classes of errors that can be made. The first class of errors are the ones that make the code unable to compile. These *compile-time* errors are of the nicer kind. They flash up at the moment they are made (at least when using an IDE), and are usually easy to fix. The second class are the *runtime* errors. Although they take a bit longer to surface, they usually manifest after some time in an illegal argument exception, null pointer exception or similar during the program execution. These kind of errors are a bit harder to handle, but at least they will show, eventually. Then there are the *behavior-changing* errors. These errors are of the worst kind. They do not show up during compilation and they do not turn on a blinking red light during runtime either. The program can seem to work perfectly fine with them in play, but the business logic can be damaged in ways that will only show up over time.

For discovering runtime errors and behavior changes when refactoring, it is essential to have good test coverage. Testing in this context means writing automated tests. Manual testing may have its uses, but when refactoring, it is automated unit testing that dominate. For discovering behavior changes it is especially important to have tests that cover potential problems, since these kind of errors does not reveal themselves.

Unit testing is not a way to *prove* that a program is correct, but it is a way to make you confident that it *probably* works as desired. In the context of test-driven development (commonly known as TDD), the tests are even a way to define how the program is *supposed* to work. It is then, by definition, working if the tests are passing.

If the test coverage for a code base is perfect, then it should, theoretically, be risk-free to perform refactorings on it. This is why automated tests and refactoring are such a great match.

Testing the code from correctness section

The worst thing that can happen when refactoring is to introduce changes to the behavior of a program, as in the example on section 1.2.11 on page 25. This example may be trivial, but the essence is clear. The only problem with the example is that it is not clear how to create automated tests for it, without changing it in intrusive ways.

Unit tests, as they are known from the different xUnit frameworks around, are only suitable to test the *result* of isolated operations. They can not easily (if at all) observe the *history* of a program.

This problem is still open.

1.3 The Project

In this section we look at the work that shall be done for this project, its building stones and some of the methodologies used.

1.3.1 Project description

The aim of this master's project will be to explore the relationship between the *Extract Method* and the *Move Method* refactorings. This will be done by composing the two into a composite refactoring. The refactoring will be called the *Extract and Move Method* refactoring.

The two primitive *Extract Method* and *Move Method* refactorings must already be implemented in a tool, so the *Extract and Move Method* refactoring is going to be built on top of those.

The composition of the *Extract Method* and *Move Method* refactorings springs naturally out of the need to move procedures closer to the data they manipulate. This composed refactoring is not well described in the literature, but it is implemented in at least one tool called CodeRush¹, that is an extension for MS Visual Studio². In CodeRush it is called *Extract Method to Type*³, but I choose to call it *Extract and Move Method*, since I feel this better communicates which primitive refactorings it is composed of.

The project will consist of implementing the *Extract and Move Method* refactoring, as well as executing it over a larger code base, as a case study. To be able to execute the refactoring automatically, I have to make it analyze code to determine the best selections to extract into new methods.

1.3.2 The premises

Before we can start manipulating source code and write a tool for doing so, we need to decide on a programming language for the code we are going to manipulate. Also, since we do not want to start from scratch by implementing primitive refactorings ourselves, we need to choose an existing tool that provides the needed refactorings. In addition to be able to perform

¹<https://help.devexpress.com/#CodeRush/CustomDocument3519>

²<http://www.visualstudio.com/>

³<https://help.devexpress.com/#CodeRush/CustomDocument6710>

changes, we need a framework for analyzing source code for the language we select.

Choosing the target language

Choosing which programming language the code that shall be manipulated shall be written in, is not a very difficult task. We choose to limit the possible languages to the object-oriented programming languages, since most of the terminology and literature regarding refactoring comes from the world of object-oriented programming. In addition, the language must have existing tool support for refactoring.

The Java programming language¹ is the dominating language when it comes to example code in the literature of refactoring, and is thus a natural choice. Java is perhaps, currently the most influential programming language in the world, with its Java Virtual Machine that runs on all of the most popular architectures and also supports dozens of other programming languages², with Scala, Clojure and Groovy as the most prominent ones. Java is currently the language that every other programming language is compared against. It is also the primary programming language for the author of this thesis.

Choosing the tools

When choosing a tool for manipulating Java, there are certain criteria that have to be met. First of all, the tool should have some existing refactoring support that this thesis can build upon. Secondly it should provide some kind of framework for parsing and analyzing Java source code. Third, it should itself be open source. This is both because of the need to be able to browse the code for the existing refactorings that is contained in the tool, and also because open source projects hold value in them selves. Another important aspect to consider is that open source projects of a certain size, usually has large communities of people connected to them, that are committed to answering questions regarding the use and misuse of the products, that to a large degree is made by the community itself.

There is a certain class of tools that meet these criteria, namely the class of *IDEs*³. These are programs that is meant to support the whole production cycle of a computer program, and the most popular IDEs that support Java, generally have quite good refactoring support.

The main contenders for this thesis is the Eclipse IDE, with the Java development tools (JDT), the IntelliJ IDEA Community Edition and the NetBeans IDE (see section 1.2.6 on page 21). Eclipse and NetBeans are both free, open source and community driven, while the IntelliJ IDEA has an open sourced community edition that is free of charge, but also offer an Ultimate Edition with an extended set of features, at additional cost. All three IDEs supports adding plugins to extend their functionality and tools

¹<https://www.java.com/>

²They compile to Java bytecode.

³*Integrated Development Environment*

that can be used to parse and analyze Java source code. But one of the IDEs stand out as a favorite, and that is the Eclipse IDE. This is the most popular [11] among them and seems to be de facto standard IDE for Java development regardless of platform.

1.3.3 The primitive refactorings

The refactorings presented here are the primitive refactorings used in this project. They are the abstract building blocks used by the *Extract and Move Method* refactoring.

The Extract Method refactoring The *Extract Method* refactoring is used to extract a fragment of code from its context and into a new method. A call to the new method is inlined where the fragment was before. It is used to break code into logical units, with names that explain their purpose.

An example of an *Extract Method* refactoring is shown in listing 4 on this page. It shows a method containing calls to the methods `foo` and `bar` of a type `X`. These statements are then extracted into the new method `fooBar`.

Before	After
<pre>class C { void method() { X x = new X(); x.foo(); x.bar(); } }</pre>	<pre>class C { void method() { X x = new X(); fooBar(x); } void fooBar(X x) { x.foo(); x.bar(); } }</pre>

Listing 4: An example of an *Extract Method* refactoring.

The Move Method refactoring The *Move Method* refactoring is used to move a method from one class to another. This can be appropriate if the method is using more features of another class than of the class which it is currently defined.

Listing 5 on the next page shows an example of this refactoring. Here a method `fooBar` is moved from the class `C` to the class `X`.

1.3.4 The Extract and Move Method refactoring

The *Extract and Move Method* refactoring is a composite refactoring composed of the primitive *Extract Method* and *Move Method* refactorings. The effect of this refactoring on source code is the same as when extracting a method and moving it to another class. Conceptually, this is done without an intermediate step. In practice, as we shall see later, an intermediate step may be necessary.

Before	After
<pre> class C { void method() { X x = new X(); fooBar(x); } void fooBar(X x) { x.foo(); x.bar(); } } class X { void foo(){/*...*/} void bar(){/*...*/} } </pre>	<pre> class C { void method() { X x = new X(); x.fooBar(); } } class X { void fooBar() { foo(); bar(); } void foo(){/*...*/} void bar(){/*...*/} } </pre>

Listing 5: An example of a *Move Method* refactoring.

An example of this composite refactoring is shown in listing 6 on this page. The example joins the examples from listing 4 and listing 5. This means that the selection consisting of the consecutive calls to the methods **foo** and **bar**, is extracted into a new method **fooBar** located in the class **X**.

Before	After
<pre> class C { void method() { X x = new X(); x.foo(); x.bar(); } } class X { void foo(){/*...*/} void bar(){/*...*/} } </pre>	<pre> class C { void method() { X x = new X(); x.fooBar(); } } class X { void fooBar() { foo(); bar(); } void foo(){/*...*/} void bar(){/*...*/} } </pre>

Listing 6: An example of the *Extract and Move Method* refactoring.

1.3.5 The Coupling Between Object Classes metric

The best known metric for measuring coupling between classes in object-oriented software is called *Coupling Between Object Classes*, usually

abbreviated as CBO. The metric is defined in the article *A Metrics Suite for Object Oriented Design* [CK94] by Chidamber and Kemerer, published in 1994.

Definition. *CBO* for a class is a count of the number of other classes to which it is coupled.

An object is coupled to another object if one of them acts on the other by using methods or instance variables of the other object. This relation goes both ways, so both outgoing and incoming uses are counted. Each coupling relationship is only considered once when measuring CBO for a class.

How can the Extract and Move Method refactoring improve CBO?

Listing 7 on the next page shows how CBO changes for a class when it is refactored with the *Extract and Move Method* refactoring. In the example we consider only the CBO value of class **C**.

Before refactoring the class **C** with the *Extract and Move Method* refactoring, it has a CBO value of 4. The class uses members of the classes **A** and **B**, which accounts for 2 of the coupling relationships of class **C**. In addition to this, it uses its variable **x** with type **X** and also the methods **foo** and **bar** declared in class **Y**, giving it a total CBO value of 4.

The after-part of the example code in listing 7 shows the result of extracting the lines 5 and 6 of class **C** into a new method **fooBar**, with a subsequent move of it to class **X**.

With respect to the CBO metric, the refactoring action accomplishes something important: It eliminates the uses of class **Y** from class **C**. This means that the class **C** is no longer coupled to **Y**, only the classes **A**, **B** and **X**. The CBO value of class **C** is therefore 3 after refactoring, while no other class have received any increase in CBO.

The example shown here is an ideal situation. Coupling is reduced for one class without any increase of coupling for another class. There is also another point that is important. It is the fact that to reduce the CBO value for a class, we need to remove *all* its uses of another class. This is done for the class **C** in listing 7 on the facing page, where all uses of class **Y** is removed by the *Extract and Move Method* refactoring.

Highlight code

1.3.6 Research questions

The main question that I seek an answer to in this thesis is:

Is it possible to automate the analysis and execution of the *Extract and Move Method* refactoring, and do so for all of the code of a larger project?

The secondary questions will then be:

Can we do this efficiently? Can we automate the analysis and execution of the refactoring so it can be run in a reasonable amount of time?

Before	After
1 <code>class C {</code>	1 <code>class C {</code>
2 <code> A a; B b;</code>	2 <code> A a; B b;</code>
3 <code> X x;</code>	3 <code> X x;</code>
4 <code> void method() {</code>	4 <code> void method() {</code>
5 <code> x.y.foo();</code>	5 <code> x.fooBar();</code>
6 <code> x.y.bar();</code>	6 <code> }</code>
7 <code> }</code>	7 <code> <i>/* Uses of A and B.</i></code>
8 <code> <i>/* Uses of A and B.</i></code>	8 <code> <i>No uses of other</i></code>
9 <code> <i>No uses of other</i></code>	9 <code> <i>classes. */</i></code>
10 <code> <i>classes. */</i></code>	10 <code> }</code>
11 <code>}</code>	11 <code></code>
12 <code>class X {</code>	12 <code>class X {</code>
13 <code> Y y;</code>	13 <code> Y y;</code>
14 <code> <i>/* No uses of C.</i></code>	14 <code> <i>/* No uses of C.</i></code>
15 <code> <i>Uses of Y. */</i></code>	15 <code> <i>Uses of Y. */</i></code>
16 <code> void fooBar() {</code>	16 <code> void fooBar() {</code>
17 <code> y.foo();</code>	17 <code> y.foo();</code>
18 <code> y.bar();</code>	18 <code> y.bar();</code>
19 <code> }</code>	19 <code> }</code>
20 <code>}</code>	20 <code>}</code>
21 <code>class Y {</code>	21 <code>class Y {</code>
22 <code> void foo(){</code>	22 <code> void foo(){</code>
23 <code> <i>/* No uses of C. */</i></code>	23 <code> void foo(){</code>
24 <code> }</code>	24 <code> <i>/* No uses of C. */</i></code>
25 <code> void bar(){</code>	25 <code> }</code>
26 <code> <i>/* No uses of C. */</i></code>	26 <code> void bar(){</code>
27 <code> }</code>	27 <code> <i>/* No uses of C. */</i></code>
28 <code>}</code>	28 <code> }</code>
29 <code></code>	29 <code>}</code>

Listing 7: An example of improving CBO. Class `C` has a CBO value of 4 before refactoring it, and 3 after.

Can we perform changes safely? Can we take actions to prevent the refactoring from breaking the code? By breaking the code we mean to either do changes that do not compile, or make changes that alter the behavior of the program.

Can we improve the quality of source code? Assuming that the refactoring is safe: Is it feasible to assure that the code we refactor actually gets better in terms of coupling?

How can the automation of the refactoring be helpful? Assuming the refactoring does in fact improve the quality of source code and is safe to use: What is the usefulness of the refactoring in a software development

setting? In what parts of the development process can the refactoring play a role?

1.3.7 Methodology

This section will present some of the methods used during the work of this thesis.

Evolutionary design

In the programming work for this project, it have tried to use a design strategy called evolutionary design, also known as continuous or incremental design [14b]. It is a software design strategy advocated by the Extreme Programming community. The essence of the strategy is that you should let the design of your program evolve naturally as your requirements change. This is seen in contrast with up-front design, where design decisions are made early in the process.

The motivation behind evolutionary design is to keep the design of software as simple as possible. This means not introducing unneeded functionality into a program. You should defer introducing flexibility into your software, until it is needed to be able to add functionality in a clean way.

Holding up design decisions, implies that the time will eventually come when decisions have to be made. The flexibility of the design then relies on the programmer's abilities to perform the necessary refactoring, and her confidence in those abilities. From my experience working on this project, I can say that this confidence is greatly enhanced by having automated tests to rely on (see section 1.3.7 on the current page).

The choice of going for evolutionary design developed naturally. As Fowler points out in his article *Is Design Dead?*, evolutionary design much resembles the "code and fix" development strategy [Fow04]. A strategy that most of us have practiced in school. This was also the case when I first started this work. I had to learn the inner workings of Eclipse and its refactoring-related plugins. That meant a lot of fumbling around with code I did not know, in a trial and error fashion. Eventually I started writing tests for my code, and my design began to evolve.

Test-driven development

As mentioned before, the project started out as a classic code and fix development process. My focus was aimed at getting something to work, rather than doing so according to best practice. This resulted in a project that got out of its starting blocks, but it was not accompanied by any tests. Hence it was soon difficult to make any code changes with the confidence that the program was still correct afterwards (assuming it was so before changing it). I always knew that I had to introduce some tests at one point, but this experience accelerated the process of leading me onto the path of testing.

I then wrote tests for the core functionality of the plugin, and thus gained more confidence in the correctness of my code. I could now perform quite drastic changes without “wetting my pants“. After this, nearly all of the semantic changes done to the business logic of the project, or the addition of new functionality, were made in a test-driven manner. This means that before performing any changes, I would define the desired functionality through a set of tests. I would then run the tests to check that they were run and that they did not pass. Then I would do any code changes necessary to make the tests pass. The definition of how the program is supposed to operate is then captured by the tests. However, this does not prove the correctness of the analysis leading to the test definitions.

1.3.8 Case study

Write

1.3.9 Dogfooding

Write

1.4 Related Work

1.4.1 Refactoring safety

This section presents a couple of approaches to improving the safety of performing refactorings. In these approaches, the problems that are addressed are not compilation problems, but behavior-altering problems that are not easily discovered during static analysis of source code. An example of such a problem is presented in section 1.2.11 on page 25.

Project “Safer Refactorings”

[Sto] is a proposal for a master’s thesis. The proposer is my supervisor, Volker Stolz from the University of Oslo.

The proposed solution for making refactorings safer, is to insert assertions into source code when refactoring it. For the example in listing 3 on page 27, that is the result of a refactoring, it is suggested that we insert an assert statement between lines 9 and 10. In this example, the assert statement would be

```
assert c.x == this;
```

which would discover the aliasing problems of this example.

“Making Program Refactoring Safer”

This is the name of an article [Soa+10] about providing a way to improve safety during refactoring. Soares et al. approaches the problem of preserving behavior during refactoring by analyzing a transformation and then generate

a test suite for it, using static analysis. These tests are then run for both the before- and after-code, and is compared to assure that they are consistent.

1.4.2 Search-based refactoring

Search-Based Refactoring: an empirical study [OC08] is a paper by Mark O’Keeffe and Mel Ó Cinnéide published in 2008. The authors present an empirical study of different algorithmic approaches to search-based refactoring.

The common approach for all these algorithms is to generate a set of changes to a program for then to use a “fitness function” to evaluate if they improve its design or not. The fitness function consists of a weighted sum of different object-oriented metrics.

Among other things, the authors conclude that even with small input programs, their solution representation is memory-intensive, at least for some algorithms. The programs they refactor on have in average 4,000 lines of code, spread over 57 classes. I.e. considerably smaller than one of the programs that will be subject to refactoring in this project.

1.4.3 The compositional paradigm of refactoring

This paradigm builds upon the observation of Vakilian et al. [Vak+12], that of the many automated refactorings existing in modern IDEs, the simplest ones are dominating the usage statistics. The report mainly focuses on Eclipse as the tool under investigation.

The paradigm is described almost as the opposite of automated composition of refactorings (see section 1.2.9 on page 24). It works by providing the programmer with easily accessible primitive refactorings. These refactorings shall be accessed via keyboard shortcuts or quick-assist menus¹ and be promptly executed, opposed to in the currently dominating wizard-based refactoring paradigm. They are meant to stimulate composing smaller refactorings into more complex changes, rather than doing a large upfront configuration of a wizard-based refactoring, before previewing and executing it. The compositional paradigm of refactoring is supposed to give control back to the programmer, by supporting him with an option of performing small rapid changes instead of large changes with a lesser degree of control. The report authors hope this will lead to fewer unsuccessful refactorings. It also could lower the bar for understanding the steps of a larger composite refactoring and thus also help in figuring out what goes wrong if one should choose to opt in on a wizard-based refactoring.

Vakilian and his associates have performed a survey of the effectiveness of the compositional paradigm versus the wizard-based one. They claim to have found evidence of that the *compositional paradigm* outperforms the *wizard-based*. It does so by reducing automation, which seem counterintuitive. Therefore they ask the question “What is an appropriate level of automation?”, and thus questions what they feel is a rush toward more automation in the software engineering community.

¹Think quick-assist with Ctrl+1 in Eclipse

Chapter 2

The search-based Extract and Move Method refactoring

In this chapter I will delve into the workings of the search-based *Extract and Move Method* refactoring. We will see the choices it must make along the way and why it chooses a text selection as a candidate for refactoring or not.

After defining some concepts, I will introduce an example that will be used throughout the chapter to illustrate how the refactoring works in some simple situations.

2.1 The inputs to the refactoring

For executing an *Extract and Move Method* refactoring, there are two simple requirements. The first thing the refactoring needs is a text selection, telling it what to extract. Its second requirement is a target for the subsequent move operation.

The extracted method must be called instead of the selection that makes up its body. Also, the method call has to be performed via a variable, since the method is not static. Therefore, the move target must be a variable in the scope of the extracted selection. The actual new location for the extracted method will be the class representing the type of the move target variable. But, since the method also must be called through a variable, it makes sense to define the move target to be either a local variable or a field in the scope of the text selection.

2.2 Defining a text selection

A text selection, in our context, is very similar to what you think of when selecting a bit of text in your editor or other text processing tool with your mouse or keyboard. It is an abstract construct that is meant to capture which specific portion of text we are about to deal with.

To be able to clearly reason about a text selection done to a portion of text in a computer file, that consist of pure text, we put up the following definition.

Definition. A *text selection* in a text file is defined by two non-negative integers, in addition to a reference to the file itself. The first integer is an offset into the file, while the second reference is the length of the text selection.

This means that the selected text consist of a number of characters equal to the length of the selection, where the first character is found at the specified offset.

2.3 Where we look for text selections

2.3.1 Text selections are found in methods

The text selections we are interested in are those that surrounds program statements. Therefore, the place we look for selections that can form candidates for an execution of the *Extract and Move Method* refactoring, is within the body of a single method.

On ignoring static methods In this project we are not analyzing static methods for candidates to the *Extract and Move Method* refactoring. The reason for this is that in the cases where we want to perform the refactoring for a selection within a static method, the first step is to extract the selection into a new method. Hence this method also become static, since it must be possible to call it from a static context. It would then be difficult to move the method to another class, make it non-static and calling it through a variable. To avoid these obstacles, we simply ignore static methods.

2.3.2 The possible text selections of a method body

The number of possible text selections that can be made from the text in a method body, are equal to all the sub-sequences of characters within it. For our purposes, analyzing program source code, we must define what it means for a text selection to be valid.

Definition. A *valid text selection* is a text selection that contains all of one or more consecutive program statements.

For a sequence of statements, the text selections that can be made from it, are equal to all its sub-sequences. Listing 9 on page 40 show an example of all the text selections that can be made from the code in listing 8 on the facing page, lines 16-18. For convenience and the clarity of this example, the text selections are represented as tuples with the start and end line of all selections: $\{(16), (17), (18), (16, 17), (16, 18), (17, 18)\}$.

Each nesting level of a method body can have many such sequences of statements. The outermost nesting level has one such sequence, and each branch contains their own sequence of statements. Listing 8 on the current

Clean	With statement sequences
<pre> 1 class C { 2 A a; B b; boolean bool; 3 4 void method(int val) { 5 if (bool) { 6 a.foo(); 7 a = new A(); 8 a.bar(); 9 } 10 11 a.foo(); 12 a.bar(); 13 14 switch (val) { 15 case 1: 16 b.a.foo(); 17 b.a.bar(); 18 break; 19 default: 20 a.foo(); 21 } 22 } 23 }</pre>	<pre> class C { A a; B b; boolean bool; void method(int val) { if (bool) { a.foo(); a = new A(); a.bar(); } a.foo(); a.bar(); switch (val) { case 1: b.a.foo(); b.a.bar(); break; default: a.foo(); } } }</pre>

Listing 8: Classes **A** and **B** are both public. The methods **foo** and **bar** are public members of class **A**.

page has a version of some code where all such sequences of statements are highlighted for a method body.

To complete our example of possible text selections, I will now list all possible text selections for the method in listing 8 on this page, by nesting level. There are 23 of them in total.

Level 1 (10 selections)

{(5, 9), (11), (12), (14, 21), (5, 11), (5, 12), (5, 21), (11, 12), (11, 21), (12, 21)}

Level 2 (13 selections)

{(6), (7), (8), (6, 7), (6, 8), (7, 8), (16), (17), (18), (16, 17), (16, 18), (17, 18), (20)}

The complexity

The complexity of how many text selections that needs to be analyzed for a body of in total n statements, is bounded by $O(n^2)$. A body of statements

```

16   b.a.foo();
17   b.a.bar();
18   break;

```

Listing 9: Example of how the text selections generator would generate text selections based on a lists of statements. Each highlighted rectangle represents a text selection.

is here all the statements in all nesting levels of a sequence of statements. A method body (or a block) is a body of statements. To prove that the complexity is bounded by $O(n^2)$, I present a couple of theorems and proves them.

Theorem. The number of text selections that need to be analyzed for each list of statements of length n , is exactly

$$\sum_{i=1}^n i = \frac{n(n+1)}{2}$$

Proof. For $n = 1$ this is trivial: $\frac{1(1+1)}{2} = \frac{2}{2} = 1$. One statement equals one selection.

For $n = 2$, you get one text selection for the first statement, one selection for the second statement, and one selection for the two of them combined. This equals three selections. $\frac{2(2+1)}{2} = \frac{6}{2} = 3$.

For $n = 3$, you get 3 selections for the two first statements, as in the case where $n = 2$. In addition you get one selection for the third statement itself, and two more statements for the combinations of it with the two previous statements. This equals six selections. $\frac{3(3+1)}{2} = \frac{12}{2} = 6$.

Assume that for $n = k$ there exists $\frac{k(k+1)}{2}$ text selections. Then we want to add selections for another statement, following the previous k statements. So, for $n = k + 1$, we get one additional selection for the statement itself. Then we get one selection for each pair of the new selection and the previous k statements. So the total number of selections will be the number of already generated selections, plus k for every pair, plus one for the statement itself: $\frac{k(k+1)}{2} + k + 1 = \frac{k(k+1)+2k+2}{2} = \frac{k(k+1)+2(k+1)}{2} = \frac{(k+1)(k+2)}{2} = \frac{(k+1)((k+1)+1)}{2} = \sum_{i=1}^{k+1} i$ \square

Theorem. The number of text selections for a body of statements is maximized if all the statements are at the same level.

Proof. Assume we have a body of, in total, k statements. Then, the sum of the lengths of all the lists of statements in the body, is also k . Let $\{l, \dots, m, (k - l - \dots - m)\}$ be the lengths of the lists of statements in the body, with $l + \dots + m < k \Rightarrow \forall i \in \{l, \dots, m\} : i < k$.

Then, the number of text selections that are generated for the k statements is

$$\begin{aligned} \frac{l(l+1)}{2} + \dots + \frac{m(m+1)}{2} + \frac{(k-l-\dots-m)((k-l-\dots-m)+1)}{2} &= \\ \frac{l^2+l}{2} + \dots + \frac{m^2+m}{2} + \frac{k^2-2kl-\dots-2km+l^2+\dots+m^2+k-l-\dots-m}{2} &= \\ \frac{2l^2-2kl+\dots+2m^2-2km+k^2+k}{2} & \end{aligned}$$

It then remains to show that this inequality holds:

$$\frac{2l^2-2kl+\dots+2m^2-2km+k^2+k}{2} < \frac{k(k+1)}{2} = \frac{k^2+k}{2}$$

By multiplication by 2 on both sides, and by removing the equal parts, we get

$$2l^2-2kl+\dots+2m^2-2km < 0$$

Since $\forall i \in \{l, \dots, m\} : i < k$, we have that $\forall i \in \{l, \dots, m\} : 2ki > 2i^2$, so all the pairs of parts on the form $2i^2 - 2ki$ are negative. In sum, the inequality holds. □

Therefore, the complexity for the number of selections that needs to be analyzed for a body of n statements is $O\left(\frac{n(n+1)}{2}\right) = O(n^2)$.

2.4 Disqualifying a selection

Certain text selections would lead to broken code if used as input to the *Extract and Move Method* refactoring. To avoid this, we have to check all text selections for such conditions before they are further analyzed. This section is therefore going to present some properties that make a selection unsuitable for our refactoring.

2.4.1 A call to a protected or package-private method

If a text selection contains a call to a protected or package-private method, it would not be safe to move it to another class. The reason for this, is that we cannot know if the called method is being overridden by some subclass of the enclosing class, or not.

Imagine that the protected method `foo` is declared in class *A*, and overridden in class *B*. The method `foo` is called from within a selection done to a method in *A*. We want to extract and move this selection to another class. The method `foo` is not public, so the *Move Method* refactoring must make it public, making the extracted method able to call it from the extracted method's new location. The problem is, that the now public method `foo` is overridden in a subclass, where it has a protected status.

This makes the compiler complain that the subclass *B* is trying to reduce the visibility of a method declared in its superclass *A*. This is not allowed in Java, and for good reasons. It would make it possible to make a subclass that could not be a substitute for its superclass.

The problem this check helps to avoid, is a little subtle. The problem does not arise in the class where the change is done, but in a class derived from it. This shows that classes acting as superclasses are especially fragile to introducing errors in the context of automated refactoring.

2.4.2 A double class instance creation

The following is a problem caused solely by the underlying *Move Method* refactoring. The problem occurs if two classes are instantiated such that the first constructor invocation is an argument to a second, and that the first constructor invocation takes an argument that is built up using a field. As an example, say that **name** is a field of the enclosing class, and we have the expression `new A(new B(name))`. If this expression is located in a selection that is moved to another class, **name** will be left untouched, instead of being prefixed with a variable of the same type as it is declared in. If **name** is the destination for the move, it is not replaced by **this**, or removed if it is a prefix to a member access (`name.member`), but it is still left by itself.

Situations like this would lead to code that will not compile. Therefore, we have to avoid them by not allowing selections to contain such double class instance creations that also contains references to fields.

2.4.3 Instantiation of non-static inner class

When a non-static inner class is instantiated, this must happen in the scope of its declaring class. This is because it must have access to the members of the declaring class. If the inner class is public, it is possible to instantiate it through an instance of its declaring class, but this is not handled by the underlying *Move Method* refactoring.

Performing a move on a method that instantiates a non-static inner class, will break the code if the instantiation is not handled properly. For this reason, selections that contains instantiations of non-static inner classes are deemed unsuitable for the *Extract and Move Method* refactoring.

2.4.4 References to enclosing instances of the enclosing class

The title of this section may be a little hard to grasp at first. What it means is that there is a (non-static) class *C* that is declared in the scope of possibly multiple other classes. And there is a statement in the body of a method declared in class *C*, that contains a reference to one or more instances of these enclosing classes of *C*.

The problem with this, is that these references may not be valid if they are moved to another class. Theoretically, some situations could easily be solved by passing, to the moved method, a reference to the instance where the problematic referenced member is declared. This should work in the case where this member is publicly accessible. This is not done in the underlying

Move Method refactoring, so it cannot be allowed in the *Extract and Move Method* refactoring either.

2.4.5 Inconsistent return statements

To verify that a text selection is consistent with respect to return statements, we must check that if a selection contains a return statement, then every possible execution path within the selection ends in either a return or a throw statement. This property is important regarding the *Extract Method* refactoring. If it holds, it means that a method could be extracted from the selection, and a call to it could be substituted for the selection. If the method has a non-void return type, then a call to it would also be a valid return point for the calling method. If its return value is of the void type, then the *Extract Method* refactoring will append an empty return statement to the back of the method call. Therefore, the analysis does not discriminate on either kinds of return statements, with or without a return value.

A throw statement is accepted anywhere a return statement is required. This is because a throw statement causes an immediate exit from the current block, together with all outer blocks in its control flow that does not catch the thrown exception.

Return statements can be either explicit or implicit. An *explicit* return statement is formed by using the **return** keyword, while an *implicit* return statement is a statement that is not formed using **return**, but must be the last statement of a method that can have any side effects. This can happen in methods with a void return type. An example is a statement that is inside one or more blocks. The last statement of a method could for instance be a synchronized statement, but the last statement that is executed in the method, and that can have any side effects, may be located inside the body of the synchronized statement.

We can start the check for this property by looking at the last statement of a selection to see if it is a return statement (explicit or implicit) or a throw statement. If this is the case, then the property holds, assuming the selected code does not contain any compilation errors. All execution paths within the selection should end in either this, or another, return or throw statement.

State somewhere that we assume no compilation errors?

If the last statement of the selection is not a return or throw, the execution of it must eventually end in one for the selection to be legal. This means that all branches of the last statement of every branch must end in a return or throw. Given this recursive definition, there are only five types of statements that are guaranteed to end in a return or throw if their child branches does. All other statements would have to be considered illegal. The first three: Block-statements, labeled statements and do-statements are all kinds of fall-through statements that always gets their body executed. Do-statements would not make much sense if written such that they always ends after the first round of execution of their body, but that is not our concern. The remaining two statements that can end in a return or throw are if-statements and try-statements.

For an if-statement, the rule is that if its then-part does not contain any return or throw statements, this is considered illegal. If the then-part does contain a return or throw, the else-part is checked. If its else-part is non-existent, or it does not contain any return or throw statements, the statement is considered illegal. If an if-statement is not considered illegal, the bodies of its two parts must be checked.

Try-statements are handled much the same way as if-statements. The body of a try-statement must contain a return or throw. The same applies to its catch clauses and finally body.

2.4.6 Ambiguous return values

The problem with ambiguous return values arise when a selection is chosen to be extracted into a new method, but it needs to return more than one value from that method.

This problem occurs in two situations. The first situation arise when there is more than one local variable that is both assigned to within a selection and also referenced after the selection. The other situation occur when there is only one such assignment, but the selection also contain return statements.

Therefore we must examine the selection for assignments to local variables that are referenced after the text selection. Then we must verify that not more than one such reference is done, or zero if any return statements are found.

2.4.7 Illegal statements

An illegal statement may be a statement that is of a type that is never allowed, or it may be a statement of a type that is only allowed if certain conditions are true.

Any use of the **super** keyword is prohibited, since its meaning is altered when moving a method to another class.

For a *break* statement, there are two situations to consider: A break statement with or without a label. If the break statement has a label, it is checked that whole of the labeled statement is inside the selection. If the break statement does not have a label attached to it, it is checked that its innermost enclosing loop or switch statement also is inside the selection.

The situation for a *continue* statement is the same as for a break statement, except that it is not allowed inside switch statements.

Regarding *assignments*, two types of assignments are allowed: Assignments to non-final variables and assignments to array access. All other assignments are regarded illegal.

Expand with more illegal statements and/or conclude that I did not have time to analyze all statement types.

2.5 Disqualifying selections from the example

Among the selections we found for the code in listing 8 on page 39, not many of them must be disqualified on the basis of containing something illegal. The only statement causing trouble is the break statement in line 18. None of the selections on nesting level 2 can contain this break statement, since the innermost switch statement is not inside any of these selections.

This means that the text selections (18), (16, 18) and (17, 18) can be excluded from further consideration, and we are left with the following selections.

Level 1 (10 selections)

{(5, 9), (11), (12), (14, 21), (5, 11), (5, 12), (5, 21), (11, 12), (11, 21), (12, 21)}

Level 2 (10 selections)

{(6), (7), (8), (6, 7), (6, 8), (7, 8), (16), (17), (16, 17), (20)}

2.6 Finding a move target

In the analysis needed to perform the *Extract and Move Method* refactoring automatically, the selection we choose is found among all the selections that has a possible move target. Therefore, the best possible move target must be found for all the candidate selections, so that we are able to sort out the selection that is best suited for the refactoring.

To find the best move target for a specific text selection, we first need to find all the possible targets. Since the target must be a local variable or a field, we are basically looking for names within the selection; names that represents references to variables.

The names we are looking for, we call prefixes. This is because we are not interested in names that occur in the middle of a dot-separated sequence of names. We are only interested in names that constitutes prefixes of other names, possibly themselves. The reason for this, is that two lexically equal names need not be referencing the same variable, if they themselves are not referenced via the same prefix. Consider the two method calls `a.x.foo()` and `b.x.foo()`. Here, the two references to `x`, in the middle of the qualified names both preceding `foo()`, are not referencing the same variable. Even though the variables may share the type, and the method `foo` thus is the same for both, we would not know through which of the variables `a` or `b` we should call the extracted method.

The possible move targets are then the prefixes that are not among a subset of the prefixes that are not valid move targets (see section 2.7 on the following page). Also, prefixes that are just simple names, and have only one occurrence, are left out. This is because they are not going to have any positive effect on coupling between classes, and are only going to increase the complexity of the code.

For finding the best move target among these safe prefixes, a simple heuristic is used. It is as simple as choosing the prefix that is most frequently

referenced within the selection.

2.7 Unfixes

The prefixes that are not valid as move targets are called unfixes.

An unfix can be a name that is assigned to within a selection. The reason that this cannot be allowed, is that the result would be an assignment to the **this** keyword, which is not valid in Java (see appendix A.1 on page 111).

Prefixes that originates from variable declarations within the same selection are also considered unfixes. This is because when a method is moved, it needs to be called through a variable. If this variable is also declared within the method that is to be moved, this obviously cannot be done.

Also considered as unfixes are variable references that are of types that are not suitable for moving methods to. This can either be because it is not physically possible to move a method to the desired class or that it will cause compilation errors by doing so.

If the type binding for a name is not resolved it is considered and unfix. The same applies to types that is only found in compiled code, so they have no underlying source that is accessible to us. (E.g. the `java.lang.String` class.)

Interfaces types are not suitable as targets. This is simply because interfaces in Java cannot contain methods with bodies. (This thesis does not deal with features of Java versions later than Java 7. Java 8 has interfaces with default implementations of methods.)

Neither are local types allowed. This accounts for both local and anonymous classes. Anonymous classes are effectively the same as interface types with respect to unfixes. Local classes could in theory be used as targets, but this is not possible due to limitations of the way the *Extract and Move Method* refactoring has to be implemented. The problem is that the refactoring is done in two steps, so the intermediate state between the two refactorings would not be legal Java code. In the intermediate step for the case where a local class is the move target, the extracted method would need to take the local class as a parameter. This new method would need to live in the scope of the declaring class of the originating method. The local class would then not be in the scope of the extracted method, thus bringing the source code into an illegal state. One could imagine that the method was extracted and moved in one operation, without an intermediate state. Then it would make sense to include variables with types of local classes in the set of legal targets, since the local classes would then be in the scopes of the method calls. If this makes any difference for software metrics that measure coupling would be a different discussion.

The last class of names that are considered unfixes are names used in null tests. These are tests that reads like this: if `<name>` equals `null` then do something. If allowing variables used in those kinds of expressions as targets for moving methods, we would end up with code containing boolean expressions like `this == null`, which would not be meaningful, since `this`

————— Before —————	————— After Extract Method —————
<pre> void declaresLocalClass() { class LocalClass { void foo() {} void bar() {} } LocalClass inst = new LocalClass(); inst.foo(); inst.bar(); } </pre>	<pre> void declaresLocalClass() { class LocalClass { void foo() {} void bar() {} } LocalClass inst = new LocalClass(); fooBar(inst); } // Intermediate step void fooBar(LocalClass inst) { inst.foo(); inst.bar(); } </pre>

Listing 10: When the *Extract and Move Method* tries to use a variable with a local type as the move target, an intermediate step is performed that is not allowed. Here: **LocalClass** is not in the scope of **fooBar** in its intermediate location.

would never be **null**.

2.8 Finding the example selections that have possible targets

We now pick up the thread from section 2.5 on page 45 where we have a set of text selections that needs to be analyzed to find out if some of them are suitable targets for the *Extract and Move Method* refactoring.

We start by analyzing the text selections for nesting level 2, because these results can be used to reason about the selections for nesting level 1. First we have all the single-statement selections.

Selections (6), (8) and (20).

All these selections have a prefix that contains a possible target, namely the variable **a**. The problem is that the prefixes are only one segment long, and their frequency counts are only 1 as well. None of these selections are therefore considered as suitable candidates for the refactoring.

Selection (7).

This selection contains the unfix **a**, and no other possible targets. The reason for **a** being an unfix is that it is assigned to within the selection. Selection (7) is therefore unsuited as a refactoring candidate.

Selections (16) and (17).

These selections both have a possible target. The target for both selections is the variable **b**. Both the prefixes have frequency 1. We denote this with the new tuples $((16), \mathbf{b.a}, f(1))$ and $((17), \mathbf{b.a}, f(1))$. They contain the selection, the prefix with the target and the frequency for this prefix.

Then we have all the text selections from level 2 that are composed of multiple statements:

Selections (6, 7), (6, 8) and (7, 8).

All these selections are disqualified for the reason that they contain the unfix **a**, due to the assignment, and no other possible move targets.

Selection (16, 17).

This selection is the last selection that is not analyzed on nesting level 2. It contains only one possible move target, and that is the variable **b**. It also contains only one prefix **b.a**, with frequency count 2. Therefore we have a new candidate $((16, 17), \mathbf{b.a}, f(2))$.

Moving on to the text selections for nesting level 1, starting with the single-statement selections:

Selection (5, 9).

This selection contains two variable references that must be analyzed to see if they are possible move candidates. The first one is the variable **bool**. This variable is of type **boolean**, that is a primary type and therefore not possible to make any changes to. The second variable is **a**. The variable **a** is an unfix in (5, 9), for the same reason as in the selections (6, 7), (7, 8) and (6, 8). So selection (5, 9) contains no possible move targets.

Selections (11) and (12).

These selections are disqualified for the same reasons as selections (6) and (8). Their prefixes are one segment long and are referenced only one time.

Selection (14, 21)

This is the switch statement from listing 8 on page 39. It contains the relevant variable references **val**, **a** and **b**. The variable **val** is a primary type, just as **bool**. The variable **a** is only found in one statement, and in a prefix with only one segment, so it is not considered to be a possible move target. The only variable left is **b**. Just as in the selection (16, 17), **b** is part of the prefix **b.a**, that has 2 appearances. We have found a new candidate $((14, 21), \mathbf{b.a}, f(2))$.

It remains to see if we get any new candidates by analyzing the multi-statement text selections for nesting level 1:

Selections (5, 11) and (5, 12).

These selections are disqualified for the same reason as (5, 9). The only possible move target **a** is an unfix.

Selection (5, 21).

This is whole of the method body in listing 8 on page 39. The variables **a**, **bool** and **val** are either an unfix or primary types. The variable **b** is the only possible move target, and we have, again, the prefix **b.a** as the center of attention. The refactoring candidate is ((5, 21), **b.a**, $f(2)$).

Selection (11, 12).

This small selection contains the prefix **a** with frequency 2, and no unfixes. The candidate is ((11, 12), **a**, $f(2)$).

Selection (11, 21)

This selection contains the selection (11, 12) in addition to the switch statement. The selection has two possible move targets. The first one is **b**, in a prefix with frequency 2. The second is **a**, although it is in a simple prefix, it is referenced 3 times, and is therefore chosen as the target for the selection. The new candidate is ((11, 21), **a**, $f(3)$).

Selection (12, 21).

This selection is similar to the previous (11, 21), only that **a** now has frequency count 2. This means that the prefixes **a** and **b.a** both have the count 2. Of the two, **b.a** is preferred, since it has more segments than **a**. Thus the candidate for this selection is ((12, 21), **b.a**, $f(2)$).

For the method in listing 8 on page 39 we therefore have the following 8 candidates for the *Extract and Move Method* refactoring: $\{((16), \mathbf{b.a}, f(1)), ((17), \mathbf{b.a}, f(1)), ((16, 17), \mathbf{b.a}, f(2)), ((14, 21), \mathbf{b.a}, f(2)), ((5, 21), \mathbf{b.a}, f(2)), ((11, 12), \mathbf{a}, f(2)), ((11, 21), \mathbf{a}, f(3)), ((12, 21), \mathbf{b.a}, f(2))\}$.

It now only remains to specify an order for these candidates, so we can choose the most suitable one to refactor.

2.9 Choosing the selection

When choosing a selection between the text selections that have possible move targets, the selections need to be ordered. The criteria below are presented in the order they are prioritized. If not one selection is favored over the other for a concrete criterion, the selections are evaluated by the next criterion.

1. The first criterion that is evaluated is whether a selection contains any unfixes or not. If selection *A* contains no unfixes, while selection *B* does, selection *A* is favored over selection *B*. This is because, if we can, we want to avoid moving such as assignments and variable declarations. This is done under the assumption that, if possible, avoiding selections containing unfixes will make the code moved a little cleaner.

2. The second criterion that is evaluated is whether a selection contains multiple possible targets or not. If selection *A* has only one possible target, and selection *B* has multiple, selection *A* is favored. If both selections have multiple possible targets, they are considered equal with respect to this criterion. The rationale for this heuristic is that we would prefer not to introduce new couplings between classes when performing the *Extract and Move Method* refactoring.
3. When evaluating this criterion, this is with the knowledge that selection *A* and *B* both have one possible target, or multiple possible targets. Then, if the move target candidate of selection *A* has a higher reference count than the target candidate of selection *B*, selection *A* is favored. The reason for this is that we would like to move the selection that gets rid of the most references to another class.
4. The last criterion is that if the frequencies of the targets chosen for both selections are equal, the selection with the target that is part of the prefix with highest number of segments is favored. This is done to favor indirection.

If none of the above mentioned criteria favors one selection over another, the selections are considered to be equally good candidates for the *Extract and Move Method* refactoring.

2.10 Concluding the example

For choosing one of the remaining selections, we need to order our candidates after the criteria in the previous section. Below we have the candidates ordered in descending order, with the “best” candidate first:

- | | |
|--|--|
| 1. ((16, 17), b.a , <i>f</i> (2)) | 5. ((11, 21), a , <i>f</i> (3)) |
| 2. ((11, 12), a , <i>f</i> (2)) | 6. ((5, 21), b.a , <i>f</i> (2)) |
| 3. ((16), b.a , <i>f</i> (1)) | 7. ((12, 21), b.a , <i>f</i> (2)) |
| 4. ((17), b.a , <i>f</i> (1)) | 8. ((14, 21), b.a , <i>f</i> (2)) |

The candidates ordered 5-8 all have unfixes in them, therefore they are ordered last. None of the candidates ordered 1-4 have multiple possible targets. The only interesting discussion is now why (16,17) was favored over (11,12). This is because the prefix **b.a** enclosing the move target of selection (16,17) has one more segment than the prefix **a** of (11,12).

The selection is now extracted into a new method **gen_123** and then moved to class **B**, since that is the type of the variable **b** that is the target from the chosen refactoring candidate. The name of the method has a randomly generated suffix. In the refactoring implementation, the extracted methods follow the pattern **generated_<long>**, where **<long>** is a pseudo-random long value. This is shortened here to make the example readable. The result after the refactoring is shown in listing 11 on the facing page.

```

1  class C {
2      A a; B b; boolean bool;
3
4      void method(int val) {
5          if (bool) {
6              a.foo();
7              a = new A();
8              a.bar();
9          }
10
11         a.foo();
12         a.bar();
13
14         switch (val) {
15             case 1:
16                 b.gen_123(this);
17                 break;
18             default:
19                 a.foo();
20         }
21     }
22 }

```

```

public class B {
    A a;

    public void gen_123(C c) {
        a.foo();
        a.bar();
    }
}

```

Listing 11: The result after refactoring the class **C** of listing 8 on page 39 with the *Extract and Move Method* refactoring with $((16, 17), b.a, f(2))$ as input.

2.11 ??

Short description of changing code and error handling?

Pointing to implementation chapter

Chapter 3

Refactorings in Eclipse JDT: Design and Shortcomings

This chapter will deal with some of the design behind refactoring support in Eclipse, and the JDT in specific. After which it will follow a section about shortcomings of the refactoring API in terms of composition of refactorings.

3.1 Design

The refactoring world of Eclipse can in general be separated into two parts: The language independent part and the part written for a specific programming language – the language that is the target of the supported refactorings.

What about the language specific part?

3.1.1 The Language Toolkit

The Language Toolkit¹, or LTK for short, is the framework that is used to implement refactorings in Eclipse. It is language independent and provides the abstractions of a refactoring and the change it generates, in the form of the classes **Refactoring**² and **Change**³.

There are also parts of the LTK that is concerned with user interaction, but they will not be discussed here, since they are of little value to us and our use of the framework. We are primarily interested in the parts that can be automated.

The Refactoring Class

The abstract class **Refactoring** is the core of the LTK framework. Every refactoring that is going to be supported by the LTK have to end up creating an instance of one of its subclasses. The main responsibilities of subclasses

¹The content of this section is a mixture of written material from <https://www.eclipse.org/articles/Article-LTK/ltk.html> and <http://www.eclipse.org/articles/article.php?file=Article-Unleashing-the-Power-of-Refactoring/index.html>, the LTK source code and my own memory.

²`org.eclipse.ltk.core.refactoring.Refactoring`

³`org.eclipse.ltk.core.refactoring.Change`

of **Refactoring** is to implement template methods for condition checking (**checkInitialConditions**¹ and **checkFinalConditions**²), in addition to the **createChange**³ method that creates and returns an instance of the **Change** class.

If the refactoring shall support that others participate in it when it is executed, the refactoring has to be a processor-based refactoring⁴. It then delegates to its given **RefactoringProcessor**⁵ for condition checking and change creation. Participating in a refactoring can be useful in cases where the changes done to programming source code affects other related resources in the workspace. This can be names or paths in configuration files, or maybe one would like to perform additional logging of changes done in the workspace.

The Change Class

This class is the base class for objects that is responsible for performing the actual workspace transformations in a refactoring. The main responsibilities for its subclasses is to implement the **perform**⁶ and **isValid**⁷ methods. The **isValid** method verifies that the change object is valid and thus can be executed by calling its **perform** method. The **perform** method performs the desired change and returns an undo change that can be executed to reverse the effect of the transformation done by its originating change object.

Executing a Refactoring

The life cycle of a refactoring generally follows two steps after creation: condition checking and change creation. By letting the refactoring object be handled by a **CheckConditionsOperation**⁸ that in turn is handled by a **CreateChangeOperation**⁹, it is assured that the change creation process is managed in a proper manner.

The actual execution of a change object has to follow a detailed life cycle. This life cycle is honored if the **CreateChangeOperation** is handled by a **PerformChangeOperation**¹⁰. If also an undo manager¹¹ is set for the **PerformChangeOperation**, the undo change is added into the undo history.

¹`org.eclipse.ltk.core.refactoring.Refactoring#checkInitialConditions()`

²`org.eclipse.ltk.core.refactoring.Refactoring#checkFinalConditions()`

³`org.eclipse.ltk.core.refactoring.Refactoring#createChange()`

⁴`org.eclipse.ltk.core.refactoring.participants.ProcessorBasedRefactoring`

⁵`org.eclipse.ltk.core.refactoring.participants.RefactoringProcessor`

⁶`org.eclipse.ltk.core.refactoring.Change#perform()`

⁷`org.eclipse.ltk.core.refactoring.Change#isValid()`

⁸`org.eclipse.ltk.core.refactoring.CheckConditionsOperation`

⁹`org.eclipse.ltk.core.refactoring.CreateChangeOperation`

¹⁰`org.eclipse.ltk.core.refactoring.PerformChangeOperation`

¹¹`org.eclipse.ltk.core.refactoring.IUndoManager`

3.2 Shortcomings

This section is introduced naturally with a conclusion: The JDT refactoring implementation does not facilitate composition of refactorings. This section will try to explain why, and also identify other shortcomings of both the usability and the readability of the JDT refactoring source code.

refine

I will begin at the end and work my way toward the composition part of this section.

3.2.1 Absence of Generics in Eclipse Source Code

This section is not only concerning the JDT refactoring API, but also large quantities of the Eclipse source code. The code shows a striking absence of the Java language feature of generics. It is hard to read a class' interface when methods return objects or takes parameters of raw types such as **List** or **Map**. This sometimes results in having to read a lot of source code to understand what is going on, instead of relying on the available interfaces. In addition, it results in a lot of ugly code, making the use of typecasting more of a rule than an exception.

3.2.2 Composite Refactorings Will Not Appear as Atomic Actions

Missing Flexibility from JDT Refactorings

The JDT refactorings are not made with composition of refactorings in mind. When a JDT refactoring is executed, it assumes that all conditions for it to be applied successfully can be found by reading source files that have been persisted to disk. They can only operate on the actual source material, and not (in-memory) copies thereof. This constitutes a major disadvantage when trying to compose refactorings, since if an exception occurs in the middle of a sequence of refactorings, it can leave the project in a state where the composite refactoring was only partially executed. It makes it hard to discard the changes done without monitoring and consulting the undo manager, an approach that is not bullet proof.

Broken Undo History

When designing a composed refactoring that is to be performed as a sequence of refactorings, you would like it to appear as a single change to the workspace. This implies that you would also like to be able to undo all the changes done by the refactoring in a single step. This is not the way it appears when a sequence of JDT refactorings is executed. It leaves the undo history filled up with individual undo actions corresponding to every single JDT refactoring in the sequence. This problem is not trivial to handle in Eclipse (see section 4.2.6 on page 63).

Chapter 4

Composite Refactorings in Eclipse

4.1 A Simple Ad Hoc Model

As pointed out in chapter 3 on page 53, the Eclipse JDT refactoring model is not very well suited for making composite refactorings. Therefore a simple model using changer objects (of type **RefaktorChanger**) is used as an abstraction layer on top of the existing Eclipse refactorings, instead of extending the **Refactoring**¹ class.

The use of an additional abstraction layer is a deliberate choice. It is due to the problem of creating a composite **Change**² that can handle text changes that interfere with each other. Thus, a **RefaktorChanger** may, or may not, take advantage of one or more existing refactorings, but it is always intended to make a change to the workspace.

4.1.1 A typical RefaktorChanger

The typical refaktor changer class has two responsibilities, checking preconditions and executing the requested changes. This is not too different from the responsibilities of an LTK refactoring, with the distinction that a refaktor changer also executes the change, while an LTK refactoring is only responsible for creating the object that can later be used to do the job.

Checking of preconditions is typically done by an **Analyzer**³. If the preconditions validate, the upcoming changes are executed by an **Executor**⁴.

4.2 The Extract and Move Method Refactoring

4.2.1 The Building Blocks

This is a composite refactoring, and hence is built up using several primitive refactorings. These basic building blocks are, as its name implies, the *Extract*

¹`org.eclipse.ltk.core.refactoring.Refactoring`

²`org.eclipse.ltk.core.refactoring.Change`

³`no.uio.ifi.refaktor.analyze.analyzers.Analyzer`

⁴`no.uio.ifi.refaktor.change.executors.Executor`

Method refactoring [Fow99] and the *Move Method* refactoring [Fow99]. In Eclipse, the implementations of these refactorings are found in the classes **ExtractMethodRefactoring**¹ and **MoveInstanceMethodProcessor**², where the last class is designed to be used together with the processor-based **MoveRefactoring**³.

The ExtractMethodRefactoring Class

This class is quite simple in its use. The only parameters it requires for construction is a compilation unit⁴, the offset into the source code where the extraction shall start, and the length of the source to be extracted. Then you have to set the method name for the new method together with its visibility and some not so interesting parameters.

The MoveInstanceMethodProcessor Class

For the *Move Method*, the processor requires a little more advanced input than the class for the *Extract Method*. For construction it requires a method handle⁵ for the method that is to be moved. Then the target for the move have to be supplied as the variable binding from a chosen variable declaration. In addition to this, one have to set some parameters regarding setters/getters, as well as delegation.

To make a working refactoring from the processor, one have to create a **MoveRefactoring** with it.

4.2.2 The ExtractAndMoveMethodChanger

The **ExtractAndMoveMethodChanger**⁶ class is a subclass of the class **RefaktorChanger**⁷. It is responsible for analyzing and finding the best target for, and also executing, a composition of the *Extract Method* and *Move Method* refactorings. This particular changer is the one of my changers that is closest to being a true LTK refactoring. It can be reworked to be one if the problems with overlapping changes are resolved. The changer requires a text selection and the name of the new method, or else a method name will be generated. The selection has to be of the type **CompilationUnitTextSelection**⁸. This class is a custom extension to **TextSelection**⁹, that in addition to the basic offset, length and similar methods, also carry an instance of the underlying compilation unit handle for the selection.

¹`org.eclipse.jdt.internal.corext.refactoring.code.ExtractMethodRefactoring`

²`org.eclipse.jdt.internal.corext.refactoring.structure.MoveInstanceMethodProcessor`

³`org.eclipse.ltk.core.refactoring.participants.MoveRefactoring`

⁴`org.eclipse.jdt.core.ICompilationUnit`

⁵`org.eclipse.jdt.core.IMethod`

⁶`no.uio.ifi.refaktor.changers.ExtractAndMoveMethodChanger`

⁷`no.uio.ifi.refaktor.changers.RefaktorChanger`

⁸`no.uio.ifi.refaktor.utils.CompilationUnitTextSelection`

⁹`org.eclipse.jface.text.TextSelection`

The ExtractAndMoveMethodAnalyzer

The analysis and precondition checking is done by the **ExtractAndMoveMethodAnalyzer**¹. First is check whether the selection is a valid selection or not, with respect to statement boundaries and that it actually contains any selections. Then it checks the legality of both extracting the selection and also moving it to another class. This checking of is performed by a range of checkers (see section 5.5 on page 72). If the selection is approved as legal, it is analyzed to find the presumably best target to move the extracted method to.

For finding the best suitable target the analyzer is using a **PrefixesCollector**² that collects all the possible candidate targets for the refactoring. All the non-candidates is found by an **UnfixesCollector**³ that collects all the targets that will give some kind of error if used. (For details about the property collectors, see section 5.4 on page 71.) All prefixes (and unfixes) are represented by a **Prefix**⁴, and they are collected into sets of prefixes. The safe prefixes is found by subtracting from the set of candidate prefixes the prefixes that is enclosing any of the unfixes. A prefix is enclosing an unfix if the unfix is in the set of its sub-prefixes. As an example, "a.b" is enclosing "a", as is "a". The safe prefixes is unified in a **PrefixSet**. If a prefix has only one occurrence, and is a simple expression, it is considered unsuitable as a move target. This occurs in statements such as "a.foo()". For such statements it bares no meaning to extract and move them. It only generates an extra method and the calling of it.

The most suitable target for the refactoring is found by finding the prefix with the most occurrences. If two prefixes have the same occurrence count, but they differ in the number of segments, the one with most segments is chosen.

The ExtractAndMoveMethodExecutor

If the analysis finds a possible target for the composite refactoring, it is executed by an **ExtractAndMoveMethodExecutor**⁵. It is composed of the two executors known as **ExtractMethodRefactoringExecutor**⁶ and **MoveMethodRefactoringExecutor**⁷. The **ExtractAndMoveMethodExecutor** is responsible for gluing the two together by feeding the **MoveMethodRefactoringExecutor** with the resources needed after executing the extract method refactoring.

¹no.uio.ifi.refaktor.analyze.analyzers.ExtractAndMoveMethodAnalyzer

²no.uio.ifi.refaktor.analyze.collectors.PrefixesCollector

³no.uio.ifi.refaktor.analyze.collectors.UnfixesCollector

⁴no.uio.ifi.refaktor.extractors.Prefix

⁵no.uio.ifi.refaktor.change.executors.ExtractAndMoveMethodExecutor

⁶no.uio.ifi.refaktor.change.executors.ExtractMethodRefactoringExecutor

⁷no.uio.ifi.refaktor.change.executors.MoveMethodRefactoringExecutor

The ExtractMethodRefactoringExecutor

This executor is responsible for creating and executing an instance of the **ExtractMethodRefactoring** class. It is also responsible for collecting some post execution resources that can be used to find the method handle for the extracted method, as well as information about its parameters, including the variable they originated from.

The MoveMethodRefactoringExecutor

This executor is responsible for creating and executing an instance of the **MoveRefactoring**. The move refactoring is a processor-based refactoring, and for the *Move Method* refactoring it is the **MoveInstanceMethodProcessor** that is used.

The handle for the method to be moved is found on the basis of the information gathered after the execution of the *Extract Method* refactoring. The only information the **ExtractMethodRefactoring** is sharing after its execution, regarding find the method handle, is the textual representation of the new method signature. Therefore it must be parsed, the strings for types of the parameters must be found and translated to a form that can be used to look up the method handle from its type handle. They have to be on the unresolved form. The name for the type is found from the original selection, since an extracted method must end up in the same type as the originating method.

Elaborate?

When analyzing a selection prior to performing the *Extract Method* refactoring, a target is chosen. It has to be a variable binding, so it is either a field or a local variable/parameter. If the target is a field, it can be used with the **MoveInstanceMethodProcessor** as it is, since the extracted method still is in its scope. But if the target is local to the originating method, the target that is to be used for the processor must be among its parameters. Thus the target must be found among the extracted method's parameters. This is done by finding the parameter information object that corresponds to the parameter that was declared on basis of the original target's variable when the method was extracted. (The extracted method must take one such parameter for each local variable that is declared outside the selection that is extracted.) To match the original target with the correct parameter information object, the key for the information object is compared to the key from the original target's binding. The source code must then be parsed to find the method declaration for the extracted method. The new target must be found by searching through the parameters of the declaration and choose the one that has the same type as the old binding from the parameter information object, as well as the same name that is provided by the parameter information object.

4.2.3 The SearchBasedExtractAndMoveMethodChanger

The **SearchBasedExtractAndMoveMethodChanger**¹ is a changer whose purpose is to automatically analyze a method, and execute the *Extract and Move Method* refactoring on it if it is a suitable candidate for the refactoring.

First, the **SearchBasedExtractAndMoveMethodAnalyzer**² is used to analyze the method. If the method is found to be a candidate, the result from the analysis is fed to the **ExtractAndMoveMethodExecutor**, whose job is to execute the refactoring (see section 4.2.2 on page 59).

The SearchBasedExtractAndMoveMethodAnalyzer

This analyzer is responsible for analyzing all the possible text selections of a method and then choose the best result out of the analysis results that is, by the analyzer, considered to be the potential candidates for the Extract and Move Method refactoring.

Before the analyzer is able to work with the text selections of a method, it needs to generate them. To do this, it parses the method to obtain a **MethodDeclaration** for it (see section 5.2.1 on page 67). Then there is a statement lists creator that creates statements lists of the different groups of statements in the body of the method declaration. A text selections generator generates text selections of all the statement lists for the analyzer to work with.

The statement lists creator is responsible for generating lists of statements for all the possible nesting levels of statements in the method. The statement lists creator is implemented as an AST visitor (see section 5.3 on page 69). It generates lists of statements by visiting all the blocks in the method declaration and stores their statements in a collection of statement lists. In addition, it visits all of the other statements that can have a statement as a child, such as the different control structures and the labeled statement.

The switch statement is the only kind of statement that is not straight forward to obtain the child statements from. It stores all of its children in a flat list. Its switch case statements are included in this list. This means that there are potential statement lists between all of these case statements. The list of statements from a switch statement is therefore traversed, and the statements between the case statements are grouped as separate lists.

The highlighted part of listing 8 on page 39 shows an example of how the statement lists creator would separate a method body into lists of statements.

The text selections generator generates text selections for each list of statements from the statement lists creator. The generator generates a text selection for every sub-sequence of statements in a list. For a sequence

¹`no.uio.ifi.refaktor.change.changers.SearchBasedExtractAndMoveMethodChanger`

²`no.uio.ifi.refaktor.analyze.analyzers.SearchBasedExtractAndMoveMethodAnalyzer`

of statements, the first statement and the last statement span out a text selection.

In practice, the text selections are calculated by only one traversal of the statement list. There is a set of generated text selections. For each statement, there is created a temporary set of selections, in addition to a text selection based on the offset and length of the statement. This text selection is added to the temporary set. Then the new selection is added with every selection from the set of generated text selections. These new selections are added to the temporary set. Then the temporary set of selections is added to the set of generated text selections. The result of adding two text selections is a new text selection spanned out by the two addends.

Finding the candidate for the refactoring is done by analyzing all the generated text selection with the **ExtractAndMoveMethodAnalyzer** (see section 4.2.2 on page 59). If the analyzer generates a useful result, an **ExtractAndMoveMethodCandidate** is created from it, that is kept in a list of potential candidates. If no candidates are found, the **NoTargetFoundException** is thrown.

Since only one of the candidates can be chosen, the analyzer must sort out which candidate to choose. The sorting is done by the static **sort** method of **Collections**. The comparison in this sorting is done by an **ExtractAndMoveMethodCandidateComparator**.

Write about the **ExtractAndMoveMethodCandidateComparator/FavorNoUnfixesCandidateComparator**

4.2.4 The Prefix Class

This class exists mainly for holding data about a prefix, such as the expression that the prefix represents and the occurrence count of the prefix within a selection. In addition to this, it has some functionality such as calculating its sub-prefixes and intersecting it with another prefix. The definition of the intersection between two prefixes is a prefix representing the longest common expression between the two.

4.2.5 The PrefixSet Class

A prefix set holds elements of type **Prefix**. It is implemented with the help of a **HashMap**¹ and contains some typical set operations, but it does not implement the **Set**² interface, since the prefix set does not need all of the functionality a **Set** requires to be implemented. In addition It needs some other functionality not found in the **Set** interface. So due to the relatively limited use of prefix sets, and that it almost always needs to be referenced as such, and not a **Set<Prefix>**, it remains as an ad hoc solution to a concrete problem.

¹`java.util.HashMap`

²`java.util.Set`

There are two ways adding prefixes to a **PrefixSet**. The first is through its **add** method. This works like one would expect from a set. It adds the prefix to the set if it does not already contain the prefix. The other way is to *register* the prefix with the set. When registering a prefix, if the set does not contain the prefix, it is just added. If the set contains the prefix, its count gets incremented. This is how the occurrence count is handled.

The prefix set also computes the set of prefixes that is not enclosing any prefixes of another set. This is kind of a set difference operation only for enclosing prefixes.

4.2.6 Hacking the Refactoring Undo History

Where to put this section?

As an attempt to make multiple subsequent changes to the workspace appear as a single action (i.e. make the undo changes appear as such), I tried to alter the undo changes¹ in the history of the refactorings.

My first impulse was to remove the, in this case, last two undo changes from the undo manager² for the Eclipse refactorings, and then add them to a composite change³ that could be added back to the manager. The interface of the undo manager does not offer a way to remove/pop the last added undo change, so a possible solution could be to decorate [Gam+95] the undo manager, to intercept and collect the undo changes before delegating to the **addUndo** method⁴ of the manager. Instead of giving it the intended undo change, a null change could be given to prevent it from making any changes if run. Then one could let the collected undo changes form a composite change to be added to the manager.

There is a technical challenge with this approach, and it relates to the undo manager, and the concrete implementation **UndoManager2**⁵. This implementation is designed in a way that it is not possible to just add an undo change, you have to do it in the context of an active operation⁶. One could imagine that it might be possible to trick the undo manager into believing that you are doing a real change, by executing a refactoring that is returning a kind of null change that is returning our composite change of undo refactorings when it is performed. But this is not the way to go.

Apart from the technical problems with this solution, there is a functional problem: If it all had worked out as planned, this would leave the undo history in a dirty state, with multiple empty undo operations corresponding to each of the sequentially executed refactoring operations, followed by a composite undo change corresponding to an empty change of the workspace for rounding of our composite refactoring. The solution to this particular problem could be to intercept the registration of the intermediate changes in the undo manager, and only register the last empty change.

¹`org.eclipse.ltk.core.refactoring.Change`

²`org.eclipse.ltk.core.refactoring.IUndoManager`

³`org.eclipse.ltk.core.refactoring.CompositeChange`

⁴`org.eclipse.ltk.core.refactoring.IUndoManager#addUndo()`

⁵`org.eclipse.ltk.internal.core.refactoring.UndoManager2`

⁶`org.eclipse.core.commands.operations.TriggeredOperations`

Unfortunately, not everything works as desired with this solution. The grouping of the undo changes into the composite change does not make the undo operation appear as an atomic operation. The undo operation is still split up into separate undo actions, corresponding to the change done by its originating refactoring. And in addition, the undo actions has to be performed separate in all the editors involved. This makes it no solution at all, but a step toward something worse.

There might be a solution to this problem, but it remains to be found. The design of the refactoring undo management is partly to be blamed for this, as it it is to complex to be easily manipulated.

Chapter 5

Analyzing Source Code in Eclipse

5.1 The Java model

The Java model of Eclipse is its internal representation of a Java project. It is light-weight, and has only limited possibilities for manipulating source code. It is typically used as a basis for the Package Explorer in Eclipse.

The elements of the Java model is only handles to the underlying elements. This means that the underlying element of a handle does not need to actually exist. Hence the user of a handle must always check that it exist by calling the **exists** method of the handle.

The handles with descriptions is listed in table 5.1 on the following page, while the hierarchy of the Java Model is shown in fig. 5.1 on page 67.

5.2 The Abstract Syntax Tree

Eclipse is following the common paradigm of using an abstract syntax tree for source code analysis and manipulation.

When parsing program source code into something that can be used as a foundation for analysis, the start of the process follows the same steps as in a compiler. This is all natural, because the way a compiler analyzes code is no different from how source manipulation programs would do it, except for some properties of code that is analyzed in the parser, and that they may be differing in what kinds of properties they analyze. Thus the process of translation source code into a structure that is suitable for analyzing, can be seen as a kind of interrupted compilation process (see fig. 5.2 on page 68).

The process starts with a *scanner*, or lexer. The job of the scanner is to read the source code and divide it into tokens for the parser. Therefore, it is also sometimes called a tokenizer. A token is a logical unit, defined in the language specification, consisting of one or more consecutive characters. In the Java language the tokens can for instance be the **this** keyword, a curly bracket **{** or a **nameToken**. It is recognized by the scanner on the basis of something equivalent of a regular expression. This part of the process is often implemented with the use of a finite automata. In fact, it is common

Table 5.1: The elements of the Java Model [Vog12].

Project Element	Java Model element	Description
Java project	IJavaProject	The Java project which contains all other objects.
Source folder / binary folder / external library	IPackageFragmentRoot	Hold source or binary files, can be a folder or a library (zip / jar file).
Each package	IPackageFragment	Each package is below the IPackageFragmentRoot , sub-packages are not leaves of the package, they are listed directed under IPackageFragmentRoot .
Java Source file	ICompilationUnit	The Source file is always below the package node.
Types / Fields / Methods	IType / IField / IMethod	Types, fields and methods.

to specify the tokens in regular expressions, that in turn is translated into a finite automata lexer. This process can be automated.

The program component used to translate a stream of tokens into something meaningful, is called a parser. A parser is fed tokens from the scanner and performs an analysis of the structure of a program. It verifies that the syntax is correct according to the grammar rules of a language, that is usually specified in a context-free grammar, and often in a variant of the Backus–Naur Form¹. The result coming from the parser is in the form of an *Abstract Syntax Tree*, AST for short. It is called *abstract*, because the structure does not contain all of the tokens produced by the scanner. It only contain logical constructs, and because it forms a tree, all kinds of parentheses and brackets are implicit in the structure. It is this AST that is used when performing the semantic analysis of the code.

As an example we can think of the expression $(5 + 7) * 2$. The root of this tree would in Eclipse be an **InfixExpression** with the operator **TIMES**, and a left operand that is also an **InfixExpression** with the operator **PLUS**. The left operand **InfixExpression**, has in turn a left operand of type **NumberLiteral** with the value "5" and a right operand **NumberLiteral** with the value "7". The root will have a right operand of type **NumberLiteral** and value "2". The AST for this expression is illustrated in fig. 5.3 on page 68.

Contrary to the Java Model, an abstract syntax tree is a heavy-weight

¹https://en.wikipedia.org/wiki/Backus-Naur_Form




Figure 5.1: The Java model of Eclipse. “{ **SomeElement** }*” means “**SomeElement** zero or more times“. For recursive structures, “...” is used.

representation of source code. It contains information about properties like type bindings for variables and variable bindings for names.

5.2.1 The AST in Eclipse

In Eclipse, every node in the AST is a child of the abstract superclass **ASTNode**¹. Every **ASTNode**, among a lot of other things, provides information about its position and length in the source code, as well as a reference to its parent and to the root of the tree.

The root of the AST is always of type **CompilationUnit**. It is not the same as an instance of an **ICompilationUnit**, which is the compilation unit handle of the Java model. The children of a **CompilationUnit** is an optional **PackageDeclaration**, zero or more nodes of type **ImportDeclaration** and all its top-level type declarations that has node types **AbstractTypeDeclaration**.

An **AbstractTypeDeclaration** can be one of the types **AnnotationTypeDeclaration**, **EnumDeclaration** or **TypeDeclaration**. The children of an **AbstractTypeDeclaration** must be a subtype of a **BodyDeclaration**.

¹org.eclipse.jdt.core.dom.ASTNode

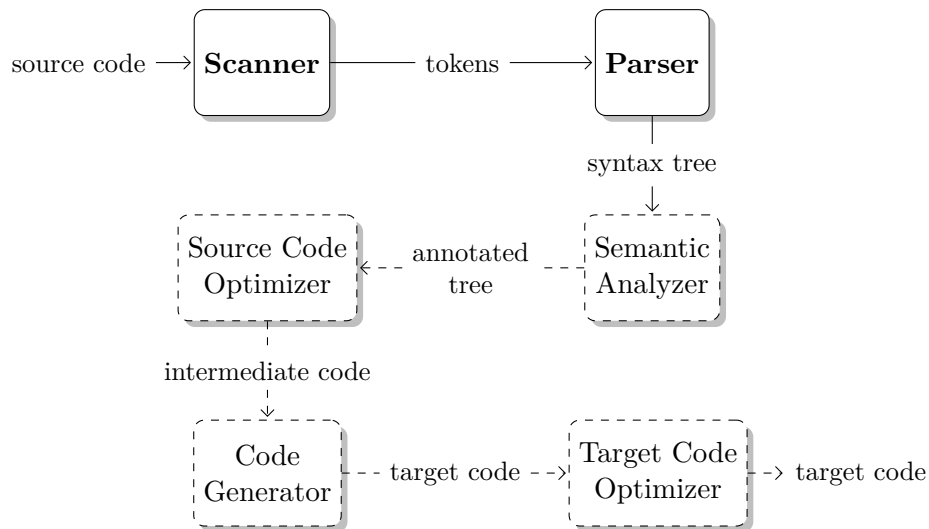


Figure 5.2: Interrupted compilation process. (Full compilation process borrowed from *Compiler construction: principles and practice* by Kenneth C. Louden [Lou97].)




Figure 5.3: The abstract syntax tree for the expression $(5 + 7) * 2$.

These subtypes are: **AnnotationTypeMemberDeclaration**, **EnumConstantDeclaration**, **FieldDeclaration**, **Initializer** and **MethodDeclaration**.

Of the body declarations, the **MethodDeclaration** is the most interesting one. Its children include lists of modifiers, type parameters, parameters and exceptions. It has a return type node and a body node. The body, if present, is of type **Block**. A **Block** is itself a **Statement**, and its children is a list of **Statement** nodes.

There are too many types of the abstract type **Statement** to list up, but there exists a subtype of **Statement** for every statement type of Java, as one would expect. This also applies to the abstract type **Expression**. However, the expression **Name** is a little special, since it is both used as an operand in compound expressions, as well as for names in type declarations and such.

There is an overview of some of the structure of an Eclipse AST in fig. 5.4 on the facing page.

Add more to the AST format tree? fig. 5.4 on the next page




Figure 5.4: The format of the abstract syntax tree in Eclipse.

5.3 The ASTVisitor

So far, the only thing that has been addressed is how the data that is going to be the basis for our analysis is structured. Another aspect of it is how we are going to traverse the AST to gather the information we need, so we can conclude about the properties we are analysing. It is of course possible to start at the top of the tree, and manually search through its nodes for the ones we are looking for, but that is a bit inconvenient. To be able to efficiently utilize such an approach, we would need to make our own framework for traversing the tree and visiting only the types of nodes we are after. Luckily, this functionality is already provided in Eclipse, by its **ASTVisitor**¹.

The Eclipse AST, together with its **ASTVisitor**, follows the *Visitor* pattern [Gam+95]. The intent of this design pattern is to facilitate extending

¹org.eclipse.jdt.core.dom.ASTVisitor

the functionality of classes without touching the classes themselves.

Let us say that there is a class hierarchy of elements. These elements all have a method `accept(Visitor visitor)`. In its simplest form, the `accept` method just calls the `visit` method of the visitor with itself as an argument, like this: `visitor.visit(this)`. For the visitors to be able to extend the functionality of all the classes in the elements hierarchy, each `Visitor` must have one visit method for each concrete class in the hierarchy. Say the hierarchy consists of the concrete classes `ConcreteElementA` and `ConcreteElementB`. Then each visitor must have the (possibly empty) methods `visit(ConcreteElementA element)` and `visit(ConcreteElementB element)`. This scenario is depicted in fig. 5.5 on this page.

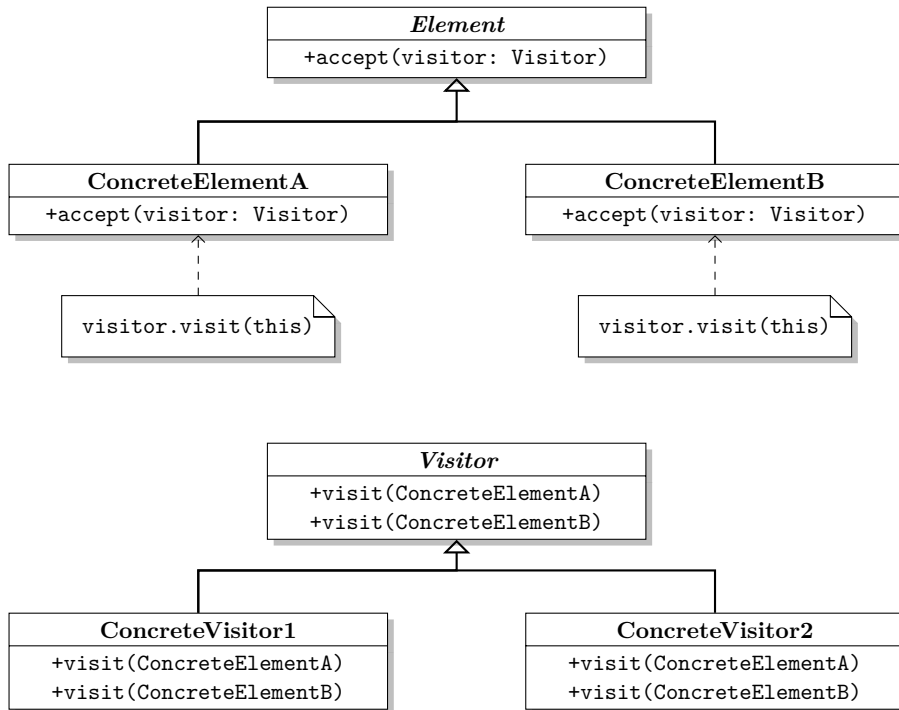


Figure 5.5: The Visitor Pattern.

The use of the visitor pattern can be appropriate when the hierarchy of elements is mostly stable, but the family of operations over its elements is constantly growing. This is clearly the case for the Eclipse AST, since the hierarchy of type `ASTNode` is very stable, but the functionality of its elements is extended every time someone needs to operate on the AST. Another aspect of the Eclipse implementation is that it is a public API, and the visitor pattern is an easy way to provide access to the nodes in the tree.

The version of the visitor pattern implemented for the AST nodes in Eclipse also provides an elegant way to traverse the tree. It does so by following the convention that every node in the tree first let the visitor visit itself, before it also makes all its children accept the visitor. The children are only visited if the visit method of their parent returns `true`. This pattern

then makes for a prefix traversal of the AST. If postfix traversal is desired, the visitors also has `endVisit` methods for each node type, that is called after the `visit` method for a node. In addition to these visit methods, there are also the methods `preVisit(ASTNode)`, `postVisit(ASTNode)` and `preVisit2(ASTNode)`. The `preVisit` method is called before the type-specific `visit` method. The `postVisit` method is called after the type-specific `endVisit`. The type specific `visit` is only called if `preVisit2` returns `true`. Overriding the `preVisit2` is also altering the behavior of `preVisit`, since the default implementation is responsible for calling it.

An example of a trivial `ASTVisitor` is shown in listing 12 on the current page.

```
public class CollectNamesVisitor extends ASTVisitor {
    Collection<Name> names = new LinkedList<Name>();

    @Override
    public boolean visit(QualifiedName node) {
        names.add(node);
        return false;
    }

    @Override
    public boolean visit(SimpleName node) {
        names.add(node);
        return true;
    }
}
```

Listing 12: An `ASTVisitor` that visits all the names in a subtree and adds them to a collection, except those names that are children of any `QualifiedName`.

5.4 Property collectors

The prefixes and unfixes are found by property collectors¹. A property collector is of the `ASTVisitor` type, and thus visits nodes of type `ASTNode` of the abstract syntax tree (see section 5.3 on page 69).

5.4.1 The PrefixesCollector

The `PrefixesCollector`² finds prefixes that makes up the basis for calculating move targets for the *Extract and Move Method* refactoring. It visits expression statements³ and creates prefixes from its expressions in the

¹`no.uio.ifi.refaktor.extractors.collectors.PropertyCollector`

²`no.uio.ifi.refaktor.extractors.collectors.PrefixesCollector`

³`org.eclipse.jdt.core.dom.ExpressionStatement`

case of method invocations. The prefixes found is registered with a prefix set, together with all its sub-prefixes.

5.4.2 The UnfixesCollector

The **UnfixesCollector**¹ finds unfixes within a selection.

Give more technical detail?

5.4.3 The ContainsReturnStatementCollector

Remove section?

The **ContainsReturnStatementCollector**² is a very simple property collector. It only visits the return statements within a selection, and can report whether it encountered a return statement or not.

5.4.4 The LastStatementCollector

The **LastStatementCollector**³ collects the last statement of a selection. It does so by only visiting the top level statements of the selection, and compares the textual end offset of each encountered statement with the end offset of the previous statement found.

5.5 Checkers

The checkers are a range of classes that checks that text selections complies with certain criteria. All checkers operates under the assumption that the code they check is free from compilation errors. If a **Checker**⁴ fails, it throws a **CheckerException**. The checkers are managed by the **LegalStatementsChecker**, which does not, in fact, implement the **Checker** interface. It does, however, run all the checkers registered with it, and reports that all statements are considered legal if no **CheckerException** is thrown. Many of the checkers either extends the **PropertyCollector** or utilizes one or more property collectors to verify some criteria. The checkers registered with the **LegalStatementsChecker** are described next. They are run in the order presented below.

5.5.1 The CallToProtectedOrPackagePrivateMethodChecker

This checker is used to check that at selection does not contain a call to a method that is protected or package-private. Such a method either has the access modifier **protected** or it has no access modifier.

The workings of the **CallToProtectedOrPackagePrivateMethodChecker** is very simple. It looks for calls to methods that are ei-

¹`no.uio.ifi.refaktor.extractors.collectors.UnfixesCollector`

²`no.uio.ifi.refaktor.analyze.collectors.ContainsReturnStatementCollector`

³`no.uio.ifi.refaktor.analyze.collectors.LastStatementCollector`

⁴`no.uio.ifi.refaktor.analyze.analyzers.Checker`

ther protected or package-private within the selection, and throws an `IllegalExpressionFoundException` if one is found.

5.5.2 The `DoubleClassInstanceCreationChecker`

The `DoubleClassInstanceCreationChecker` checks that there are no double class instance creations where the inner constructor call take and argument that is built up using field references.

The checker visits all nodes of type `ClassInstanceCreation` within a selection. For all of these nodes, if its parent also is a class instance creation, it accepts a visitor that throws a `IllegalExpressionFoundException` if it encounters a name that is a field reference.

5.5.3 The `InstantiationOfNonStaticInnerClassChecker`

The `InstantiationOfNonStaticInnerClassChecker` checks that selections does not contain instantiations of non-static inner classes. The `MoveInstanceMethodProcessor` in Eclipse does not handle such instantiations gracefully when moving a method. This problem is also related to bug...

[File Eclipse bug report](#)

5.5.4 The `EnclosingInstanceReferenceChecker`

The purpose of this checker is to verify that the names in a text selection are not referencing any enclosing instances. In theory, the underlying problem could be solved in some situations, but our dependency on the `MoveInstanceMethodProcessor` prevents this.

The `EnclosingInstanceReferenceChecker`¹ is a modified version of the `EnclosingInstanceReferenceFinder`² from the `MoveInstanceMethodProcessor`. Wherever the `EnclosingInstanceReferenceFinder` would create a fatal error status, the checker will throw a `CheckerException`.

The checker works by first finding all of the enclosing types of a selection. Thereafter, it visits all the simple names of the selection to check that they are not references to variables or methods declared in any of the enclosing types. In addition, the checker visits `this`-expressions to verify that no such expressions are qualified with any name.

5.5.5 The `ReturnStatementsChecker`

The checker for return statements is meant to verify that a text selection is consistent regarding return statements.

If the selection is free from return statements, then the checker validates. So this is the first thing the checker investigates.

¹`no.uio.ifi.refaktor.analyze.analyzers.EnclosingInstanceReferenceChecker`

²`org.eclipse.jdt.internal.corext.refactoring.structure.MoveInstanceMethodProcessor.EnclosingInstanceReferenceFinder`

If the checker proceeds any further, it is because the selection contains one or more return statements. The next test is therefore to check if the last statement of the selection ends in either a return or a throw statement. The responsibility for checking that the last statement of the selection eventually ends in a return or throw statement, is put on the **LastStatementOfSelectionEndsInReturnOrThrowChecker**. For every node visited, if the node is a statement, it does a test to see if the statement is a return, a throw or if it is an implicit return statement. If this is the case, no further checking is done. This checking is done in the **preVisit2** method (see section 5.3 on page 69). If the node is not of a type that is being handled by its type-specific visit method, the checker performs a simple test. If the node being visited is not the last statement of its parent that is also enclosed by the selection, an **IllegalStatementFoundException** is thrown. This ensures that all statements are taken care of, one way or the other. It also ensures that the checker is conservative in the way it checks for legality of the selection.

To examine if a statement is an implicit return statement, the checker first finds the last statement declared in its enclosing method. If this statement is the same as the one under investigation, it is considered an implicit return statement. If the statements are not the same, the checker does a search to see if the statement examined is also the last statement of the method that can be reached. This includes the last statement of a block statement, a labeled statement, a synchronized statement or a try statement, that in turn is the last statement enclosed by one of the statement types listed. This search goes through all the parents of a statement until a statement is found that is not one of the mentioned acceptable parent statements. If the search ends in a method declaration, then the statement is considered to be the last reachable statement of the method, and thus it is an implicit return statement.

There are two kinds of statements that are handled explicitly: If-statements and try-statements. Block, labeled and do-statements are handled by fall-through to the other two.

If-statements are handled explicitly by overriding their type-specific visit method. If the then-part does not contain any return or throw statements an **IllegalStatementFoundException** is thrown. If it does contain a return or throw, its else-part is checked. If the else-part is non-existent, or it does not contain any return or throw statements an exception is thrown. If no exception is thrown while visiting the if-statement, its children are visited.

A try-statement is checked very similar to an if-statement. Its body must contain a return or throw. The same applies to its catch clauses and finally body. Failure to validate produces an **IllegalStatementFoundException**.

If the checker does not complain at any point, the selection is considered valid with respect to return statements.

5.5.6 The AmbiguousReturnValueChecker

This checker verifies that there are no ambiguous return values in a selection.

First, the checker needs to collect some data. Those data are the binding

keys for all simple names that are assigned to within the selection, including variable declarations, but excluding fields. The checker also collects whether there exists a return statement in the selection or not. No further checks of return statements are needed, since, at this point, the selection is already checked for illegal return statements (see section 5.5.5 on page 73).

After the binding keys of the assignees are collected, the checker searches the part of the enclosing method that is after the selection for references whose binding keys are among the collected keys. If more than one unique referral is found, or only one referral is found, but the selection also contains a return statement, we have a situation with an ambiguous return value, and an exception is thrown.

5.5.7 The IllegalStatementsChecker

This checker is designed to check for illegal statements.

Notice that labels in break and continue statements needs some special treatment. Since a label does not have any binding information, we have to search upwards in the AST to find the **LabeledStatement** that corresponds to the label from the break or continue statement, and check that it is contained in the selection. If the break or continue statement does not have a label attached to it, it is checked that its innermost enclosing loop or switch statement (break statements only) also is contained in the selection.

Follow the development in the semantics section. . .

Chapter 6

Technicalities

6.1 Source code organization

All the parts of this master's project is under version control with Git¹.

The software written is organized as some Eclipse plugins. Writing a plugin is the natural way to utilize the API of Eclipse. This also makes it possible to provide a user interface to manually run operations on selections in program source code or whole projects/packages.

When writing a plugin in Eclipse, one has access to resources such as the current workspace, the open editor and the current selection.

The thesis work is contained in the following Eclipse projects:

no.uio.ifi.refaktor

This is the main Eclipse plugin project, and contains all of the business logic for the plugin.

no.uio.ifi.refaktor.tests

This project contains the tests for the main plugin.

no.uio.ifi.refaktor.examples

Contains example code used in testing. It also contains code for managing this example code, such as creating an Eclipse project from it before a test run.

no.uio.ifi.refaktor.benchmark

This project contains code for running search based versions of the composite refactoring over selected Eclipse projects.

no.uio.ifi.refaktor.releng

Contains the rmap, queries and target definitions needed by by Buckminster on the Jenkins continuous integration server.

¹<http://git-scm.com/>

6.1.1 The `no.uio.ifi.refaktor` project

`no.uio.ifi.refaktor.analyze`

This package, and its sub-packages, contains code that is used for analyzing Java source code. The most important sub-packages are presented below.

`no.uio.ifi.refaktor.analyze.analyzers`

This package contains source code analyzers. These are usually responsible for analyzing text selections or running specialized analyzers for different kinds of entities. Their structure are often hierarchical. This means that you have an analyzer for text selections, that in turn is utilized by an analyzer that analyzes all the selections of a method. Then there are analyzers for analyzing all the methods of a type, all the types of a compilation unit, all the compilation units of a package, and, at last, all of the packages in a project.

`no.uio.ifi.refaktor.analyze.checkers`

A package containing checkers. The checkers are classes used to validate that a selection can be further analyzed and chosen as a candidate for a refactoring. Invalidating properties can be such as usage of inner classes or the need for multiple return values.

`no.uio.ifi.refaktor.analyze.collectors`

This package contains the property collectors. Collectors are used to gather properties from a text selection. This is mostly properties regarding referenced names and their occurrences. It is these properties that makes up the basis for finding the best candidates for a refactoring.

`no.uio.ifi.refaktor.change`

This package, and its sub-packages, contains functionality for manipulate source code.

`no.uio.ifi.refaktor.change.changers`

This package contains source code changers. They are used to glue together the analysis of source code and the actual execution of the changes.

`no.uio.ifi.refaktor.change.executors`

The executors that are responsible for making concrete changes are found in this package. They are mostly used to create and execute one or more Eclipse refactorings.

`no.uio.ifi.refaktor.change.processors`

Contains a refactoring processor for the *Move Method* refactoring. The code is stolen and modified to fix a bug. The related bug is described in appendix A.2 on page 111.

no.uio.ifi.refaktor.handlers

This package contains handlers for the commands defined in the plugin manifest.

no.uio.ifi.refaktor.prefix

This package contains the **Prefix** type that is the data representation of the prefixes found by the **PrefixesCollector**. It also contains the prefix set for storing and working with prefixes.

no.uio.ifi.refaktor.statistics

The package contains statistics functionality. Its heart is the statistics aspect that is responsible for gathering statistics during the execution of the *Extract and Move Method* refactoring.

no.uio.ifi.refaktor.statistics.reports

This package contains a simple framework for generating reports from the statistics data generated by the aspect. Currently, the only available report type is a simple text report.

no.uio.ifi.refaktor.textselection

This package contains the two custom text selections that are used extensively throughout the project. One of them is just a subclass of the other, to support the use of the memento pattern to optimize the memory usage during benchmarking.

no.uio.ifi.refaktor.debugging

The package contains a debug utility class. In addition to this, the package **no.uio.ifi.refaktor.utils.aspects** contains a couple of aspects used for debugging purposes.

no.uio.ifi.refaktor.utils

Utility package that contains all the functionality that has to do with parsing of source code. It also has utility classes for looking up handles to methods and types et cetera.

no.uio.ifi.refaktor.utils.caching

This package contains the caching manager for compilation units, along with classes for different caching strategies.

no.uio.ifi.refaktor.utils.nullobjects

Contains classes for creating different null objects. Most of the classes is used to represent null objects of different handle types. These null objects are returned from various utility classes instead of returning a **null** value when other values are not available.

6.2 Continuous integration

The continuous integration server Jenkins¹ has been set up for the project². It is used as a way to run tests and perform code coverage analysis.

To be able to build the Eclipse plugins and run tests for them with Jenkins, the component assembly project Buckminster³ is used, through its plugin for Jenkins. Buckminster provides for a way to specify the resources needed for building a project and where and how to find them. Buckminster also handles the setup of a target environment to run the tests in. All this is needed because the code to build depends on an Eclipse installation with various plugins.

6.2.1 Problems with AspectJ

The Buckminster build worked fine until introducing AspectJ into the project. When building projects using AspectJ, there are some additional steps that needs to be performed. First of all, the aspects themselves must be compiled. Then the aspects needs to be woven with the classes they affect. This demands a process that does multiple passes over the source code.

When using AspectJ with Eclipse, the specialized compilation and the weaving can be handled by the AspectJ Development Tools⁴. This works all fine, but it complicates things when trying to build a project depending on Eclipse plugins outside of Eclipse. There is supposed to be a way to specify a compiler adapter for javac, together with the file extensions for the file types it shall operate. The AspectJ compiler adapter is called **Ajc11CompilerAdapter**⁵, and it works with files that has the extensions ***.java** and ***.aj**. I tried to setup this in the build properties file for the project containing the aspects, but to no avail. The project containing the aspects does not seem to be built at all, and the projects that depends on it complains that they cannot find certain classes.

I then managed to write an Ant⁶ build file that utilizes the AspectJ compiler adapter, for the **no.uio.ifi.refaktor** plugin. The problem was then that it could no longer take advantage of the environment set up by Buckminster. The solution to this particular problem was of a “hacky” nature. It involves exporting the plugin dependencies for the project to an Ant build file, and copy the exported path into the existing build script. But then the Ant script needs to know where the local Eclipse installation is located. This is no problem when building on a local machine, but to utilize the setup done by Buckminster is a problem still unsolved. To get the classpath for the build setup correctly, and here comes the most “hacky” part of the solution, the Ant script has a target for copying the

¹<http://jenkins-ci.org/>

²A work mostly done by the supervisor.

³<http://www.eclipse.org/buckminster/>

⁴<https://www.eclipse.org/ajdt/>

⁵`org.aspectj.tools.ant.taskdefs.Ajc11CompilerAdapter`

⁶<https://ant.apache.org/>

classpath elements into a directory relative to the project directory and checking it into Git. When no **ECLIPSE_HOME** property is set while running Ant, the script uses the copied plugins instead of the ones provided by the Eclipse installation when building the project. This obviously creates some problems with maintaining the list of dependencies in the Ant file, as well as remembering to copy the plugins every time the list of dependencies change.

The Ant script described above is run by Jenkins before the Buckminster setup and build. When setup like this, the Buckminster build succeeds for the projects not using AspectJ, and the tests are run as normal. This is all good, but it feels a little scary, since the reason for Buckminster not working with AspectJ is still unknown.

The problems with building with AspectJ on the Jenkins server lasted for a while, before they were solved. This is reflected in the “Test Result Trend” and “Code Coverage Trend” reported by Jenkins.

Chapter 7

Benchmarking

This part of the master's project is located in the Eclipse project `no.uio.ifi.refaktor.benchmark`. The purpose of it is to run the equivalent of the `SearchBasedExtractAndMoveMethodChanger` (see section 4.2.3 on page 61) over a larger software project, both to test its robustness but also its effect on different software metrics.

7.1 The benchmark setup

The benchmark itself is set up as a JUnit test case. This is a convenient setup, and utilizes the JUnit Plugin Test Launcher. This provides us with a fully functional Eclipse workbench. Most importantly, this gives us access to the Java Model of Eclipse (see section 5.1 on page 65).

7.1.1 The `ProjectImporter`

The Java project that is going to be used as the data for the benchmark, must be imported into the JUnit workspace. This is done by the `ProjectImporter`¹. The importer require the absolute path to the project description file. It is named `.project` and is located at the root of the project directory.

The project description is loaded to find the name of the project to be imported. The project that shall be the destination for the import is created in the workspace, on the base of the name from the description. Then an import operation is created, based on both the source and destination information. The import operation is run to perform the import.

I have found no simple API call to accomplish what the importer does, which tells me that it may not be too many people performing this particular action. The solution to the problem was found on Stack Overflow². It contains enough dirty details to be considered inconvenient to use, if not wrapping it in a class like my `ProjectImporter`. One would probably have to delve into the source code for the import wizard to find out how the import operation works, if no one had already done it.

¹`no.uio.ifi.refaktor.benchmark.ProjectImporter`

²<https://stackoverflow.com/questions/12401297>

7.2 Statistics

Statistics for the analysis and changes is captured by the **StatisticsAspect**¹. This an *aspect* written in AspectJ.

7.2.1 AspectJ

AspectJ² is an extension to the Java language, and facilitates combining aspect-oriented programming with the object-oriented programming in Java.

Aspect-oriented programming is a programming paradigm that is meant to isolate so-called *cross-cutting concerns* into their own modules. These cross-cutting concerns are functionalities that spans over multiple classes, but may not belong naturally in any of them. It can be functionality that does not concern the business logic of an application, and thus may be a burden when entangled with parts of the source code it does not really belong. Examples include logging, debugging, optimization and security.

Aspects are interacting with other modules by defining advices. The concept of an *advice* is known from both aspect-oriented and functional programming [14a]. It is a function that modifies another function when the latter is run. An advice in AspectJ is somewhat similar to a method in Java. It is meant to alter the behavior of other methods, and contains a body that is executed when it is applied.

An advice can be applied at a defined *pointcut*. A pointcut picks out one or more *join points*. A join point is a well-defined point in the execution of a program. It can occur when calling a method defined for a particular class, when calling all methods with the same name, accessing/assigning to a particular field of a given class and so on. An advice can be declared to run both before, after returning from a pointcut, when there is thrown an exception in the pointcut or after the pointcut either returns or throws an exception. In addition to picking out join points, a pointcut can also bind variables from its context, so they can be accessed in the body of an advice. An example of a pointcut and an advice is found in listing 13 on the facing page.

7.2.2 The Statistics class

The statistics aspect stores statistical information in an object of type **Statistics**. As of now, the aspect needs to be initialized at the point in time where it is desired that it starts its data gathering. At any point in time the statistics aspect can be queried for a snapshot of the current statistics.

The **Statistics** class also include functionality for generating a report of its gathered statistics. The report can be given either as a string or it can be written to a file.

¹`no.uio.ifi.refaktor.aspects.StatisticsAspect`

²<http://eclipse.org/aspectj/>

```

pointcut methodAnalyze(
    SearchBasedExtractAndMoveMethodAnalyzer analyzer) :
    call(* SearchBasedExtractAndMoveMethodAnalyzer.analyze())
    && target(analyzer);

after(SearchBasedExtractAndMoveMethodAnalyzer analyzer) :
    methodAnalyze(analyzer) {
    statistics.methodCount++;
    debugPrintMethodAnalysisProgress(analyzer.method);
}

```

Listing 13: An example of a pointcut named `methodAnalyze`, and an advice defined to be applied after it has occurred.

7.2.3 Advices

The statistics aspect contains advices for gathering statistical data from different parts of the benchmarking process. It captures statistics from both the analysis part and the execution part of the composite *Extract and Move Method* refactoring.

For the analysis part, there are advices to count the number of text selections analyzed and the number of methods, types, compilation units and packages analyzed. There are also advices that counts for how many of the methods there is found a selection that is a candidate for the refactoring, and for how many methods there is not.

There exists advices for counting both the successful and unsuccessful executions of all the refactorings. Both for the *Extract Method* and *Move Method* refactorings in isolation, as well as for the combination of them.

7.3 Optimizations

When looking for optimizations to make for the benchmarking process, I used the VisualVM¹ profiler for the Java Virtual Machine to both profile the application and also to make memory dumps of its heap.

7.3.1 Caching

When profiling the benchmark process before making any optimizations, it early became apparent that the parsing of source code was a place to direct attention towards. This discovery was done when only *analyzing* source code, before trying to do any *manipulation* of it. Caching of the parsed ASTs seemed like the best way to save some time, as expected. With only a simple cache of the most recently used AST, the analysis time was speeded up by a factor of around 20. This number depends a little upon which type of system the analysis is run.

¹<http://visualvm.java.net/>

The caching is managed by a cache manager, that now, by default, utilizes the not so well known feature of Java called a *soft reference*. Soft references are best explained in the context of weak references. A *weak reference* is a reference to an object instance that is only guaranteed to persist as long as there is a *strong reference* or a soft reference referring the same object. If no such reference is found, its referred object is garbage collected. A strong reference is basically the same as a regular Java reference. A soft reference has the same guarantees as a weak reference when it comes to its relation to strong references, but it is not necessarily garbage collected whenever there exists no strong references to it. A soft reference *may* reside in memory as long as the JVM has enough free memory in the heap. A soft reference will therefore usually perform better than a weak reference when used for simple caching and similar tasks. The way to use a soft/weak reference is to ask for its referent. The return value then has to be tested to check that it is not `null`. For the basic usage of soft references, see listing 14 on the current page. For a more thorough explanation of weak references in general, see [Nic06].

```
// Strong reference
Object strongRef = new Object();

// Soft reference
SoftReference<Object> softRef =
    new SoftReference<Object>(new Object());

// Using the soft reference
Object obj = softRef.get();
if (obj != null) {
    // Use object here
}
```

Listing 14: Showing the basic usage of soft references. Weak references is used the same way. (The references are part of the `java.lang.ref` package.)

The cache based on soft references has no limit for how many ASTs it caches. It is generally not advisable to keep references to ASTs for prolonged periods of time, since they are expensive structures to hold on to. For regular plugin development, Eclipse recommends not creating more than one AST at a time to limit memory consumption. Since the benchmarking has nothing to do with user experience, and throughput is everything, these advices are intentionally ignored. This means that during the benchmarking process, the target Eclipse application may very well work close to its memory limit for the heap space for long periods during the benchmark.

7.3.2 Candidates stored as mementos

When performing large scale analysis of source code for finding candidates to the *Extract and Move Method* refactoring, memory is an issue. One

of the inputs to the refactoring is a variable binding. This variable binding indirectly retains a whole AST. Since ASTs are large structures, this quickly leads to an **OutOfMemoryError** if trying to analyze a large project without optimizing how we store the candidates data. This means that the JVM cannot allocate more memory for our benchmark, and it exists disgracefully.

A possible solution could be to just allow the JVM to allocate even more memory, but this is not a dependable solution. The allocated memory could easily supersede the physical memory of a machine, and that would make the benchmark go really slow.

Thus, the candidates data must be stored in another format. Therefore, we use the memento pattern to store the variable binding information. This is done in a way that makes it possible to retrieve the variable binding at a later point. The data that is stored to achieve this, is the key to the original variable binding. In addition to the key, we know which method and text selection the variable is referenced in, so that we can find it by parsing the source code and search for it when it is needed.

7.4 Handling failures

write

Chapter 8

Case Studies

In this chapter I am going to present a few case studies. This is done to give an impression of how the search-based *Extract and Move Method* refactoring performs when giving it a larger project to take on. I will try to answer where it lacks, in terms of completeness, as well as showing its effect on refactored source code.

The first and primary case, is refactoring source code from the Eclipse JDT UI project. The project is chosen because it is a real project, still in development, with a large code base that is written by many different people through several years. The code is installed in thousands of Eclipse applications worldwide, and must be seen as a good representative for professionally written Java source code. It is also the home for most of the JDT refactoring code.

For the second case, the *Extract and Move Method* refactoring is fed the `no.uio.ifi.refaktor` project. This is done as a variation of the “dogfooding” methodology, where you use your own tools to do your job, also referred to as “eating your own dog food” [Har06].

8.1 The tools

For conducting these experiments, three tools are used. Two of the “tools” both uses Eclipse as their platform. The first is our own tool, written to be able to run the *Extract and Move Method* refactoring as a batch process, analyzing and refactoring many methods after each other. The second is JUnit, that is used for running the projects own unit tests on the target code both before and after it is refactored. The last tool that is used is a code quality management tool, called SonarQube. It can be used to perform different tasks for assuring code quality, but we are only going to take advantage of one of its main features, namely Quality profiles.

A quality profile is used to define a set of coding rules that a project is supposed to comply with. Failure to following these rules will be recorded as so-called “issues”, marked as having one of several degrees of severities, ranging from “info” to “blocker”, where the latter one is the most severe. The measurements done for these case studies are therefore not presented as fine-grained software metrics results, but rather as the number of issues

for each defined rule.

In addition to the coding rules defined through quality profiles, SonarQube calculates the complexity of source code. The metric that is used is cyclomatic complexity, developed by Thomas J. McCabe in 1976 [McC76]. In this metric, functions have an initial complexity of 1, and whenever the control flow of a function splits, the complexity increases by one¹. SonarQube discriminates between functions and accessors. Accessors are methods that are recognized as setters or getters. Accessors are not counted in the complexity analysis.

8.2 The SonarQube quality profile

The quality profile that is used with SonarQube in these case studies has got the name IFI Refaktor Case Study (version 6). The rules defined in the profile are chosen because they are the available rules found in SonarQube that measures complexity and coupling. Now follows a description of the rules in the quality profile. The values that are set for these rules are listed in table 8.1 on the facing page.

Avoid too complex class is a rule that measures cyclomatic complexity for every statement in the body of a class, except for setters and getter. The threshold value set is its default value of 200.

Classes should not be coupled to too many other classes is a rule that measures how many other classes a class depends upon. It does not count the dependencies of nested classes. It is meant to promote the Single Responsibility Principle. The metric for the rule resembles the CBO metric that is described in section 1.3.5 on page 31, but is only considering outgoing dependencies. The max value for the rule is chosen on the background of an empirical study by Raed Shatnawi, that concludes that the number 9 is the most useful threshold for the CBO metric [Sha10]. This study is also performed on Eclipse source code, so this threshold value should be particularly well suited for the Eclipse JDT UI case in this chapter.

Control flow statements . . . should not be nested too deeply is a rule that is meant to counter “Spaghetti code”. It measures the nesting level of if, for, while, switch and try statements. The nesting levels start at 1. The max value set is its default value of 3.

Methods should not be too complex is a rule that measures cyclomatic complexity the same way as the “Avoid too complex class” rule. The max value used is 10, which “seems like a reasonable, but not magical, upper limit“ [McC76].

Methods should not have too many lines is a rule that simply measures the number of lines in methods. The threshold value of 20 is

¹<http://docs.codehaus.org/display/SONAR/Metric+definitions>

used for this metric. This is based on my own subjective opinions, as the default value of 100 seems a bit too loose.

NPath Complexity is a rule that measures the number of possible execution paths through a function. The value used is the default value of 200, that seems like a recognized threshold for this metric.

Too many methods is a rule that measures the number of methods in a class. The threshold value used is the default value of 10.

Table 8.1: The IFI Refaktor Case Study quality profile (version 6).

Rule	Max value
Avoid too complex class	200
Classes should not be coupled to too many other classes (Single Responsibility Principle)	9
Control flow statements ... should not be nested too deeply	3
Methods should not be too complex	10
Methods should not have too many lines	20
NPath Complexity	200
Too many methods	10

8.3 The input

A precondition for the source code that is going to be the target for a series of *Extract and Move Method* refactorings, is that it is organized as an Eclipse project. It is also assumed that the code is free from compilation errors.

8.4 The experiment

For a given project, the first job that is done, is to refactor its source code. The refactoring batch job produces three things: The refactored project, statistics gathered during the execution of the series of refactorings, and an error log describing any errors happening during this execution. See chapter 7 on page 83 for more information about how the refactorings are performed.

After the refactoring process is done, the before- and after-code is analyzed with SonarQube. The analysis results are then stored in a database and displayed through a SonarQube server with a web interface.

How long are these results going to be publicly available?

The before- and after-code is also tested with their own unit tests. This is done to discover any changes in the semantic behavior of the refactored code, within the limits of these tests.

8.5 Case 1: The Eclipse JDT UI project

This case is the ultimate test for our *Extract and Move Method* refactoring. The target source code is massive. With its over 300,000 lines of code and over 25,000 methods, it is formidable task to perform automated changes on it. There should be plenty of situations where things can go wrong, and, as we shall see later, they do.

I will start by presenting some statistics from the refactoring execution, before I pick apart the SonarQube analysis and conclude by commenting on the results from the unit tests. The configuration for the experiment is specified in table 8.2 on the current page.

Table 8.2: Configuration for Case 1.

Benchmark data	
Launch configuration	CaseStudy.launch
Project	no.uio.ifi.refaktor.benchmark
Repository	gitolite@git.uio.no:ifi-stolz-refaktor
Commit	43c16c04520746edd75f8dc2a1935781d3d9de6c
Input data	
Project	org.eclipse.jdt.ui
Repository	git://git.eclipse.org/gitroot/jdt/eclipse.jdt.ui.git
Commit	f218388fea6d4ec1da7ce22432726c244888bb6b
Branch	R3_8_maintenance
Tests suites	org.eclipse.jdt.ui.tests.AutomatedSuite, org.eclipse.jdt.ui.tests.refactoring.all.- AllAllRefactoringTests

8.5.1 Statistics

The statistics gathered during the refactoring execution is presented in table 8.3 on page 101.

Execution time

I consider the total execution time of approximately 1.5 hours as being acceptable. It clearly makes the batch process unsuitable for doing any on-demand analysis or changes, but it is good enough for running periodic jobs, like over-night analysis.

As the statistics show, 75% of the total time goes into making the actual code changes. The time consumers are here the primitive *Extract Method* and *Move Method* refactorings. Included in the change time is the parsing and precondition checking done by the refactorings, as well as

textual changes done to files on disk. All this parsing and disk access is time-consuming, and constitute a large part of the change time.

In comparison, the pure analysis time, used to find suitable candidates, only make up for 15% of the total time consumed. This includes analyzing almost 600,000 text selections, while the number of attempted executions of the *Extract and Move Method* refactoring are only about 2,500. So the number of executed primitive refactorings are approximately 5,000. Assuming the time used on miscellaneous tasks are used mostly for parsing source code for the analysis, we can say that the time used for analyzing code is at most 25% of the total time. This means that for every primitive refactoring executed, we can analyze around 360 text selections. So, with an average of about 21 text selections per method, it is reasonable to say that we can analyze over 15 methods in the time it takes to perform a primitive refactoring.

Refactoring candidates

Out of the 27,667 methods that were analyzed, 2,552 methods contained selections that were considered candidates for the *Extract and Move Method* refactoring. This is roughly 9% off the methods in the project. These 9% of the methods had on average 14.4 text selections that were considered possible refactoring candidates.

Executed refactorings

2,469 out of 2,552 attempts on executing the *Extract and Move Method* refactoring were successful, giving a success rate of 96.7%. The failure rate of 3.3% stem from situations where the analysis finds a candidate selection, but the change execution fails. This failure could be an exception that was thrown, and the refactoring aborts. It could also be the precondition checking for one of the primitive refactorings that gives us an error status, meaning that if the refactoring proceeds, the code will contain compilation errors afterwards, forcing the composite refactoring to abort. This means that if the *Extract Method* refactoring fails, no attempt is done for the *Move Method* refactoring.

Out of the 2,552 *Extract Method* refactorings that were attempted executed, 69 of them failed. This give a failure rate of 2.7% for the primitive refactoring. In comparison, the *Move Method* refactoring had a failure rate of 0.6 % of the 2,483 attempts on the refactoring.

Redundant information? Put in benchmark chapter?

8.5.2 SonarQube analysis

Results from the SonarQube analysis are shown in table 8.4 on page 102.

Diversity in the number of entities analyzed

The analysis performed by SonarQube is reporting fewer methods than found by the pre-refactoring analysis. SonarQube discriminates between functions (methods) and accessors, so the 1,296 accessors play a part in this

calculation. SonarQube also has the same definition as our plugin when it comes to how a class is defined. Therefore it seems like SonarQube misses 277 classes that our plugin handles. This can explain why the SonarQube report differs from our numbers by approximately 2,500 methods,

Complexity

On all complexity rules that work on the method level, the number of issues decreases with between 3.1% and 6.5% from before to after the refactoring. The average complexity of a method decreases from 3.6 to 3.3, which is an improvement of about 8.3%. So, on the method level, the refactoring must be said to have a slightly positive impact.

The improvement in complexity on the method level is somewhat traded for complexity on the class level. The complexity per class metric is worsened by 3% from before to after. The issues for the “Too many methods” rule also increase by 14.5%. These numbers indicate that the refactoring makes quite a lot of the classes a little more complex overall. This is the expected outcome, since the *Extract and Move Method* refactoring introduces almost 2,500 new methods into the project.

The only number that can save the refactoring’s impact on complexity on the class level, is the “Avoid too complex class” rule. It improves with 2.5%, thus indicating that the complexity is moderately better distributed between the classes after the refactoring than before.

Coupling

One of the hopes when starting this project, was to be able to make a refactoring that could lower the coupling between classes. Better complexity at the method level is a not very unexpected byproduct of dividing methods into smaller parts. Lowering the coupling on the other hand, is a far greater task. This is also reflected in the results for the only coupling rule defined in the SonarQube quality profile, namely the “Classes should not be coupled to too many other classes (Single Responsibility Principle)” rule.

The number of issues for the coupling rule is 1,098 before the refactoring, and 1,199 afterwards. This is an increase in issues of 9.2%, and a blow for this project. These numbers can be interpreted two ways. The first possibility is that our assumptions are wrong, and that increasing indirection does not decrease coupling between classes. The other possibility is that our analysis and choices of candidate text selections are not good enough. I vote for the second possibility. (Voting against the public opinion may also be a little bold.)

What probably happens is, that many of the times the *Extract and Move Method* refactoring is performed, the *Move Method* refactoring “drags” with it references to classes that are unknown to the method destination. If it happens to be so lucky that it removes a dependency from one class, it might as well introduce three new dependencies to another class. In those situations that a class does not know about the originating class of a moved method, the *Move Method* refactoring most certainly will introduce

a dependency. This is because there is a bug¹ in the refactoring, making it pass an instance of the originating class as a reference to the moved method, regardless of whether the reference is used in the method body or not.

There is also the possibility that the heuristics used to find candidate text selections are not good enough, they most certainly are not. I wish I had more time to fine-tune them, and to complete the analysis part of the project, but this is simply not the case. This becomes even clearer when analyzing the unit test results for the after-code.

Totals

On the bright side, the total number of issues is lower after the refactoring than it was before. Before the refactoring, the total number of issues is 8,270, and after it is 8,155. An improvement of only 1.4%.

Then SonarQube tells me that the total complexity has increased by 2.9%, and that the (more questionable) “technical debt” has increased from 1,003.4 to 1,032.7 days, also a deterioration of 2.9%. Although these numbers are similar, no correlation has been found between them.

8.5.3 Unit tests

The tests that have been run for the Eclipse JDT UI project, are the tests in the test suites specified as the main test suites on the JDT UI wiki page on how to contribute to the project².

Before the refactoring

Running the tests for the before-code of Eclipse JDT UI yielded 4 errors and 3 failures for the **AutomatedSuite** test suite (2,007 test cases), and 2 errors and 3 failures for the **AllAllRefactoringTests** test suite (3,816 test cases).

After the refactoring

The test results for the after-code of the Eclipse JDT UI project is another story. The reason for this is that during the setup for the unit tests, Eclipse now reports that the project contains 322 fatal errors, and a lot of errors that probably follows from these. This is another blow for this master’s project.

It has now been shown that the *Extract and Move Method* refactoring, in its current state, produces code that does not compile. Had these errors originated from only one bug, it would not have been much of a problem, but this is not the case. By only looking at some random compilation problems in the refactored code, I came up with at least four different bugs that caused those problems. I then stopped looking for more, since some of the bugs would take more time to fix than I could justify using on them at this point.

write bug reports?

¹https://bugs.eclipse.org/bugs/show_bug.cgi?id=228635

²https://wiki.eclipse.org/JDT_UI/How_to_Contribute#Unit_Testing

The only thing that can be said in my defence, is that all the compilation errors could have been avoided if the types of situations that causes them were properly handled by the primitive refactorings, that again are supported by the Eclipse JDT UI project. All of the four randomly found bugs that I mentioned before, are also weaknesses of the *Move Method* refactoring. If the primitive refactorings had detected the up-coming errors in their precondition checking phase, the refactorings would have been aborted, since this is how the *Extract and Move Method* refactoring handles such situations.

Of course, taking all possible situations into account is an immense task. This is one of the reasons for the failure. A complete analysis is too big of a task for this master's project to handle. Looking at it now, this comes as no surprise, since the task is obviously also too big for the creators of the primitive *Move Method* refactoring. This shows that the underlying primitive refactorings are not complete enough to be fully relied upon for avoiding compilation errors.

Considering all these problems, it is difficult to know how to interpret the unit test results from after refactoring the Eclipse JDT UI. The **AutomatedSuite** reported 565 errors and 5 failures. Three of the failures were the same as reported before the refactoring took place, so two of them are new. For these two cases it is not immediately apparent what makes them behave differently. The program is so complex that to analyze it to find this out, we might need more powerful methods than just manually analyzing its source code. This is somewhat characteristic for imperative programming: The programs are often hard to analyze and understand.

For the **AllAllRefactoringTests** test suite, the three failures are gone, but the two errors have grown to 2,257 errors. I will not try to analyze those errors.

What I can say, is that it is likely that the *Extract and Move Method* refactoring has introduced some unintended behavioral changes. Let us say that the refactoring introduces at least two behavior-altering changes for every 2,500 executions. More than that is difficult to say about the behavior-preserving properties of the *Extract and Move Method* refactoring, at this point.

8.5.4 Conclusions

After automatically analyzing and executing the *Extract and Move Method* refactoring for all the methods in the Eclipse JDT UI project, the results does not look that promising. For this case, the refactoring seems almost unusable as it is now. The error rate and measurements done tells us this.

The refactoring makes the code a little less complex at the method level. But this is merely a side effect of extracting methods, and holds little scientific value. When it comes to the overall complexity, it is increased, although it is slightly better spread among the classes.

The analysis done before the *Extract and Move Method* refactoring, is currently not complete enough to make the refactoring useful. It introduces too many errors in the code, and the code may change it's behavior. It also

remains to prove that large scale refactoring with it can decrease coupling between classes. A better analysis may prove this, but in its present state, the opposite is the fact. The coupling measurements done by SonarQube shows this.

On the bright side, the performance of the refactoring process is not that bad. It shows that it is possible to make a tool the way we do, if we can make the tool do anything useful. As long as the analysis phase is not going to involve anything that uses too much disk access, a lot of analysis can be done in a reasonable amount of time.

The time used on performing the actual changes excludes a trial and error approach with the tools used in this master's project. In a trial and error approach, you could for instance be using the primitive refactorings used in this project to refactor code, and only then make decisions based on the effect, possibly shown by traditional software metrics. The problem with the approach taken in this project, compared to a trial and error approach, is that using heuristics beforehand is much more complicated. But on the other hand, a trial and error approach would still need to face the challenges of producing code that does compile without errors. If using refactorings that could produce in-memory changes, a trial and error approach could be made more efficient.

8.6 Case 2: The `no.uio.ifi.refaktor` project

In this case we will see a form of the “dogfooding” methodology used, when refactoring our own `no.uio.ifi.refaktor` project with the *Extract and Move Method* refactoring.

In this case I will try to point out some differences from case 1, and how they impact the execution of the benchmark. The `refaktor` project is 39 times smaller than the Eclipse JDT UI project, measured in lines of code. This will make things a bit more transparent. It will therefore be interesting to see if this case can shed light on any aspect of our project that were lost in the larger case 1.

The configuration for the experiment is specified in table 8.5 on page 103.

8.6.1 Statistics

The statistics gathered during the refactoring execution is presented in table 8.6 on page 104.

Differences

There are some differences between the two projects that make them a little difficult to compare by performance.

Different complexity. Although the JDT UI project is 39 times greater than the `refaktor` project in terms of lines of code, it is only about 26 times its size measured in numbers of methods. This means that the methods in the `refaktor` project are smaller in average than in the JDT project. This

is also reflected in the SonarQube report, where the complexity per method for the JDT project is 3.6, while the refactor project has a complexity per method of 2.1.

Number of selections per method. The analysis for the JDT project processed 21 text selections per method in average. This number for the refactor project is only 8 selections per method analyzed. This is a direct consequence of smaller methods.

Different candidates to methods ratio. The differences in how the projects are factored are also reflected in the ratios for how many methods that are chosen as candidates compared to the total number of methods analyzed. For the JDT project, 9% of the methods were considered to be candidates, while for the refactor project, only 5% of the methods were chosen.

The average number of possible candidate selection. For the methods that are chosen as candidates, the average number of possible candidate selections for these methods differ quite much. For the JDT project, the number of possible candidate selections for these methods was 14.44 selections per method, while the candidate methods in the refactor project had only 3.91 candidate selections to choose from, in average.

Execution time

The differences in complexity, and the different candidate methods to total number of methods ratios, is shown in the distributions of the execution times. For the JDT project, 75% of the total time was used on the actual changes, while for the refactor project, this number was only 63%.

For the JDT project, the benchmark used on average 0.21 seconds per method in the project, while for the refactor project it used only 0.07 seconds per method. So the process used 3 times as much time per method for the JDT project than for the refactor project.

While the JDT project is 39 times larger than the refactor project measured in lines of code, the benchmark used about 79 times as long time on it than for the refactor project. Relatively, this is about twice as long.

Since the details of these execution times are not that relevant to this master's project, only their magnitude, I will leave them here.

Executed refactorings

For the composite *Extract and Move Method* refactoring performed in case 2, 53 successful attempts out of 58 gives a success rate of 91.4%. This is 5.3 percentage points worse than for case 1.

8.6.2 SonarQube analysis

Results from the SonarQube analysis is shown in table 8.7 on page 105.

Not much is to be said about these results. The trends in complexity and coupling are the same. We end up a little worse after the refactoring process than before.

8.6.3 Unit tests

The tests used for this case are the same that has been developed throughout the master's project.

The code that was refactored for this case suffered from some of the problems discovered in case 1. This means that the after-code for case 2 also contained compilation errors, but they were not as many. The code contained only 6 errors that made the code not compile.

All of the errors made, originated from the same bug. It is a bug that happens in situation where a class instance creation is moved from between packages, and the class for the instance is package-private. The *Move Method* refactoring does not detect that there will be a visibility problem, and neither does it promote the package-private class to be public.

Since the errors were easy to fix manually, I corrected them and ran the unit tests as planned. Before the refactoring, all tests passed. All tests also passed after the refactoring, with the six error corrections. Since the corrections done is not of a kind that could make the behavior of the program change, it is likely that the refactorings done to the `no.uio.ifi.refaktor` project did not change its behavior. This is also supported by the informal experiment presented next.

8.6.4 An informal experiment

To complete the task of “eating my own dog food”, I conducted an informal experiment where I used the refactored version of the `no.uio.ifi.refaktor` project, with the corrections, to again refactor “itself”.

The experiment produced code containing the same six errors as after the previous experiment. I also compared the after-code from the two experiments with a diff-tool. The only differences found were different method names. This is expected, since the method names are randomly generated by the *Extract and Move Method* refactoring.

The outcome of this simple experiment makes me more confident that the *Extract and Move Method* refactoring made only behavior-preserving changes to the `no.uio.ifi.refaktor` project, apart from the compilation errors.

8.6.5 Conclusions

The differences in complexity between the Eclipse JDT UI project and the `no.uio.ifi.refaktor` project, clearly influenced the differences in their execution times. This is mostly because fewer of the methods were chosen to be refactored for the refactor project than for the JDT project. What this makes difficult, is to know if there are any severe performance penalties associated with refactoring on a large project compared to a small one.

The trends in the SonarQube analysis are the same for this case as for the previous one. This gives more confidence in the these results.

By refactoring our own code and using it again to refactor our code, we showed that it is possible to write an automated composite refactoring that works for many cases. That it probably did not alter the behavior of a smaller project shows us nothing more than that though, and might just be a coincidence.

8.7 Summary

Write? Or wrap up in final conclusions?

Table 8.3: Statistics after batch refactoring the Eclipse JDT UI project with the *Extract and Move Method* refactoring.

Time used	
Total time	98m38s
Analysis time	14m41s (15%)
Change time	74m20s (75%)
Miscellaneous tasks	9m37s (10%)
Numbers of each type of entity analyzed	
Packages	110
Compilation units	2,097
Types	3,152
Methods	27,667
Text selections	591,500
Numbers for <i>Extract and Move Method</i> refactoring candidates	
Methods chosen as candidates	2,552
Methods NOT chosen as candidates	25,115
Candidate selections (multiple per method)	36,843
<i>Extract and Move Method</i> refactorings executed	
Fully executed	2,469
Not fully executed	83
Total attempts	2,552
Primitive refactorings executed	
<i>Extract Method</i> refactorings	
Performed	2,483
Not performed	69
Total attempts	2,552
<i>Move Method</i> refactorings	
Performed	2469
Not performed	14
Total attempts	2,483

Table 8.4: Results for analyzing the Eclipse JDT UI project, before and after the refactoring, with SonarQube and the IFI Refaktor Case Study quality profile. (Bold numbers are better.)

Number of issues for each rule	Before	After
Avoid too complex class	81	79
Classes should not be coupled to too many other classes (Single Responsibility Principle)	1,098	1,199
Control flow statements . . . should not be nested too deeply	1,375	1,285
Methods should not be too complex	1,518	1,452
Methods should not have too many lines	3,396	3,291
NPath Complexity	348	329
Too many methods	454	520
Total number of issues	8,270	8,155
Complexity		
Per function	3.6	3.3
Per class	29.5	30.4
Per file	44.0	45.3
Total complexity	84,765	87,257
Numbers of each type of entity analyzed		
Files	1,926	1,926
Classes	2,875	2,875
Functions	23,744	26,332
Accessors	1,296	1,019
Statements	162,768	165,145
Lines of code	320,941	329,112
Technical debt (in days)	1,003.4	1,032.7

Table 8.5: Configuration for Case 2.

Benchmark data	
Launch configuration	CaseStudyDogfooding.launch
Project	no.uio.ifi.refaktor.benchmark
Repository	gitolite@git.uio.no:ifi-stolz-refaktor
Commit	43c16c04520746edd75f8dc2a1935781d3d9de6c
Input data	
Project	no.uio.ifi.refaktor
Repository	gitolite@git.uio.no:ifi-stolz-refaktor
Commit	43c16c04520746edd75f8dc2a1935781d3d9de6c
Branch	master
Test configuration	no.uio.ifi.refaktor.tests/ExtractTest.launch

Table 8.6: Statistics after batch refactoring the `no.uio.ifi.refaktor` project with the *Extract and Move Method* refactoring.

Time used	
Total time	1m15s
Analysis time	0m18s (24%)
Change time	0m47s (63%)
Miscellaneous tasks	0m10s (14%)
Numbers of each type of entity analyzed	
Packages	33
Compilation units	154
Types	168
Methods	1,070
Text selections	8,609
Numbers for <i>Extract and Move Method</i> refactoring candidates	
Methods chosen as candidates	58
Methods NOT chosen as candidates	1,012
Candidate selections (multiple per method)	227
<i>Extract and Move Method</i> refactorings executed	
Fully executed	53
Not fully executed	5
Total attempts	58
Primitive refactorings executed	
<i>Extract Method</i> refactorings	
Performed	56
Not performed	2
Total attempts	58
<i>Move Method</i> refactorings	
Performed	53
Not performed	3
Total attempts	56

Table 8.7: Results for analyzing the `no.uio.ifi.refaktor` project, before and after the refactoring, with SonarQube and the IFI Refaktor Case Study quality profile. (Bold numbers are better.)

Number of issues for each rule	Before	After
Avoid too complex class	1	1
Classes should not be coupled to too many other classes (Single Responsibility Principle)	29	34
Control flow statements ... should not be nested too deeply	24	21
Methods should not be too complex	17	15
Methods should not have too many lines	41	40
NPath Complexity	3	3
Too many methods	13	15
Total number of issues	128	129
Complexity		
Per function	2.1	2.1
Per class	12.5	12.9
Per file	13.8	14.2
Total complexity	2,089	2,148
Numbers of each type of entity analyzed		
Files	151	151
Classes	167	167
Functions	987	1,045
Accessors	35	30
Statements	3,355	3,416
Lines of code	8,238	8,460
Technical debt (in days)	19.0	20.7

Chapter 9

Conclusions and Future Work

This chapter will conclude this master's thesis. It will try to give justified answers to the research questions posed (see section 1.3.6 on page 32) and present some future work that could be done to take this project to the next level.

9.1 Conclusions

One of the motivations for this thesis was to create a fully automated composite refactoring that could be used to make program source code better in terms of coupling between classes. Earlier, in section 1.3.5, it was shown that a composition of the *Extract Method* and the *Move Method* refactorings reduces the coupling between two classes in an ideal situation. The Eclipse IDE implements both these refactorings, as well as providing a framework for analyzing source code, so it was considered a suitable tool to build upon for our project.

The search-based *Extract and Move Method* refactoring was created by utilizing the analysis and refactoring support of Eclipse, and a small framework was built for executing large scale refactoring with it. The refactoring was set up to analyze and execute changes on the Eclipse JDT UI project. Statistics was gathered during this process and the resulting code was analyzed through SonarQube. The project's own unit tests were also performed to find out if our refactoring introduces any behavior-altering changes in the code it refactor.

Answering the main research question The first and greatest challenge was to find out if the *Extract and Move Method* refactoring could be automated, in all tasks ranging from analysis to executing changes. It is now confirmed that this can be done, since it has been implemented as a part of the work done for this project. It has also been shown that the refactoring can be used to refactor large code bases, through the case study done on the Eclipse JDT UI project.

If we ask if using the existing Eclipse refactorings for this task is *easy*, this is another question. The refactorings provided by the JDT UI project are clearly not meant to be combined in any way. The preconditions for one refactoring are not always easily retrievable after the execution of another. Also, the refactorings are all assuming that the code they shall refactor is textualized. This means that the source code must be parsed between the executions of each refactoring. Another problem with this dependency on textual changes is that you cannot make a composition of two refactorings appear as one change if their changes overlap. This will make the undo-history of the refactoring show two changes instead of one, and is not nice for usability if the refactoring would be used as an on-demand refactoring in an IDE.

Apart from the problems with implementing the actual refactoring, the analysis framework is quite nicely solved in Eclipse. The AST generated when parsing source code supports using visitors to traverse it, and this works without problems.

Is the refactoring efficient enough? By being able to process over 300,000 pure lines of code in about 1.5 hours on a mid-level laptop computer, the search-based *Extract and Move Method* refactoring must be said to perform reasonably well.

What about breaking the source code? The case studies showed that our safety measures that rely on the precondition checking of the existing primitive refactorings are not good enough in practice. If we were going to assure that code we change compiles, we would need to consider all possible situations where the refactoring could fail and search for them in our analysis. It is an open question if this is even feasible. Our analysis is incomplete, and so is the analysis for the *Extract Method* and the *Move Method* refactorings.

Our refactoring does not take any precautions to preserve behavior. A few running and failing unit test for the JDT UI project after the refactoring indicate that our refactoring probably causes some changes to the way a program behaves.

check out test report

Is the quality of the code improved? For coupling, there is no evidence that the refactoring improves the quality of source code. Shall we believe the SonarQube analysis from the case studies, our refactoring makes classes more coupled after the refactoring than before, in the general case. This is probably because our analysis and heuristics for finding the best candidates for the refactoring are not adequate.

Is the refactoring useful? In its present state, the refactoring cannot be said to be very useful. It generates too many compilation errors for it to fall into that category. On the other hand, if the problems with the search-based *Extract and Move Method* refactoring were to be solved it could be useful in some situations.

If the refactoring was perfected, it could of course be used as a regular on-demand automated refactoring on a per method base (or per class, package or project).

As it is now, the refactoring is not very well suited to be set to perform unattended refactoring. But if we could find a way to filter out the changes that create compilation errors, we could use the refactoring to look for improvement points in a software project. This process could for instance be scheduled to run at regular times, possibly after a nightly build or the like. Then the results could be made available, and an administrator could be set to review them and choose whether or not they should be performed.

9.2 Future work

Find out if a complete analysis is feasible

Complete the analysis

Make refactorings safer (behavior)

Improve heuristics/introduce metrics

Appendix A

Eclipse Bugs Submitted

A.1 Eclipse bug 420726: Code is broken when moving a method that is assigning to the parameter that is also the move destination

This bug was found when analyzing what kinds of names that were to be considered as *unfixes* (see section 5.4.2 on page 72).

The bug The bug emerges when trying to move a method from one class to another, and when the target for the move (must be a variable, local or field) is both a parameter variable and also is assigned to within the method body. Eclipse allows this to happen, although it is the sure path to a compilation error. This is because we would then have an assignment to a **this** expression, which is not allowed in Java. The submitted bug report can be found on https://bugs.eclipse.org/bugs/show_bug.cgi?id=420726.

The solution The solution to this problem is to add all simple names that are assigned to in a method body to the set of unfixes.

A.2 Eclipse bug 429416: IAE when moving method from anonymous class

I discovered this bug during a batch change on the `org.eclipse.jdt.ui` project.

The bug This bug surfaces when trying to use the *Move Method* refactoring to move a method from an anonymous class to another class. This happens both for my simulation as well as in Eclipse, through the user interface. It only occurs when Eclipse analyzes the program and finds it necessary to pass an instance of the originating class as a parameter to the moved method. I.e. it want to pass a **this** expression. The execution ends in an `IllegalArgumentException`¹ in

¹`java.lang.IllegalArgumentException`

`SimpleName`¹ and its `setIdentifier(String)` method. The simple name is attempted created in the method `createInlinedMethodInvocation`² so the `MoveInstanceMethodProcessor` was early a clear suspect.

The `createInlinedMethodInvocation` is the method that creates a method invocation where the previous invocation to the method that was moved was located. From its code it can be read that when a `this` expression is going to be passed in to the invocation, it shall be qualified with the name of the original method's declaring class, if the declaring class is either an anonymous class or a member class. The problem with this, is that an anonymous class does not have a name, hence the term *anonymous* class! Therefore, when its name, an empty string, is passed into `newSimpleName`³ it all ends in an `IllegalArgumentException`. The submitted bug report can be found on https://bugs.eclipse.org/bugs/show_bug.cgi?id=429416.

How I solved the problem Since the `MoveInstanceMethodProcessor` is instantiated in the `MoveMethodRefactoringExecutor`⁴, and only need to be a `MoveProcessor`⁵, I was able to copy the code for the original move processor and modify it so that it works better for me. It is now called `ModifiedMoveInstanceMethodProcessor`⁶. The only modification done (in addition to some imports and suppression of warnings), is in the `createInlinedMethodInvocation`. When the declaring class of the method to move is anonymous, the `this` expression in the parameter list is not qualified with the declaring class' (empty) name.

A.3 Eclipse bug 429954: Extracting statement with reference to local type breaks code

The bug was discovered when doing some changes to the way unfixes is computed.

The bug The problem is that Eclipse is allowing selections that references variables of local types to be extracted. When this happens the code is broken, since the extracted method must take a parameter of a local type that is not in the methods scope. The problem is illustrated in listing 10 on page 47, but there in another setting. The submitted bug report can be found on https://bugs.eclipse.org/bugs/show_bug.cgi?id=429954.

Actions taken There are no actions directly springing out of this bug, since the Extract Method refactoring cannot be meant to be this way. This is handled on the analysis stage of our *Extract and Move Method* refactoring.

¹`org.eclipse.jdt.core.dom SimpleName`

²`org.eclipse.jdt.internal.corext.refactoring.structure.`

`MoveInstanceMethodProcessor#createInlinedMethodInvocation()`

³`org.eclipse.jdt.core.dom.AST#newSimpleName()`

⁴`no.uio.ifi.refaktor.change.executors.MoveMethodRefactoringExecutor`

⁵`org.eclipse.ltk.core.refactoring.participants.MoveProcessor`

⁶`no.uio.ifi.refaktor.change.processors.ModifiedMoveInstanceMethodProcessor`

So names representing variables of local types is considered unfixes (see section 5.4.2 on page 72).

write more when fixing this in legal statements checker

Glossary

design pattern A design pattern is a named abstraction, that is meant to solve a general design problem. It describes the key aspects of a common problem and identifies its participators and how they collaborate. 16

enclosing class An enclosing class is the class that surrounds any specific piece of code that is written in the inner scope of this class. 41

Extract Class The *Extract Class* refactoring works by creating a class, for then to move members from another class to that class and access them from the old class via a reference to the new class. 22

memento pattern The memento pattern is a software design pattern that is used to capture an object's internal state so that it can be restored to this state later [Gam+95]. 87

profiler A profiler is a program for analyzing performance within an application. It is used to analyze memory consumption, processing time and frequency of procedure calls and such. 85

profiling is to run a computer program through a profiler/with a profiler attached. 24, 85

software obfuscation makes source code harder to read and analyze, while preserving its semantics. 16

xUnit framework An xUnit framework is a framework for writing unit tests for a computer program. It follows the patterns known from the JUnit framework for Java [Fow]. 28

References

- [11] *JAVA EE Productivity Report 2011*. Survey. 2011. URL: http://zeroturnaround.com/wp-content/uploads/2010/11/Java_EE_Productivity_Report_2011_finalv2.pdf.
- [14a] *Advice (programming)*. In: *Wikipedia, the free encyclopedia*. Page Version ID: 462233199. Mar. 14, 2014. URL: [https://en.wikipedia.org/w/index.php?title=Advice_\(programming\)&oldid=462233199](https://en.wikipedia.org/w/index.php?title=Advice_(programming)&oldid=462233199) (visited on 03/21/2014).
- [14b] *Continuous design*. In: *Wikipedia, the free encyclopedia*. Page Version ID: 544105069. Apr. 8, 2014. URL: https://en.wikipedia.org/w/index.php?title=Continuous_design&oldid=544105069 (visited on 04/09/2014).
- [Bro04] Leo Brodie. *Thinking Forth*. 3rd ed. 2004. URL: <http://thinking-forth.sourceforge.net/>.
- [CK94] S.R. Chidamber and C.F. Kemerer. “A Metrics Suite for Object Oriented Design.” In: *IEEE Transactions on Software Engineering* 20.6 (June 1994), pp. 476–493.
- [Dem02] Serge Demeyer. “Maintainability Versus Performance: What’s the Effect of Introducing Polymorphism?” In: *ICSE’2003* (2002).
- [Fow] Martin Fowler. *Xunit*. URL: <http://www.martinfowler.com/bliki/Xunit.html> (visited on 03/27/2014).
- [Fow01] Martin Fowler. *Crossing Refactoring’s Rubicon*. 2001. URL: <http://martinfowler.com/articles/refactoringRubicon.html> (visited on 02/09/2014).
- [Fow03] Martin Fowler. *EtymologyOfRefactoring*. Sept. 10, 2003. URL: <http://martinfowler.com/bliki/EtymologyOfRefactoring.html> (visited on 03/20/2014).
- [Fow04] Martin Fowler. *Is Design Dead?* 2004. URL: <http://martinfowler.com/articles/designDead.html> (visited on 04/09/2014).
- [Fow99] Martin Fowler. *Refactoring: improving the design of existing code*. Reading, MA: Addison-Wesley, 1999.
- [Gam+95] Erich Gamma et al. *Design patterns: elements of reusable object-oriented software*. Reading, MA: Addison-Wesley, 1995.

- [Har06] W. Harrison. “Eating Your Own Dog Food.” In: *IEEE Software* 23.3 (May 2006), pp. 5–7.
- [Ker05] Joshua Kerievsky. *Refactoring to patterns*. Boston: Addison-Wesley, 2005.
- [Lou97] Kenneth C Louden. *Compiler construction: principles and practice*. Boston: PWS Pub. Co., 1997.
- [MC09] Robert C Martin and James O Coplien. *Clean code: a handbook of agile software craftsmanship*. Upper Saddle River, NJ [etc.]: Prentice Hall, 2009.
- [McC76] T.J. McCabe. “A Complexity Measure.” In: *IEEE Transactions on Software Engineering* SE-2.4 (Dec. 1976), pp. 308–320.
- [Mey88] Bertrand Meyer. *Object-oriented software construction*. Prentice-Hall, 1988.
- [Mil56] George A. Miller. “The magical number seven, plus or minus two: some limits on our capacity for processing information.” In: *Psychological Review* 63.2 (1956), pp. 81–97.
- [Nic06] Ethan Nicholas. *Understanding Weak References*. Java.net. May 4, 2006. URL: <https://weblogs.java.net/blog/2006/05/04/understanding-weak-references> (visited on 03/20/2014).
- [OC08] Mark O’Keeffe and Mel Ó Cinnéide. “Search-based Refactoring: An Empirical Study.” In: *J. Softw. Maint. Evol.* 20.5 (Sept. 2008), 345–364. URL: <http://dx.doi.org/10.1002/smr.v20:5>.
- [Opd92] William F. Opdyke. “Refactoring Object-oriented Frameworks.” UMI Order No. GAX93-05645. Champaign, IL, USA: University of Illinois at Urbana-Champaign, 1992.
- [RBJ97] Don Roberts, John Brant, and Ralph Johnson. “A Refactoring Tool for Smalltalk.” In: *Theor. Pract. Object Syst.* 3.4 (Oct. 1997), 253–263.
- [Sha10] R. Shatnawi. “A Quantitative Investigation of the Acceptable Risk Levels of Object-Oriented Metrics in Open-Source Systems.” In: *IEEE Transactions on Software Engineering* 36.2 (Mar. 2010), pp. 216–225.
- [Soa+10] G. Soares et al. “Making Program Refactoring Safer.” In: *IEEE Software* 27.4 (Aug. 2010), pp. 52–57.
- [Sto] Volker Stolz. *Project "Safer Refactorings"*. URL: http://heim.ifi.uio.no/stolz/Volker_Stolz_-_UiO/Project_Safer_Refactorings.html (visited on 04/26/2014).
- [Vak+12] Mohsen Vakilian et al. *A Compositional Paradigm of Automating Refactorings*. May 2012. URL: <https://www.ideals.illinois.edu/bitstream/handle/2142/30851/VakilianETAL2012Compositional.pdf?sequence=4>.

- [VJ12] Mohsen Vakilian and Ralph Johnson. *Composite Refactorings: The Next Refactoring Rubicons*. University of Illinois at Urbana-Champaign, 2012. URL: <https://www.ideals.illinois.edu/bitstream/handle/2142/35678/2012-WRT.pdf?sequence=2>.
- [Vog12] Lars Vogel. *Eclipse JDT - Abstract Syntax Tree (AST) and the Java Model - Tutorial*. vogella. Aug. 8, 2012. URL: <http://www.vogella.com/tutorials/EclipseJDT/article.html> (visited on 04/20/2014).

Todo list

Remove all todos (including list) before delivery/printing!!!	
Can be done by removing “draft” from documentclass. . . .	3
Write abstract	3
Structure. Write later...	15
Highlight code	32
Write	35
Write	35
State somewhere that we assume no compilation errors?	43
Expand with more illegal statements and/or conclude that I did not have time to analyze all statement types.	44
Short description of changing code and error handling?	51
Pointing to implementation chapter	51
What about the language specific part?	53
refine	55
Elaborate?	60
Write about the ExtractAndMoveMethodCandidateComparator/Fa- vorNoUnfixesCandidateComparator	62
Where to put this section?	63
Add more to the AST format tree? fig. 5.4 on page 69	68
Give more technical detail?	72
Remove section?	72
File Eclipse bug report	73
Follow the development in the semantics section...	75
write	87
How long are these results going to be publicly available?	91
Redundant information? Put in benchmark chapter?	93
write bug reports?	95
Write? Or wrap up in final conclusions?	100
check out test report	108
Find out if a complete analysis is feasible	109
Complete the analysis	109
Make refactorings safer (behavior)	109
Improve heuristics/introduce metrics	109
write more when fixing this in legal statements checker	113