

A Model Driven Approach for Refactoring Heterogeneous Software Artefacts

Keith Dooley

A dissertation submitted in partial fulfilment
of the requirements for the degree of
Master of Science
at
Maynooth University.



Faculty of Science & Engineering, Department of Computer Science

February 2016

Head of Department: Dr. Adam C. Winstanley

Supervisors: Dr. Rosemary Monahan, Dr. James F. Power

Contents

1	Introduction	7
1.1	Motivation	7
1.2	Thesis Statement	8
1.3	Contributions	9
1.4	Organisation of Dissertation	10
2	Towards a Model Driven Approach to Refactoring	12
2.1	Research Questions	12
2.2	Refactoring	12
2.2.1	Refactoring Operations	13
2.2.2	When to Refactor & Considerations for Developing a Refactoring Engine	16
2.2.3	Manual, Automated and Search-based Refactoring	17
2.3	Formal Specification & Design by Contract	18
2.4	Model Driven Engineering	20
2.4.1	Metamodel Hierarchy and Transformations	20
2.4.2	MDE Environments	20
2.4.3	Domain Specific (Transformation) Languages	22
2.4.4	Modelling Java Programs	23
2.5	Software Metrics	23
2.5.1	Project Metrics	24
2.5.2	Design Metrics	24
2.5.3	Validation of Metrics	25
2.5.4	Metric Tools	25
2.5.5	Metrics to Identify Refactorings	25
2.5.6	Summary	26
3	RefDstl: A Language for Specifying and Scripting Refactorings	27
3.1	Use Cases	27
3.2	Characteristics of a Good Programming Language	28
3.3	Design of a DSTL for Defining and Scripting Refactorings	29
3.3.1	Choosing a Suitable Paradigm	30
3.3.2	Language Constructs	31
3.4	RefDstl Implementation	32
3.4.1	Parser Generators and Grammars	33

3.4.2	A Grammar for RefDstl	33
3.4.3	Substituting RefDstl	34
3.5	Defining a Standard Library of Java Refactorings	34
3.6	A Small Case Study	34
3.7	Summary	41
4	A Model Driven Framework to Interpret RefDstl Scripts	43
4.1	High-Level Overview of the RefDstl Interpreter Framework	43
4.2	Enriching the RefDstl Language via Script Handles	49
4.2.1	Embedding a New Handle	49
4.2.2	Supporting Debugging in RefDstl Scripts	50
4.2.3	Checking Object Types	50
4.2.4	Accessing the Program Model	50
4.2.5	Accessing Old Variables	51
4.3	Custom Transformations into the RefDstl Framework	51
4.4	Disadvantage of Our Approach	52
4.5	Summary	53
5	A Case Study on Automatically Refactoring ‘God’ Classes	54
5.1	Purpose of Experiment	54
5.2	Summary of Approach	55
5.3	Identifying a Corpus for the Experiment	55
5.4	Content of the Qualitas Corpus	56
5.5	An Analysis Tool for the Qualitas Corpus	57
5.5.1	Persisting the Metadata	58
5.5.2	Analysing the Classes	58
5.6	Selecting Refactoring Candidates	60
5.6.1	Creating the MoDisco Models	60
5.7	An Extract Class Refactoring	63
5.8	Results	64
5.8.1	Predicting CBO	64
5.8.2	Time Taken	67
5.9	Discussion	68
5.10	Summary	69
6	Conclusion	70
6.1	Summary of Contributions	70
6.2	Future Work	71
	Acronyms	72
	A RefDstl Language Grammar for Xtext	74
	B RefDstl Standard Library	80

List of Tables

4.1	The interface that must be implemented by a class to integrate it into the RefDstl framework as a script handle.	50
4.2	The methods that must be implemented by the transformation interface.	52
5.1	Systems and number of versions provided in the Qualitas Corpus version 20130901e.	58
5.2	API for interacting with the clustering plugin.	63

List of Figures

1.1	An overview demonstrating how the components of the refactoring system presented in this dissertation interact.	11
2.1	This graph shows how the size of the Ant software system grows twelve-fold in terms of lines-of-code between versions 1.1 and 1.8.4.	14
2.2	Primitive refactoring operations according to Opdyke (1992).	16
2.3	An example showing the four layers of the metamodel hierarchy.	21
3.1	A UML class diagram demonstrating a poorly designed system to be refactored.	36
3.2	A UML class diagram that represents the system presented in Figure 3.1 refactored to obtain maximum inheritance.	37
3.3	The UML class diagram from Figure 3.1 refactored for minimum methods per class.	40
3.4	The UML class diagram from Figure 3.1 refactored using the optimisations calculated by the Dearthóir tool.	42
4.1	An overview of the components in the RefDstl interpreter framework.	44
5.1	The number of classes per system.	57
5.2	An entity relationship diagram illustrating the main entities and attributes persisted by the static analysis tool. Other entities, attributes and keys have been omitted for simplicity.	59
5.3	SQL query to discover ‘God’ classes in all versions.	61
5.4	SQL query to discover classes for the study.	61
5.5	The percentage of ‘God’ classes for all versions of each of the software systems in the corpus.	62
5.6	RefDstl sourcecode used to perform the extract class refactoring on an identified ‘God’ class.	65
5.7	Pearson’s correlation between the predicted CBO metric and the actual value	66
5.8	The distribution of the original, predicted post refactoring and actual post refactoring CBO measurements for the software systems used in the study.	67
5.9	Pearson’s correlation between the size of the model being refactored and the time taken to refactor the model in nanoseconds.	68

Abstract

Refactoring is the process of transforming a software system to improve its overall structure while preserving its observable behaviour. Refactoring engines are normally used to perform these transformations for efficiency and in order to avoid introducing behavioural changes into the program due to human error. Although these engines do not verify that behaviour is preserved, it is widely accepted that automated transformations are less likely to introduce errors in comparison to manual refactoring. Despite the advantages provided by refactoring engines they fall foul of certain weaknesses.

Here we hypothesise that Model Driven Engineering can be used to produce improved refactoring engines that are less vulnerable to those weaknesses. We develop a Domain Specific Transformation Language for defining new composite refactorings from a set of built-in primitives and to script their application. We also develop an interpreter for the language, effectively providing an operational semantics, in the guise of an extensible transformation framework. We evaluate our approach with a case study examining the correlation between actual and predicted measurements of the Coupling Between Objects metric for classes that undergo the extract class refactoring. The results show that our approach is promising.

Chapter 1

Introduction

This chapter presents the motivation for the research presented in this dissertation. It states explicitly our thesis and provides an overview of our contributions. It concludes with a ‘roadmap’ that discusses the organisation of the remaining chapters.

1.1 Motivation

Opdyke (1992) presented refactoring as a methodical approach to improve the *structure* of a software system’s source code while preserving its observable behaviour. It is often carried out by developers as an intermediary process, so that adding new functionality to the system is less laborious, or during preventative maintenance, so that the codebase remains comprehensible. Refactoring is an inevitable process for software that often changes, such as what Lehman and Fernández-Ramil (2006) categorise as ‘E’ type software. The process of refactoring involves applying structural transformations to the code and consists of two stages. The first stage is checking that a set of conditions hold and the second stage is actually performing the transformation. The purpose of the conditions is to ensure that the refactoring preserves the program’s behaviour. The set of conditions applicable to a refactoring varies for each transformation and for the programming language that is used to implement the system under refactoring. We call atomic transformations *primitive refactorings* and when composed these form *composite refactorings*.¹

Although manual refactoring is possible, and done often, it is generally accepted that it is more susceptible to undesired side effects due to human error. Many of these cannot be detected by the compiler so, although the system compiles, its behaviour will have changed. This leads to unpredictable behaviour that is costly to debug and rectify. If the errors are not detected until the system is released to the client then there may be other, potentially irreversible, consequences such as data loss.

Refactoring engines support the refactoring process and offer a semi-automated approach. Interaction between the developer and the refactoring engine is minimal. The developer guides the engine regarding *what* refactoring to apply and *where* it is to be applied. There is no longer a need for the developer to check that the necessary conditions are satisfied or that all the relevant areas of the codebase are updated. Despite the benefits afforded by refactoring engines they fall foul of these shortcomings, which we refer to throughout this dissertation as the four Is²:

¹It will be clear from the context when we use the word refactoring to mean the activity of refactoring or a specific transformation.

²These are numbered so later when we refer to I1, for example, we mean the first in this list.

1. **Inaccurate** — The development of a complete and sound set (i.e. all conditions are considered and all conditions are correct) of refactoring preconditions is a non-trivial task and is dependent on the programming language being refactored. This claim is supported by Brett et al. (2007) who have shown that even the most widely used refactoring engines are susceptible to having overly weak or strong preconditions. There is no support for altering the preconditions once the refactoring engine has been compiled so users cannot fix erroneous conditions in ‘the field’.
2. **Inextensible** — Refactoring engines offer limited refactorings to the user. Users cannot define their own refactorings and save them for later application. Popular refactoring engines offer only the primitive refactoring operations, which the user must compose time and time again to perform more useful refactorings.
3. **Inflexible** — Most refactoring engines are capable of refactoring just a single software artefact: source code. Nowadays this is no longer sufficient because software engineers work in an environment that consists of multiple heterogeneous artefacts. For example, there are artefacts that represent high level design documentation such as Unified Modelling Language (UML) class diagrams and artefacts that might even be involved during the system’s runtime, such as database or Extensible Markup Language schemata. These are often interdependent and refactoring one requires changes to the others. Refactoring engines are limited in their capabilities in this respect. Some refactoring engines, such as that provided by the Eclipse IDE, are able to update some aspects of JavaDoc annotations while refactoring Java source code. However they do little more beyond this.
4. **Improvident** — Ge and Murphy-Hill (2014) suggest some developers ignore refactoring engines because they do not trust the changes refactoring engines make. We believe they continue to refactor manually because it enables them to see changes being made stepwise. This means that if the refactoring appears to be having a negative impact on other aspects of the system the developer can change his or her course of action. Common refactoring engines restrict the developer from doing this and changes are made irrespective of undesired effects. For example, refactoring engines will allow a developer to stupidly move all of a system’s functionality into a single class despite the fact this degrades the quality of the system’s design.

Our thesis, outlined next, addresses these shortcomings.

1.2 Thesis Statement

We believe there is a need to research and implement ‘next generation’ refactoring engines that are better suited for contemporary software engineering environments populated by heterogeneous artefacts. This sets the stage for our thesis:

Improvements to current refactoring engines can be made through an approach that ties together Model Driven Engineering (MDE) and formal methods (specifically Design by Contract (DBC)³) and encourages developers to define their own reusable

³We note that the term ‘design by contract’ is trademarked but we use it in this dissertation when referring to approaches that share the same concepts.

and shareable libraries of refactorings.

We now briefly remark on the rationale for our thesis but delve into greater detail in Chapter 2.

MDE is an approach to software engineering where development is focussed on creating shareable models represented in some core format, which can be converted via model transformations into other representations. This might include source code. This is in contrast to earlier approaches where the goal is to represent classes of real world entities directly in source code. Interestingly, however, MDE has parallels with earlier Computer Aided Software Engineering (CASE) approaches. In the CASE approach developers modelled a system using a visual language such as UML and then forward engineered those models into source code. Changes made to the source code would then be reflected in the UML diagrams. The purpose was to keep the documentation and the code synchronised. The difference between MDE and CASE lies with CASE being only concerned with two representations. The general approach being taken remains the same. That is to say, there is a further step in the direction of developers working at higher levels of abstraction. MDE is discussed in greater detail in the following chapter.

DBC is a methodology for code development that has been successfully used to improve software reliability. Comparisons are commonly drawn between the approach taken in DBC and transactions in the business world. In the DBC approach, there exists a contract between the client who uses code in an arbitrary library and the provider of that library. The contract describes what conditions the client is obliged to fulfil in order for the code in the library to execute correctly. The correct execution of the code is also described in terms of conditions. The approach we propose to refactoring requires that developers define their own refactorings and the conditions that must be true about the *structure* of the program in order for the refactoring to be successfully applied. We briefly discuss how this will be done when discussing our contributions in the next section but a more in depth discussion is provided in the subsequent chapter.

1.3 Contributions

The work presented in this dissertation does not offer a complete solution to developing the next generation refactoring engines. However, we make the following contributions in order to evaluate our thesis. These are provided as prototype tools and a complete environment with the necessary set up and configuration is available by installing the VirtualBox disk image stored on the electronic medium disseminated with this dissertation.

- RefDstl Language — We have developed a prototype DSTL that can be used to specify and script refactorings on models of software systems in terms of a set of basic primitive operations. The language and requisite tools, such as a parser and text editor with syntax highlighting and content assist, are produced using MDE technologies. The notion of a DSTL also has roots in MDE. Design by contract features heavily in the language. The language targets the inaccuracy (11) and inextensible (12) shortcomings raised above. Although developing the necessary conditions is a difficult feat, we can at least give the end users the opportunity to weaken or strengthen the conditions being checked by the refactoring engine in ‘the field’ rather than having to look into the entire source code of the engine. Furthermore, refactorings written in the language can be saved to a script for later use.

- **RefDstl Transformation Framework** — This framework serves as the interpreter for RefDstl scripts. However, it is richer in capability than traditional interpreters. It is an extensible framework that allows others to plug in their own transformations (which are driven by refactorings) and allows artefacts other than source code to be refactored simultaneously. We have provided two noteworthy plug-ins for the framework among others. The first refactors the software system; it relies on the MoDisco⁴ metamodel for representing Java systems. A model of the system under refactoring, conforming to the MoDisco metamodel, is transformed during the refactoring. When the framework is equipped with just this transformation then it behaves as a regular refactoring engine. The second plug-in enriches the framework with a notion of code quality. It allows developers refactoring to include requirements about the system’s code quality in their refactoring conditions. This tackles the problem of improvidence (I4) described above. This plug-in uses the approach by McQuillan (2011). We perform a model transformation from the MoDisco model to McQuillan’s metrics metamodel. This happens in real time while the system model is being refactored. We allow for plug-ins to be reasoned about in the refactoring conditions, which allows the developer who defined the refactoring to prevent refactorings being applied where the quality of the system (determined by metric measurements) deteriorates. By allowing further transformations to be plugged into the framework we permit flexibility. This addresses I3.
- **Standard Library of Primitive Refactorings** — We provide a standard library of primitive refactoring operations based on the same described by Opdyke (1992). It does not directly address the shortcomings addressed above but the need for it will become apparent later.

Figure 1.1 depicts a visualisation of how our framework operates. The portion developed by us includes what is inside the box entitled ‘RefDstl Transformation Framework’ and the RefDstl metamodel. We did not develop the MoDisco metamodel or contribute to the MoDisco project in any way. The diagram shows that the process of refactoring begins with a Java software system that is transformed into a model that conforms to the MoDisco metamodel. This can be done within an Eclipse Integrated Development Environment (IDE) equipped with the MoDisco plug-ins. The resulting model is then consumed by the RefDstl framework along with a script written in the RefDstl scripting language that conforms to the RefDstl language metamodel. The framework applies a sequence of transformations to the model of the system under refactoring. This yields one or more resulting output models, the most important of which is probably the refactored system that, again, conforms to the MoDisco metamodel. The refactored system can be transformed back into Java source code using a model-to-text transformation. However, this is outside the scope of this dissertation since we concern ourselves only with models that can be manipulated algorithmically.

We discuss the implementation and evaluation of these components throughout the dissertation. The ‘roadmap’ presented in the next section describes the structure of this document.

1.4 Organisation of Dissertation

The sequel is organised as follows:

⁴<https://eclipse.org/MoDisco/>

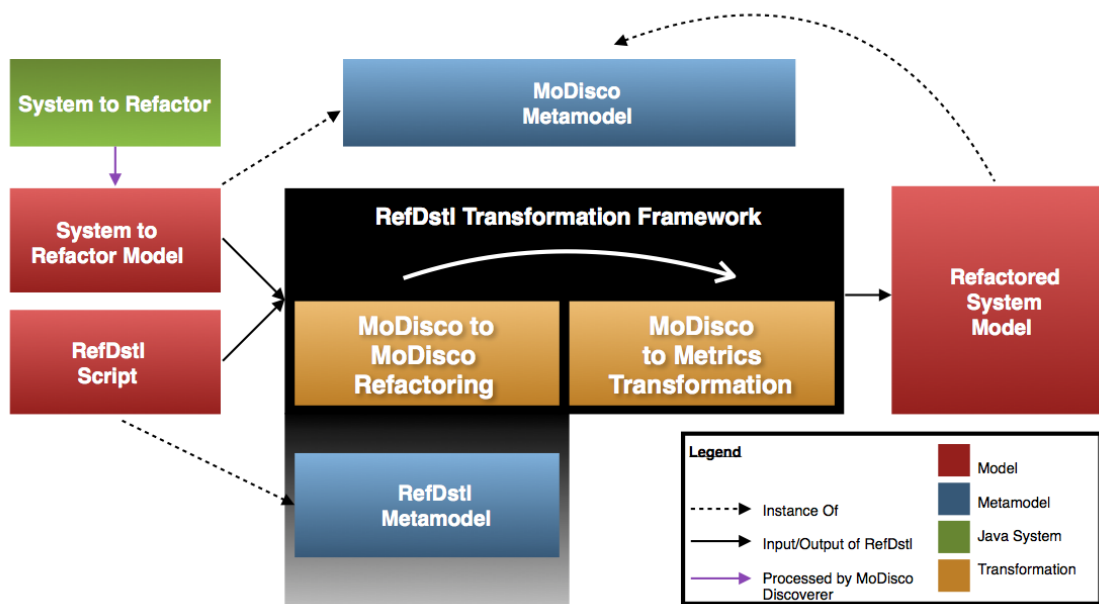


Figure 1.1: An overview demonstrating how the components of the refactoring system presented in this dissertation interact.

In Chapter 2, the necessary background material is provided. It addresses some of the assumptions that have been made thus far including the relevance of refactoring in software engineering circles. We also expand on our discussion of MDE and DBC. Furthermore, we discuss an alternative approach to ensure software quality, i.e. software metrics. That chapter comes to a conclusion with a discussion about how other researchers have used metrics to identify where refactoring should occur. This is followed in Chapter 3 with a discussion regarding the implementation of the RefDstl language. We discuss potential use cases for the language as well as language design goals. The implementation of the language is described and the chapter ends with language examples and a brief summary of the RefDstl standard library. In Chapter 4, we discuss the design and implementation of the RefDstl framework for interpreting RefDstl scripts as well as how to develop plug-ins that expand its capabilities. We evaluate the efficacy of our approach in Chapter 5 with a case study over a corpus of open-source Java software. Chapter 6 concludes the dissertation with a summary of our findings and a promise to investigate further avenues of research that arise from this body of work.

The appendices at the end contain supporting material. A list of acronyms is provided on page 71. In the electronic version of this document, the important terms and acronyms are hyperlinked to their expanded form. Appendix A contains the grammar of the RefDstl language to resolve any ambiguities concerning what is and is not a RefDstl program. Appendix B contains the full text of the RefDstl standard library.

Chapter 2

Towards a Model Driven Approach to Refactoring

In Chapter 1, we stated our thesis that MDE and DBC could be combined to address the shortcomings of current refactoring engines that we termed the ‘four Is’. While our ultimate goal in this dissertation is to demonstrate the validity of this thesis, there are three research questions that we address prior to implementing a prototype sufficient for evaluation. We open this chapter by establishing those questions, which paves the way for the literature review that follows. This review allows us to identify weaknesses concerning the state-of-the-art and prevents us reinventing the wheel. The literature review covers the topics of: refactoring, DBC, MDE, and software metrics. This chapter also affords us the opportunity to introduce nomenclature and definitions.

2.1 Research Questions

The following questions drive the remainder of this chapter and the answers form the rationale for the implementation choices of our solution:

- RQ1 What approaches have been taken to ensure that a system’s behaviour is preserved post refactoring?
- RQ2 What approaches have been taken to allow software engineers to define their own reusable refactorings?
- RQ3 What considerations are paid to the overall quality of a software system while refactoring?

The connection between these questions and the ‘four Is’ should be apparent. RQ1 is related to I1, while RQ2 relates to I2. Both I3 and I4 are associated with RQ3.

We now examine the literature and begin with the topic of refactoring to address RQ1.

2.2 Refactoring

During its lifetime, a software system grows considerably to adopt and adapt to changing and new requirements. Evidence supporting this claim is offered by the data illustrated in Figure 2.1, which was collated from data gathered by Tempero et al. (2010). It shows that after passing twenty-three versions the Ant¹ software system grows twelve-fold from 20,832 to 255,690 lines of code. This is not an isolated event. All but one of the fifteen systems considered in the data gathered by

¹<http://ant.apache.org>

Tempero et al. (2010) exhibits growth as it evolves. The ‘black sheep’ is JGraph², which reduces by an insignificant 176 lines of code.

Maintenance activities occur so frequently with software that it is estimated by Glass (2001) that they account for between 40% — 80% of software costs. In order to minimise the time, and consequently cost, of implementing changes in these rapidly growing systems it is highly important that the code is comprehensible to the responsible engineers. Approaches to software engineering can make writing comprehensible code quite a feat because many development approaches are iterative; from early approaches such as the spiral model by Boehm (1988) to more recent approaches such as extreme programming by Beck (1999). These iterative processes result in code being added repeatedly until the structure of the codebase begins to rot.

One solution to code rot is to engage in total productive maintenance (an idea admonished by Martin (2003)) — the code should be re-worked frequently to improve its structure. This however can result in regressions when defects are erroneously introduced. Refactoring is a methodical activity for improving code structure during which, code transformations are applied. However, these transformations are restricted to a set (*called refactorings*) that are widely believed to preserve observable behaviour.

Recent literature concerning refactoring has focussed on automating approaches to refactoring: the development of algorithms for identifying regions of the code that require refactoring, development and evaluation of refactoring engines that automatically apply refactorings and search based refactoring techniques that view the task of finding the best refactoring candidate as a search-space problem. Other work, closely related to RQ1, has focussed on proving the correctness of refactoring operations with respect to the observable behaviour of the program. These topics are surveyed in this section and an introduction to common refactoring operations is presented.

2.2.1 Refactoring Operations

Refactoring operations are the transformations applied to source code to improve its structure with the caveat that they preserve the observable behaviour of the program. Some of these operations can be viewed as atomic, which we call ‘primitive’ or ‘elementary’ refactorings while more complex and generally more useful molecular operations are known as ‘composite’ refactorings.

Primitive Operations & Behaviour Presentation

Primitive operations are the smallest atomic refactoring operations that can be applied to a program. The original body of work published on refactoring by Opdyke (1992) specified twenty-three primitives. These are listed in Figure 2.2. The intention of each should be self explanatory. However, the method for applying each of these is dependent on the programming language being refactored and how the program is represented. For example, if the program is represented as an abstract syntax tree then the refactoring will be performed using tree manipulations.

Not all primitive operations are relevant to every language. For example, Java omits pointers so the ‘convert instance variable to pointer’ refactoring is inapplicable. The operations were described by Opdyke (1992) with conditions that must be met regarding the structure of the program if the refactoring is to preserve behaviour. These are called preconditions. We draw attention

²<http://www.jgraph.com>

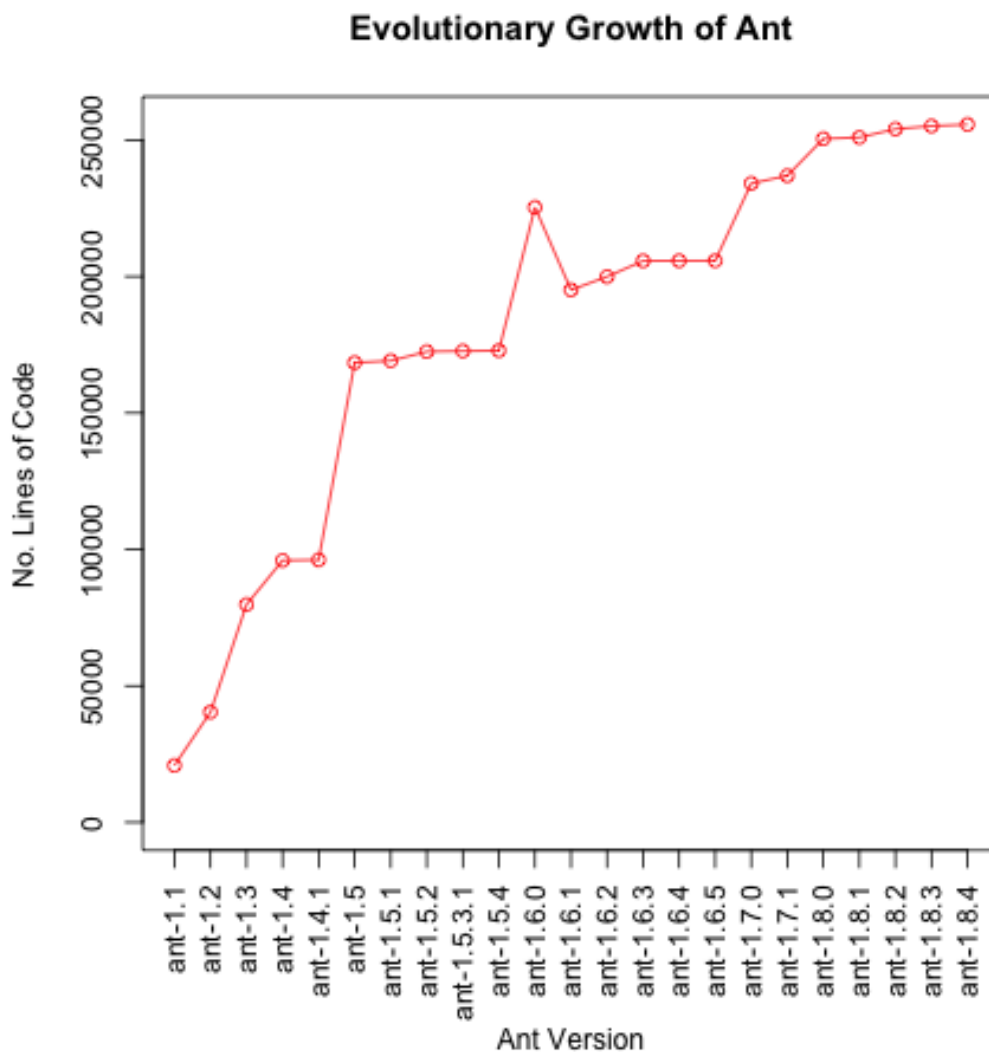


Figure 2.1: This graph shows how the size of the Ant software system grows twelve-fold in terms of lines-of-code between versions 1.1 and 1.8.4.

to the fact that the preconditions vary depending on the programming language being refactored, as do Mens et al. (2005). Furthermore, Mens et al. (2005) raise the point that it has never been demonstrated for mainstream programming languages that these preconditions are sufficient to guarantee that program behaviour is preserved for all refactorings. They claim that such a proof is impossible and we refrain from accepting this as a challenge here.

There have been attempts made with some degree of success to show that refactoring object-oriented *specification* languages (that have well defined semantics) does indeed preserve behaviour. For example, Carvalho Júnior et al. (2007) use CafeObj³ to prove the correctness of extract, inline, and move method refactorings as well as self-encapsulate field in a sequential subset of Java named Refinement Object-Oriented Language (ROOL). The purpose of (ROOL) is to reason about ‘object-oriented programs and specifications’. They encode the rules of the grammar

³<http://www.cafeobj.org>

of ROOL in CafeObj as well as the refactoring rules. CafeObj then ‘mechanises’ the refactoring proofs via rewriting.

Garrido and Meseguer (2006) take a similar approach using the Maude⁴ algebraic rewriting system by Clavel et al. (2003). They provide an equational semantics in Maude for three Java refactorings and formal proofs for two: pull–up and push–down field.

Mens et al. (2005) use typed labeled graphs to represent programs and show how program refactorings can be represented as graph transformations where it is possible to prove that refactorings exhibit properties such as access, update and call preservation.

Soares et al. (2010) take a less formal approach to demonstrating that refactoring preserves behaviour. They developed the SafeRefactor tool. It performs a static analysis to identify methods that have common signatures in both the source and target programs. It then generates test programs to test these common methods using the Randoop⁵ program and, if it finds any changes, it reports that a behavioural change has occurred.

Overbey et al. (2016) extend their earlier work (Overbey and Johnson, 2011) and use an approach to check behaviour preservation that they call ‘differential precondition checking’. They concede that their approach is neither sound nor complete but they highlight that the core algorithm used in their approach is language independent and could be abstracted into a library for use in various refactoring engines. Unlike Soares et al. (2010), their approach allows for behaviour preservation to be tested before the refactoring is performed. Their approach involves representing the program under refactoring as a program graph (an idea borrowed from Mens et al. (2005)), which they explain as an abstract syntax tree with extra edges to represent semantic information, for example variable access. The refactoring is then simulated and an updated program graph is produced. The semantic edges of the two graphs are then compared for *expected* differences. We do not describe this further here but we believe that identifying what these differences should be is one of the weaknesses to their approach.

From the literature, it appears that proving refactorings preserve behaviour is an unsolved research problem. We believe that while this remains the case, refactoring engines should at least *move* their conditions away from the core code of the engine and allow these conditions to be altered by users ‘in the field’.

Composite Operations

As the name suggests, composite refactorings are formed by composing primitive refactorings. For example, the extract class refactoring that features in the catalogue of refactorings by Fowler (1999) consists of a composition of create class, move field and move method primitives.

It is obvious that if each of the primitive refactorings are guaranteed to preserve program behaviour then composing any of these also preserves behaviour. However, Mens and Tourwe (2004) highlight work by Tokuda (2001) who describes certain compositions of refactorings as transactional. Transactional refactorings are those that when composed constitute a valid refactoring but if stopped at any point before completion then program behaviour would *not* be preserved and would therefore be invalid. They provide an example that they call ‘delegate method across object boundary’, which is the composition of a move method primitive and a create method accessor

⁴<http://maude.cs.illinois.edu>

⁵<http://randoop.github.io/randoop/>

- | | | |
|---------------------------------|--|--|
| • Create empty class. | • Change type. | • Replace statement list with function call. |
| • Create member variable. | • Change access control mode. | • Inline function call. |
| • Create member function. | • Add function argument. | • Change superclass. |
| • Delete unreferenced class. | • Delete function argument. | • Move member variable to superclass. |
| • Delete unreferenced variable. | • Reorder function arguments. | • Move member variable to subclass. |
| • Delete member function. | • Add function body. | • Abstract access to member variable. |
| • Change class name. | • Delete function body. | • Convert code segment to function. |
| • Change variable name. | • Convert instance variable to pointer. | • Move class. |
| • Change member function name. | • Convert variable references to function calls. | |

Figure 2.2: Primitive refactoring operations according to Opdyke (1992).

primitive (which was not listed by Opdyke (1992) and demonstrates that his list is incomplete). They provide a scenario that consists of a method a in class C , which is called by method b (also in class C). Method a is to be moved to class D . However, the move method primitive would not work in this situation because one of its preconditions requires that the method to be moved is not being referenced by any other methods. However, suppose we ignore this condition and carry out the refactoring along with the create method accessor refactoring. The program will be left in a state with method a located in class D and a new method a' in C , which is a delegate to the moved method a . This constitutes a valid refactoring when performed as a whole.

Choosing the correct refactoring to apply, whether it is a simple primitive or a composed sequence of primitives is a matter for the developer. Likewise, deciding when to refactor is also a decision to be made by the developer. However, there are some subjective notions to guide the developer along, such as code ‘smells’, which are discussed next. Software metrics also provide a more sensible approach. Their use is discussed in subsection 2.2.3 and more generally in section 2.5.

2.2.2 When to Refactor & Considerations for Developing a Refactoring Engine

Developers take two contrasting approaches for deciding when to refactor (Murphy-Hill and Black, 2008).

1. Developers who choose *root canal* refactoring restructure their source code after it has decayed.
2. On the other hand, developers who opt for *floss refactoring* improve their code structure in tandem with normal development activities.

Murphy-Hill et al. (2012) conclude that the latter is the favoured approach by developers.

Both approaches rely on indicators that suggest *when* and *what* to refactor. These indicators are referred to as ‘smells’ in the code. Fowler (1999) describes twenty-two such ‘smells’. These include ‘feature envy’ where the code in one class has a high number of dependencies on data in another class.

Tools known as ‘smell detectors’ have been developed to detect code ‘smells’, which can be used to determine when and what to refactor.

Simon et al. (2001) create a visual ‘smell detector’⁶ to help developers identify move method, move attribute and inline/extract class opportunities. They extended the Crocodile tool (a language independent metrics tool) to show graphs where entities that are close according to a distance metric are presented close in Euclidean space. Their tool does not actually perform the refactoring but the refactoring should be performed manually on elements that are distant from each other.

Fokaefs and Tsantalis (2011) developed a plug-in for the Eclipse IDE named JDeodorant⁷. It detects the ‘God’ class code ‘smell’ and suggests remedying refactorings. The authors have used it to identify opportunities for applying the extract class composite refactoring, which is used as a solution for eliminating the ‘God’ class ‘smell’ (where one class contains too much data or behaviour). They evaluate the tool in an empirical study involving version 5.3 of JHotDraw. They evaluated their approach using a ‘professional software quality assessor’. Although they do not discuss what qualifies her/him with such a coveted title, s/he supported the suggested refactorings by the tool and said that nine out of the sixteen tools suggested that refactoring improved the code with respect to its comprehensibility.

Murphy-Hill and Black (2008) highlight that the problem with ‘smell detectors’, such as JDeodorant, is that they are unsuitable for developers who choose to floss refactor. They argue that because floss refactoring is an activity that is repeatedly done, tools that require explicit interaction are less useful. They suggest seven habits of good ‘smell detectors’. The first of these is availability: the tool should report ‘smells’ as early as possible and with as little involvement as possible from the developer. They also support unobtrusiveness: the tool should work without blocking the developer. They have implemented the tool ‘Ambient View’ that takes into consideration these and the remaining five suggested good habits.

We next discuss automated approaches to refactoring beyond ‘smell detection’.

2.2.3 Manual, Automated and Search-based Refactoring

Investigations by Murphy-Hill et al. (2012) suggest that refactoring tools are under utilised. From their study, they found that 89% of 145 refactorings (supported by tools) were performed manually. The developers in this study were ‘toolsmiths’ who produce the refactoring tools so they are unquestionably aware of the existence of appropriate aides. Manual refactoring can be carried out by carefully following the steps laid out in refactoring catalogues such as Fowler (1999). However, this approach is not recommended by us as manual refactoring has the potential to introduce defects due to human error. For example, suppose in some method there exists a local variable that shadows a global variable. If the developer renames that local variable in all but one place then the program will compile and execute but will use the incorrect variable at some point. We believe that refactoring should at the very least be done using refactoring tools such as the

⁶A misnomer perhaps but, nonetheless, an apt description.

⁷<https://marketplace.eclipse.org/content/jdeodorant>

refactoring tool provided in the Eclipse IDE. Although these tools should be more reliable than performing the refactoring steps manually, there are other automated approaches that can further increase efficiency, for example, search based refactoring.

O’Keeffe and Ó Cinnéide (2006, 2008a,b) have developed a search based approach to refactoring and implemented the CODEIMP automated design improvement tool. It is inspired by Harman and Clark (2004) who consider software metrics as fitness functions. It takes a Java 1.4 program and performs a search of the space of applicable refactorings and refactors the Abstract Syntax Tree according to the results. Moghadam and Ó Cinnéide (2011) later improved on this work by updating CODEIMP to be applicable to Java 6 source code.

We believe this full level of automation might be a step too far for many developers. We settle for a compromise that would allow for refactorings to be scripted. With this in mind, and also our suggestion that conditions should be extracted from the core refactoring code, we take a look at DBC. This relates to RQ2.

2.3 Formal Specification & Design by Contract

Safety critical systems require a level of rigour beyond the capabilities of traditional software testing techniques such as those described by Myers et al. (2011). For this reason, rigorous approaches to software development were introduced including Design by Contract.

In the formal specification approach to developing software, the system’s critical components are modelled using specification languages such as the algebraic rewriting language Maude (Clavel et al., 2003). Properties of the system can then be reasoned about rigorously. For example, Listing 2.1 shows a specification written in Maude for a pelican–crossing⁸. The proofs at the end of the script demonstrate the safety of the system in the sense that it will not signal pedestrians to cross at the same time cars are instructed to proceed. It also demonstrates that the specification provides both cars and vehicles an opportunity to go. However, fairness is *not* reasoned about. Whilst Maude is described as an executable specification language, it is unlikely that it would be used to implement real world software. This is particularly true in this example, since pelican crossings would require an implementation language more suited for embedded systems, such as ‘C’. This creates a chasm between the specification and the implementation.

To address this issue, Meyer (1997) introduced Design by Contract (DBC). It allows for parts of the program code to be annotated with statements that express the intent of the code. Static analysis tools reason about the code with respect to the specification and can determine (in many cases) if the implementation implies the specification. The specifications in DBC are commonly referred to as contracts because they describe the obligations of both the client using the code and of the developer who provides the code.

```

1 mod PELICAN-CROSSING is
2   protecting BOOL .
3
4   sort Globals .
5
6   vars C P : Bool .

```

⁸This is a translation of the specification written in the Event–B language and published in the Rodin handbook (Jastram, 2014).

```

7
8  op invariants : Bool Bool → Globals [ctor] .
9
10 cr1 [set_peds_go] : invariants(C, P) => invariants(C, true)
11 if not C and not (P and C) .
12
13 cr1 [set_peds_stop] : invariants(C, P) => invariants(C, false)
14 if not (P and C) .
15
16 cr1 [set_cars_go] : invariants(C, P) => invariants(true, P)
17 if not P and not (P and C) .
18
19 cr1 [set_cars_stop] : invariants(C, P) => invariants(false, P)
20 if not (P and C) .
21 endm
22
23 *** If the following search returns no results then we are given some assurance
24 *** that the system is safe. It is not an absolute guarantee, the search might
25 *** have timed out before finding a result for example.
26 search in PELICAN-CROSSING
27   : invariants(false, false) =>* invariants(true, true) .
28
29 *** If the following search returns non--empty results then we are guaranteed
30 *** that the specification allows for cars to have some opportunity to go.
31 search in PELICAN-CROSSING
32   : invariants(false, false) =>* invariants(true, false) .
33
34 *** If the following search returns non--empty results then pedestrians are
35 *** guaranteed to be given an opportunity to go.
36 search in PELICAN-CROSSING
37   : invariants(false, false) =>* invariants(false, true) .

```

Listing 2.1: A Maude specification for a pelican crossing system.

Contracts for specifying behaviour in object-oriented programming languages such as Java (using JML by Leavens et al. (1999)) involve formalising:

- Preconditions — These constrain the values that may be passed as arguments to the method. If the constraints are not held then the client cannot be guaranteed that the method will behave in the intended manner. Preconditions can be computed by working backward through the program using Hoare logic (Hoare, 1969).
- Postconditions — These make guarantees about the result of methods. Postconditions are only valid when the preconditions are met and are implied by the body of code in the method.
- Invariants — These are placed on classes or within loops. Invariants on loops exist only to facilitate the theorem prover in proving that the loop behaves correctly. Class invariants describe constraints about the state of the object at any time during its lifetime. Invariants are not considered in the remainder of this dissertation.

We believe that DBC goes some way to address RQ2. However, we still require an approach for representing programs being refactored. We believe that abstract syntax trees are too verbose

for this purpose and we identify MDE as a solution. We are further motivated by MDE because it is also suitable for building a language to specify and script refactorings. We discuss MDE next.

2.4 Model Driven Engineering

The influence of modelling can be seen in many aspects of software engineering. At the high level, management use (descriptive) models called processes to describe how teams of engineers should structure their work activities. At a lower level, software engineers create models to ‘represent something by something else’ (Muller et al., 2012) such as classes to represent real world concepts. The benefits are apparent. Models allow concepts to be communicated clearly between developers or to serve as specifications for parts of a system to be implemented (we call these prescriptive models). Model Driven Engineering (MDE) builds on the advantages of using models by making the design and implementation of models the core activity of software engineering. MDE practitioners hold the view that ‘everything is a model’ (Bézivin, 2005). But what is a model? We accept the view expressed by Stachowiak (1973) (in German but translated in Muller et al. (2012)) that a model is *some thing* based on an *original* for a specific purpose that captures only the important properties that are relevant to that purpose.

It is often desirable to transform one model to another. This is done through a model-to-model transformation. For example, suppose we have a UML class diagram, the ability to transform this into a relational database schema, for example, or an XML schema is desirable. We might also want to transform the same UML class diagram into ‘skeleton’ Java code. This can be done with model-to-text transformations.

The success of modelling hinges on being able to describe models effectively. For this, the Meta Object Facility (MOF) is used. We briefly describe the architecture of MOF next.

2.4.1 Metamodel Hierarchy and Transformations

In order to describe a model we require an appropriate language. The four layer metamodel hierarchy exists to facilitate this. At the highest layer (M3) is the meta metamodel. It describes the kind of things that can exist in the M2 layer. We can think of the M2 layer as the actual language, for example, UML is located in this layer. In the M1 layer, we use the language of M2 to describe things, for example we might describe some class diagram. The M0 layer contains instances of what is described in M1. For example, it could be an object in the computer’s memory that conforms to a class diagram described in M1. An object at M0 is *not* a Java class. A Java class also exists at M1, and is expressed by the Java language at M2 which is an instance of the grammar for Java at the M3 layer.

There is an interesting advantage to this hierarchy in that the M3 layer describes the minimal requirements for any language to be able to express the lower layers. For example, consider the example shown in Figure 2.3 (reproduced from (Nastov, 2013)). In this example, if the implementation language being used is powerful enough to express the concept of class at the M3 layer then the language is sufficiently powerful to describe the notions in the lower layers. This is true due to the ‘instance of’ relationship that crosses the boundaries of the layers.

However, having languages to describe models is pointless without concrete tools. MDE environments have been developed for this reason. We discuss these next.

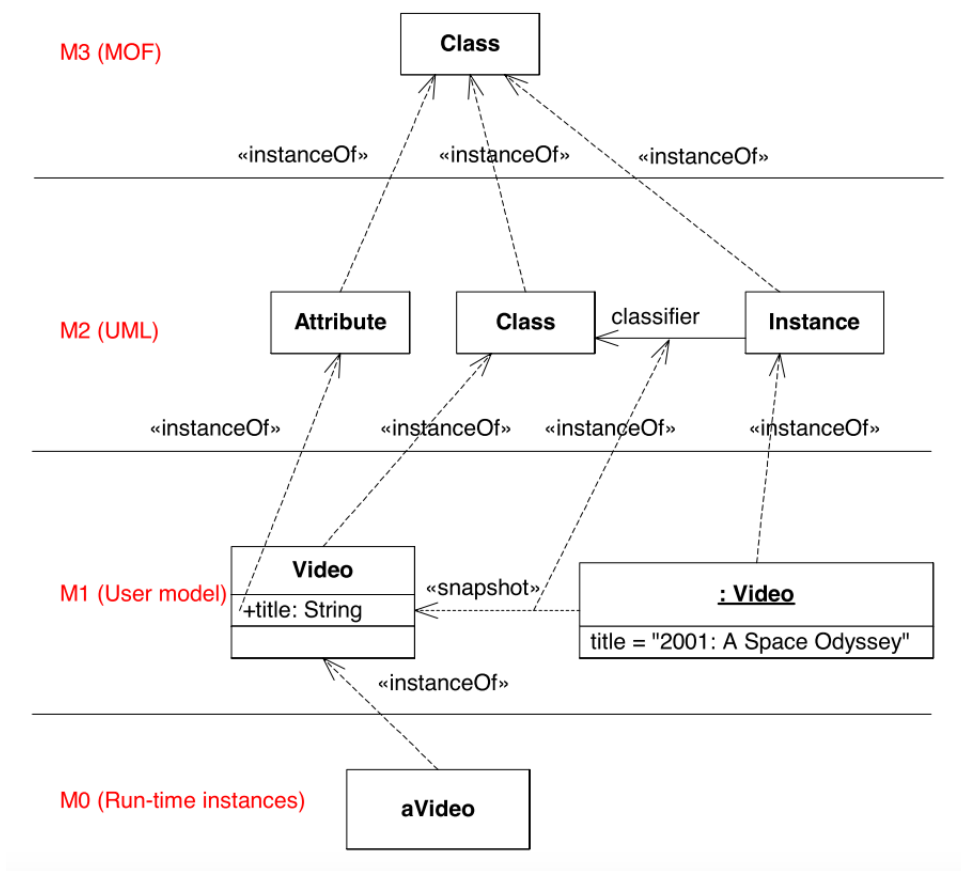


Figure 2.3: An example showing the four layers of the metamodel hierarchy.

2.4.2 MDE Environments

Frameworks have been developed specifically to accommodate MDE. The Eclipse Modelling Framework (EMF) (Steinberg et al., 2009) has been developed by the Eclipse foundation. It allows developers to build domain models either in Java, as XML schema definitions or as UML class diagrams. The models are mapped to a core modelling format called Ecore. This centralised model can be mapped to other model formats such as a generator model. The generator model can be used to produce skeleton Java code, code for testing and a plugin for Eclipse which provides an editor that has syntax highlighting and content assist.

Simulink by MathWorks⁹ is an alternative environment for working with models. It differs from EMF with respect to its target audience. Simulink is suitable for engineers developing embedded systems. It provides a block diagram environment that allows code generation and simulation MathWorks. We do not mention more about it here as we are concerned with frameworks for *software* engineers.

MDE has been used to implement languages for modelling specific domains. We call these DSLs and discuss these next.

2.4.3 Domain Specific (Transformation) Languages

A domain specific language DSL is a language developed for a specific domain or task. The greatest difference between a DSL and a general purpose language is the vocabulary used. The keywords in a DSL will be words from the domain for which the DSL is designed. In a programming language, keywords will refer to programming constructs. Familiar examples of DSLs include the \LaTeX language used for document preparation. Although it is Turing complete, its intended purpose is to describe the content and structure of articles, letters, presentation slides and other publishings. The \LaTeX package for \LaTeX (to produce diagrams such as Figure 5.2) is also a DSL so, hence, we can have one DSL embedded inside another.

DSLs are classified as being either:

1. Internal — An internal DSL is one that is developed within an already existing programming language. This is much the same as what traditionally have been called *libraries* but with a very focussed purpose. For example, Thompson (2011) presents a simple DSL for geometric manipulations of ASCII ‘art’ embedded in the Haskell¹⁰ language. Internal DSLs are simpler to implement. For example, there is no need to implement a parser or interpreter for the language. It also has the benefit that users of DSL are not required to learn a new programming language syntax and can continue to use their preferred and familiar tool-chains (IDES, text editors, etc.). We could also consider XML schema definitions, or XSDs, as crude internal DSLs since they describe sets of specific XML documents in XML.
2. External — External DSLs on the other hand are standalone languages, which come with their own advantages. Users of these languages do not need to have experience working with general programming languages. They only need an understanding of the domain in which they are working. This is beneficial because they are isolated from what could be obscure error messages while using general purpose languages. The implementation effort

⁹<http://uk.mathworks.com/products/simulink/>

¹⁰<https://www.haskell.org>

however increases. The language designer must develop their own parser and supporting tools such as an interpreter. However, language frameworks such as Xtext¹¹ and MPS¹² have been developed that relieve most of these burdens from the language designer.

Specific DSLs have been developed to transform models. For example, the Object Management Group (OMG) (OMG, 2015) specify:

1. Query/View/Transformation Operational — QVTO is an imperative language *specification* for transforming models. One implementation is ATL (Jouault et al., 2008). Although it has imperative features it also can be used declaratively, so, it also adheres to the QVTR specification.
2. Query/View/Transformation Relational — QVTR is a declarative language for performing model transformations. It relies on pattern matching to perform the transformations. The user describes a pattern that is matched on the input and this provides a rule for what the output should be.

Both QVTO and QVTR offer a general approach for describing model transformations. However, in some circumstances one might want to constrain the transformations that can be applied (as is the case with refactoring). We coin the term Domain Specific Transformation Language (DSTL) to refer to languages that exist to apply transformations under constraints.

2.4.4 Modelling Java Programs

The MoDisco¹³ project has been created to help developers to modernise legacy software systems. MoDisco comes with a Java model in Ecore format for representing Java systems as well as knowledge discoverers that can be used to transform existing Eclipse Java projects into MoDisco models. It also comes with supporting tools such as a model browser for exploring the models visually in the Eclipse IDE and transformations to produce UML diagrams.

In this section, we have seen that models provide a practical approach to represent software source code and build languages (DSLs). It has also been employed in the area of software metrics, which relates to RQ3. We discuss this next.

2.5 Software Metrics

All branches of engineering rely on measurement to provide meaningful feedback. This can be used to assess quality or to make predictions (Meneely et al., 2012). Software engineering is no exception. For example, a software engineering team developing a system might test for the mean time between crashes. This can be used to evaluate the suitability of their product for specific scenarios.

At a lower level, the engineering team might evaluate the quality of their design. This can be used as an indicator to decide how easily the system can be maintained in the future. These measurements might even form part of the contract between the engineering team and the client.

¹¹<https://eclipse.org/Xtext/>

¹²<https://www.jetbrains.com/mps/>

¹³<https://eclipse.org/MoDisco/>

Measurements could even be used as performance indicators that help management decide how to attribute specific portions of the system's implementation to specific engineers. For example, management might identify certain engineers as being better at implementing algorithms rather than object oriented design and vice versa. This trade off can be used to decide to whom tasks should be assigned.

Software metrics come in two varieties: project metrics and design metrics. We discuss both next before remarking on their validity and available suites developed for their use.

2.5.1 Project Metrics

Project metrics are high level metrics used to make management decisions. Examples include those published by Lorenz and Kidd (1994):

- Average person days per class — Suppose management is estimating the cost of staffing a new project and the high level design has already been completed so that management has an estimate of the number of classes that will be needed by the new system. Management can use data from previous projects to estimate the amount of time it will take to complete the system.
- Number of scenario scripts — A scenario script is a use case that documents actions that are performed, who they are performed by and who they affect. The number of use cases will indicate the size of the application and consequently the length of time required to implement it. Historical data is also required in this case.

Although we could devise project metrics to predict the time it would take developers to improve a system's structure, it is beyond the scope of this work and we do not address the issue further.

2.5.2 Design Metrics

Design metrics are used to evaluate the quality of an implementation. Three aspects of the system that are commonly measured include:

- Cohesion — Many software systems in existence are built using object-oriented languages that require a system to be split into modules called classes. A well designed object oriented system consists of classes that are semantically cohesive, i.e. each class represents just one meaningful concept. For example, a cohesive implementation of a 'person' class might have the person's name, age, address and phone number, but when loosely related details are added, such as bank account number and balance, then that class is no longer a cohesive unit. Many metrics related to cohesion consider the relationships between the methods and the attributes. A highly cohesive class should have methods that access a high number of the attributes. The exception to this rule of course are accessor and mutator methods. The literature contains a plethora of cohesion metrics that including the common LCOM metrics (Chidamber and Kemerer, 1991, 1994; Hitz and Montazeri, 1995) as well as the TCC and LCC metrics of Bieman and Kang (1995).
- Coupling — Coupling comes in two 'flavours'. The afferent coupling of one class is a measure of how dependent other classes are on it. High afferent coupling indicates that a

class' interface is too public. On the other hand, and heavily related, is efferent coupling. This is a measure of how tied one class is to another. If a class C has high efferent coupling then it is considered to be poorly designed because it means that if the class or classes it relies on change then those changes must be reflected in C . Various coupling measurements have been proposed for this purpose including the CBO metric published by Chidamber and Kemerer (1991).

- Complexity — Most software engineers would agree that a program consisting of just sequential lines of code without iterations or conditions is less likely to contain defects than one with a high number of if-statements and loops. The rationale of course is that the developer might have forgotten to include a condition for some special case or written an incorrect guard on a loop. Programs with branches are said to be more *complex* than sequential programs. Various metrics have been developed to quantify how complex a program is. For example, McCabe (1976) published the cyclomatic complexity metric.

2.5.3 Validation of Metrics

Meneely et al. (2012) acknowledge that software engineers and researchers do not have a formal set of rules for what constitutes metric validation. They conducted a systematic literature review and produced a list of forty-seven criteria identified from twenty papers that matched their criteria. We do not repeat the criteria here but these range from *a priori validation* that requires the author of the metric to hypothesise first what attribute being measured is statistically significant in explaining some phenomenon. They also require that in order for the metric to be valid it must be meaningful, in the sense that it allows for a manager or software engineer to make some decision based on the value of the metric. They refer to this as actionability. A more obvious validation criterion includes 'definition validity' which requires that the metric is stated correctly so that it can be implemented. The original mathematical definition of the LCOM1 metric failed this criteria as it always evaluated to zero.

2.5.4 Metric Tools

The MOOSE metrics suite¹⁴ is a metrics tool to help developers evaluate and browse their code. Evaluation is supported through polymetric views including those described by Lanza et al. (2005). MOOSE is not tied to a specific programming language. Programs are represented internally using the FAMIX metamodel. Any language can be supported by providing a transformation from source code in the language to FAMIX.

McQuillan (2011) also takes a model driven approach to calculate software metrics. She developed a metamodel that describes the measurable components of a software system. This is utilised by another metamodel, which defines the actual operations for calculating metric measurements. These are annotated with OCL expressions¹⁵ that provide the operations with semantics.

We highlight the disparity between metric tools and note that the same metric can yield different results in different tools. This is due to tool developers interpreting the metric definitions differently. This disparity has been confirmed in a study by Lincke et al. (2008).

¹⁴<http://www.moosetechnology.org>

¹⁵<http://www.omg.org/spec/OCL/>

Other than evaluating code quality, attempts have been made to use software metrics to identify where refactorings have occurred in subsequent versions of software. We discuss this interesting usage briefly next.

2.5.5 Metrics to Identify Refactorings

Refactoring detection is an area of research that involves identifying where refactorings have occurred in subsequent versions of software. The motivation for this research is to develop tools that can show developers working on large open scale projects, for example, where their changes have moved after code restructuring. A better approach would be for the refactoring engine to record and replay the changes in front of the developer, similar to what Henkel and Diwan (2005) achieve. However, their approach assumes that developers refactor with compatible refactoring engines. We know this is not the case and point out that some developers still prefer to use text editors such as Emacs¹⁶ rather than full IDEs with refactoring support.

Schneider et al. (2010) have evaluated the potential of using software metrics to identify refactorings. Their hypothesis is that, between two classes, a small change in the Euclidean distance of fifty-four different metrics is likely to be the result of refactorings and larger changes are likely to be the result of feature additions. They evaluated their approach using three versions of the Struts web-framework. They concluded that metrics were not a promising approach for determining if a change to code was a refactoring or not.

2.5.6 Summary

The purpose of this chapter has been to examine the literature to investigate the state-of-the-art and to address the three research questions listed in section 2.1.

In response to RQ1, the literature states that in general it is not possible to prove that refactoring preserves observable behaviour for all refactorings (Mens et al., 2005). However, we have identified some successful attempts to do so for formal specification languages with well defined semantics. Less ‘water tight’ approaches have also been taken such as Soares et al. (2010) who resorts to random testing. We believe that a better approach (for now) would be to allow engineers using refactoring engines to alter preconditions or postconditions for refactorings easily in the field. Thus any conditions they identify as being over or under specified can be corrected immediately.

With regard to RQ2, very little effort has been made to allow developers to define their own reusable refactorings. The more well known refactoring engines such as that provided by the Eclipse IDE allow for refactorings to be composed but only manually. Larger refactorings cannot be saved for later application. Furthermore, software source code appears to get a lot of attention in refactoring circles but very little focus is placed on refactoring other artefacts such as UML diagrams. We believe that a scriptable refactoring engine that considers alternative artefacts is required to appease contemporary software engineering efforts.

In relation to RQ3 there is also very little consideration paid to overall quality while refactoring. The exception being the work of O’Keeffe and Ó Cinnéide (2006, 2008a,b) whose CODEIMP tool searches for the optimum refactoring to apply.

Clearly, there is scope for improvement with current generation refactoring tools. We believe

¹⁶<https://www.gnu.org/software/emacs/>

by taking an MDE approach to refactoring we can develop a refactoring engine that can refactor heterogeneous artefacts while paying consideration to quality as well as making it scriptable. Moreover, we believe that we can even implement such a system using MDE principles. We begin to do so in the next chapter as we develop a language to define and script refactorings.

Chapter 3

RefDstl: A Language for Specifying and Scripting Refactorings

In this chapter, we take our first stride towards developing the RefDstl system: we develop the RefDstl language for specifying and scripting refactorings. We begin with a brief discussion regarding use cases for the language and who we anticipate will use it. We discuss these issues in order to steer the language's design, which we subsequently discuss. This includes choosing a language paradigm and developing the abstract and concrete syntax. The semantics of the language is not discussed, although, it should be self explanatory. The implementation of an interpreter is described in Chapter 4 that effectively provides an operational semantics. As the chapter comes to an end, we introduce our standard library of primitive refactoring operations, written in the RefDstl language, and we also provide a small example on using the language to refactor classes from a fictitious video game.

3.1 Use Cases

We develop a language that can be used to define new refactorings through the composition of a set of primitives provided in a standard library. The language also allows for the application of the refactorings to be scripted using embedded commands that instruct the interpreter regarding what system is to be refactored.

The primary arena where we envisage this language being used is in software production environments. For example, in industry different organisations have their own coding guidelines such as that all class attributes must be self-encapsulated (i.e. access to attributes should be through the accessor and mutator methods even in the body of the class where the attributes are defined). Different developers usually have their own style of programming that can conflict with the organisation's guidelines. The developer could manually check that their code meets the guidelines but this is tedious and not a productive use of time. Instead, the organisation could write a script in RefDstl with refactorings that corrects non-conforming code (such as applying the self encapsulate field refactoring). While many refactoring engines can already perform a self encapsulate field refactoring they require the user to work through a graphical interface. Our scripting language allows for a higher degree of automation. We concede that our language is not sufficiently powerful to describe all coding guidelines. In particular, lower level guidelines that describe the structure of the text are not possible with RefDstl. For example, it would not be possible to specify that a space must occur between each operator and its operands or that chain brackets are 'lined up'.

We point the reader interested in such low level changes to the work on ‘delta oriented programming’ by Schaefer et al. (2010). We believe their approach could be easily extended with regular expressions that would allow these lower level changes to be performed.

We believe this language also has potential for use in academic circles. For example, consider the research area of refactoring detection. Dig et al. (2006) and others identify refactorings that have occurred between subsequent versions of software. Their motivation is to allow developers to see where their code contributions go after a program undergoes refactoring. However, the associated tools that each of these researchers develop implement different metamodels for internally describing refactorings. We view this situation as unsatisfactory. It is difficult for the outputs of the different tools to be automatically compared. It also prevents users from harnessing the power of multiple tools. They should be able to merge results to provide more accurate or complete results. The language that we present here offers a solution. Although the language is text based, which might be cumbersome to have a machine produce, it has an underlying Ecore metamodel. We suggest that researchers could programmatically instantiate that metamodel and a general model-to-text transformation be provided to transform the model into a textual RefDstl script.

Of course, nobody will use our language if it is not ‘good’. We discuss the properties of ‘good’ programming languages next.

3.2 Characteristics of a Good Programming Language

We have identified the following properties that a programming language should have in order for it to gain acceptance. We bear these in mind in the subsequent sections as we discuss our implementation of the RefDstl language.

- **Easy to learn** — When faced with a new programming language, a programmer evaluates if the gains from learning how to utilise the language effectively outweigh the initial learning cost. A language that is easier to learn has a higher chance of being adopted when in competition with more difficult alternatives. Clearly, we should aim to make RefDstl simple to increase its chances of adoption. RefDstl has a higher opportunity of adoption if it follows a simple design. To obtain insight into what makes one programming language easier to learn over others, we looked toward the area of computer science education. A study by Milne and Rowe (2002) involved a questionnaire of second year students in the process of learning C++ at the University of Dundee as well as academic staff teaching programming languages from around the United Kingdom. They concluded that the programming constructs that students find most difficult to understand relate to what is happening in the machine’s memory during execution. Although students will probably learn to adapt by the end of their studies, we bear their conclusions in mind and suggest that the source of the problem is students having to grasp the concept of mutable state of program variables unlike variables more familiar to them from mathematics.
- **Ease of use** — A programming language should have a clear and consistent semantics that allows the developer to produce programs or scripts with no side effects. This is in line with guidelines developed by Karsai et al. (2009), who advocate that (domain specific)

languages should be consistent and have a defined purpose; if a feature does not contribute to that purpose then it should be removed.

- Power of expression — Languages should be powerful in expression to allow programmers to succinctly describe their objective. RefDsl must at the very least be capable of describing primitive and composite refactorings and also allow describing where the refactorings are to be applied. Ideally, the language should allow for the developer to easily specify multiple locations where refactoring is to be applied. This is in contrast to graphical approaches to refactoring where the developer has to select each location for application individually.
- Tool support — Tool support for a programming language is as important as the language itself. Good tools aide developers to complete tasks quicker, which can lower development costs. For a general purpose programming language, the bare minimum that should be provided is a compiler or interpreter. This is true of some domain specific languages as well, including RefDsl. Some DSLs require tools to transform programs in the DSL to general purpose languages. For example, SQL queries can be turned into ‘C’ calls in order to evaluate them. In addition, tools that help to resolve programming errors by providing feedback to the programmer or help to make the code more comprehensible are also desirable. Debuggers, for example, allow the programmer to step through the program to identify points in the program where unexpected behaviour occurs. The development of a debugger for RefDsl is beyond the scope of this work. However, we do provide an editor that offers syntax highlighting and content assist in the Eclipse IDE as well as an execution environment.

3.3 Design of a DSTL for Defining and Scripting Refactorings

The following cascading stages form a framework for designing and implementing new programming languages. We address each of these in turn while discussing the creation of the RefDsl language in the subsections that follow.

1. Choosing a suitable paradigm for the language — The paradigm of a language defines its high level design principles. For example, does the programming language place emphasis on objects or actions? The chosen paradigm influences how the language is expressed in text, i.e. its syntax. For example, if the language permits objects then the language must allow a way of defining/instantiating objects.
2. Defining a syntax of the language — The syntax defines the structure of the language. For example, what are the keywords of the language and what order must they be composed to form valid ‘sentences’. The syntax of a language is expressed as a grammar that consists of production rules. The grammar file is then input to a parser generator, which produces a parser. The parser is essentially a language recogniser: It accepts only those inputs that conform to the grammar specification. The syntax chosen is naturally dependent on the constructs available in the language.
3. Specifying a semantics of the language — The elements of the language’s syntax must have an explicit meaning. For example, a common question arises in language design

when invoking methods/procedures/functions in the language, that is, how are the arguments passed? They could be passed by value or by reference, for example. The language might even offer syntax that allows the programmer to switch between the two.

4. Implementing an execution environment — Once the semantics of the language have been defined then software for interpreting the language must be developed. Alternatives include producing compilers, which transform the source code program into an executable that is usually a binary executed directly on the machine or an intermediate file that is executed on a virtual machine, which is the case with Java or the .NET family of languages.

3.3.1 Choosing a Suitable Paradigm

There are various programming paradigms in existence such as imperative/declarative or object-oriented/procedural. Programming languages can even fall into multiple paradigms.

- Imperative — These programming languages usually include low level constructs that allow the programmer to explicitly control the behaviour of the underlying hardware. For example, they might have the concepts of variables and assignments to manipulate the state of the machine, branching statements that allow the machine to move from one state to another when certain conditions are met, or iterative constructs that move the machine through a sequence of states repeatedly until other conditions are met. The languages that find most use today are imperative in nature, such as ‘C’, C++, and Java.
- Declarative — These are higher level languages and programmers specify *what* they want done rather than *how* they want it done. SQL is, perhaps, the most widely used declarative programming language that is used for data definition and manipulation in relational database systems. Declarative languages can be quicker to program with and often require fewer lines of code. This is especially true when designed for domain specific purposes, such as the case with SQL. Consider how an SQL query that joins two tables would need to be written in ‘C’ with a loop that populates some data structure with the Cartesian product of the two tables followed by a loop that iterates through the structure to remove rows where the appropriate columns do not match under conditions. General purpose declarative programming languages also exist such as Haskell. The productivity gains that come with using declarative languages often come at the cost of expressibility. For example, often developers will find it easier to describe an algorithm imperatively using variables (rather than bindings that appear in declarative languages such as Haskell), loops and if statements.
- Object–Oriented — In the Object–Oriented (OO) paradigm, the focus of development is on writing reusable modules of code called classes. A class is a container for data which is stored in variables called fields (or attributes) and methods which operate on that data. A good class is highly cohesive, i.e. each method in the class reads from or writes to a high proportion of the class’ fields. It is also loosely coupled: it depends very little on other classes and allows other classes to depend very little on it.
- Procedural — Procedural programming places the focus of development on writing procedures of code, which are grouped lists of instructions. Procedural languages heavily contrast

with OO languages in that the procedural languages place emphasis on actions to be done rather than on the data.

In choosing a paradigm for RefDsl, we decided to take a balanced approach by providing a language that allows the programmer to declaratively describe what refactorings to apply but enriching it with a ‘for’ looping construct to permit repetition in the language and a conditional ‘if’ construct. It would have been possible for us to provide a ‘map’ and ‘filter’ function, which are suitable alternatives and more the ‘norm’ in declarative languages but our choice is driven by what we believe is most familiar to the target end users. It would be possible to implement different variants of RefDsl following different paradigms so that a user could choose the paradigm most familiar to them. Our reliance on MDE technology (discussed in section 3.4) to implement RefDsl means that the same interpreter could have been used to interpret scripts written in all of the variants by performing model-to-model transformations to some ‘baseline’ paradigm.

3.3.2 Language Constructs

We discuss the constructs provided by RefDsl in this section, which allows us to develop an appropriate syntax later in section 3.4.

- Modularisation via libraries of procedures — For the obvious benefit of code re-use, we implemented the concept of ‘libraries’ into the language. Each script is itself a module that is uniquely identifiable based on its location in the file system. This path is the script’s identifier that can be used to import it. Unlike other programming languages that include visibility control features (such as private methods in Java) that permit for elements of a module to be marked as hidden to external modules, all elements of a script are visible to any other script importing it. A script is imported by using the ‘imports’ keyword at the beginning of the script. Multiple scripts can be imported by writing imports statements one line after the other.
- Modularisation of scripts via procedures — We believe in the guidelines of structured program that advocate structuring code into blocks called procedures that perform one well defined task. RefDsl supports the notion of procedures for this purpose. Procedures are exactly that; they are not functions. A procedure never returns a value.
- Preconditions and Postconditions — Following the practices of design by contract, RefDsl allows for procedures to be annotated with pre- and post- conditions, which are expressions that can reason about the program being refactored in its representation as a MoDisco model. Plugins can be developed for RefDsl that provide hooks (which we call script handles), which further enhances what can be reasoned about in these conditions. We have incorporated a plugin with a hook that allows the conditions to reason about the program represented as an instance of the metrics metamodel developed by McQuillan (2011).

These conditions serve a dual purpose. The first is that they ensure the refactoring primitive operations work correctly. It might seem unusual to ensure the correctness of primitive operations in a script rather than in the transformation engine but this provides a benefit. As agreed by Overbey and Johnson (2011), the formulation of refactoring conditions is actually the most difficult part of developing a refactoring engine.

By separating the conditions from the code, it makes it convenient to debug overly strong preconditions (as they can be commented out). A standard library of refactorings is provided with the RefDsl system so the end user will be unaware of this. The second purpose is that it allows RefDsl end users to define their own refactorings with their own conditions, that can be used to verify the correctness of the refactorings with respect to properties that they specify.

Unlike some implementations of DBC, RefDsl does not provide syntax that would allow for the total correctness of the program to be verified. It only permits checking for ‘partial correctness’. This means that we are limited to verifying *if* the program completes then the conditions are guaranteed to have held. We make no guarantee that the script will finish executing. To add total correctness to the language, we would need to include loop variants¹. Pre and postconditions are simple boolean formulae built using the familiar boolean connectives as well as existential or universal quantifiers.

- Procedure Invocation — Within the body of a procedure, the user must have a way to invoke other procedures and pass parameters. We provide the ‘call’ keyword for this purpose.
- Branching — The language provides branching constructs to the user. We limit this to ‘if’ statements without ‘else’ blocks. ‘Else’ blocks could be easily added at a later point but during the development of the standard library of refactorings (in Appendix B) their inclusion was found to be unnecessary. The language currently includes the necessary boolean operations to work around this if it poses difficulties to the user.

For the sake of maintaining consistency in the language, we permit the same format in ‘if’ conditions for describing guards as is used in the formulae in pre and post conditions. This means that guards in RefDsl can be more expressive than guards in other languages, such as Java 7, which do not permit quantification in conditional guards.

- Iteration — The language also provides an iteration construct for looping over collections of primitive refactoring instructions.
- Refactoring primitives — These statements manipulate the program under refactoring, such as renaming fields, methods etc.

3.4 RefDsl Implementation

Up until now when discussing the language’s syntax we have been referring to the ‘concrete’ syntax, i.e. what the programmer types as a program. This is too high level and cumbersome to translate by compilers or understand by interpreters. So an alternative, simpler view of the syntax is used. This is known as the abstract syntax, which Völter et al. (2013) refers to as a ‘data structure that holds the core information of the program’.² The abstract syntax forgets the insignificant parts of the program. For example, it drops characters that are used as delimiters such as whitespace (to

¹In contrast to loop invariants that describe a property of the program’s state that is true upon each iteration of the loop, a loop variant is a value that must decrease upon each iteration.

²We believe the data structure is a representation of the abstract syntax and not the abstract syntax itself but that distinction bears no consequence here.

separate terms), semi-colons (often used to separate statements) or brackets (to create precedence in complex expressions). These can be omitted without losing the meaning of the program because they only serve as hints to the lexer (which separates the terms in the program text) and the parser as to what parts of the program have a higher level of precedence. The output from a parser is a parse tree with typed nodes. Interpreting or compiling the program involves traversing the tree to generate machine code instructions or interpreting the nodes of the tree and performing some action on a virtual machine (as is done by an interpreter).

In this section, we discuss our use of the Xtext language framework for developing RefDsl.

3.4.1 Parser Generators and Grammars

A parser generator is a tool that takes as input a specification of the concrete syntax and produces as output the code for a parser. When compiled this can be used to parse programs that fully meet the language specification. Moreover, if the parser is provided with input that does not conform to the specification of the language then it will halt before it finishes parsing the program. ‘User friendly’ parsers also provide feedback to help the programmer to locate the problem with their program such as line number or character the parser realised a problem had occurred.

With traditional parser generators, like JavaCC³, it is the language designer’s responsibility to write the classes that represent the abstract syntax. He or she must also embed code into the concrete language specification that builds up the abstract syntax tree as the parser reads through the program text. This is done by placing code snippets inside the grammar specification.

More recent language frameworks have been developed that make it easier to develop languages (specifically DSLs). These make it easier and less time consuming to develop languages because they automatically derive the abstract syntax based on the grammar of the language.

3.4.2 A Grammar for RefDsl

Grammars specify the syntax of programming languages. They are regularly provided in Extended Backus–Naur form (EBNF). In Backus Naur Form (BNF), a grammar consists of these components (Aho et al., 2006):

- Terminal symbols — These are keywords and other symbols in the language such as symbols for built in operators, brackets and braces for describing lists or grouping statement or symbols to denote separation of statements such as semi-colons.
- Non-terminals — A non-terminal is a variable in the grammar that is associated to a production rule. There must be exactly one non-terminal to mark the beginning of the grammar.
- Production rules — A production rule is a mapping from a non-terminal to a sequence of terminals or non-terminals. The mapping is non-deterministic as a non-terminal can be mapped to more than one sequence. The rule is divided into a head and a body. What is on the left of the rule is known as the head and the body is on the right.

EBNF extends BNF by allowing elements of regular expressions to appear in production rules. This does not make BNF more powerful but it is a convenience.

³<https://javacc.java.net>

The full grammar for the RefDstl language is presented in Appendix A. We used this grammar to generate a parser and metamodel for the RefDstl language with Xtext. Auxiliary tools such as a plug-in editor for Eclipse were also produced.

3.4.3 Substituting RefDstl

There is an additional and somewhat subtle benefit from using the Xtext toolset to develop our language. That is to say that when an Xtext grammar file is compiled it produces an Ecore metamodel. The parser that Xtext also produces builds instances of this metamodel that represent the parsed language. However, the metamodel is readily available for use by tools. We propose that researchers and developers who work on tools that involve building in-memory models of refactorings or indeed serialising them to a file should use our metamodel. This allows those researchers and developers to take advantage of some of MDE's main benefits, i.e. tool re-use and portability. There is a downside however to instantiating the metamodel programmatically, i.e. the metamodel can be instantiated in a way that would never be possible using text input to the generated parser.

3.5 Defining a Standard Library of Java Refactorings

We have thus far described the goals and design guidelines for RefDstl. With those goals in mind we described a list of constructs required and specified the syntax of the language. We now provide a standard library of refactorings, which we use in the sample scripts used in the following section. The code that makes up the standard library is presented in Appendix B. It shows how to use the language, in particular how to make procedure calls and write preconditions over the model representing the program being refactored. All of the primitive refactoring instructions that the language offers are covered in the library; they are in the bodies of the numerous procedures.

On the surface, it might seem unusual that we wrap the instructions of our language inside procedures that we encourage others to use. The reason for this is that the instructions do not have any preconditions associated with them. The interpreter will execute the instructions regardless of the state of the program being refactored. The only way in RefDstl to annotate the instructions with preconditions is by placing them inside procedures as we have done. Users of RefDstl can freely replace our library with their own if they find ours to be insufficient for their needs. It is worth noting that our library does not come with any postconditions. These can be included by either editing the library itself or by extending it. We omitted postconditions as we were trying to adhere to the primitive refactorings given by Fowler (1999).

3.6 A Small Case Study

In this section, we use RefDstl to repeat a small case study performed by O'Keeffe and Ó Cinnéide (2003).

A UML class diagram for the original design of a fictitious video game is shown in Figure 3.1. At the top of the class hierarchy is the `Weapon` class with a private attribute called `power` that is accessible to the five child classes via the public `getPower` method. Its immediate children, `MissileWeapon` and `MeleeWeapon`, both have an attribute called `range`. We might say that these two child classes exhibit the 'duplicate code smell'. The re-course here would be to 'pull up' the `range` field and the `getRange` method to the `Weapon` class to aim for the design presented in Figure 3.2. O'Keeffe and Ó Cinnéide (2003) call this maximising inheritance.

The RefDsl script for performing this refactoring is listed in Listing 3.1. At the beginning of the script, we import the standard library of refactoring primitives, which were discussed in the previous section. The next two lines instruct the interpreter about what model is to be refactored and where the resulting MoDisco model of the refactored system is to be persisted. The declaration of the `main` procedure follows. This is where the script begins its execution. It must be at the top of the script following the input/output directives. In this script, the main procedure delegates its task to a procedure called `safeRefactor`. This delegation may seem wasteful but the main procedure cannot be annotated with pre or post conditions so placing the refactoring instructions inside another procedure allows us to work around this. Future versions will remove this restriction. Note that in the RefDsl scripts that the tilde symbol is used in postconditions to mean ‘the version of this class prior to the refactoring’.

The definition of `safeRefactor` declares six preconditions identifiable by the `requires` keyword. All of these preconditions serve a similar purpose: to ensure that only existing classes are refactored. If any of the classes do not exist then the interpreter will halt executing the script. The preconditions take advantage of the `$utils` script handle that has a method provided named `classExists`. It returns a boolean value to indicate if a class by a particular fully qualified name is declared in the project undergoing refactoring. Other, more interesting, preconditions are annotated on the procedures that are defined in the standard library of refactorings (see Appendix B) and used in this script. Those preconditions are not repeated in Listing 3.1 since the purpose of the standard library is to minimise repetition of preconditions.

The definition of `safeRefactor` declares twelve postconditions that are indicated by the `ensures` keyword. These reason about the available methods and methods inherited metrics, which are calculated using the `am` and `mi` operations from the McQuillan (2011) metrics meta-model that is accessed using the `$metrics` script handle. The first postcondition⁴ states that the `Weapon` class will inherit the same amount of methods when the `getRange` method is pulled up. Following this, we state that `MissileWeapon` and `MeleeWeapon` will inherit more methods after the refactoring. This is to be expected since we are pulling a method from each of these classes and locating it in the `Weapon` class. The classes `Bow`, `Javelin` and `Sword` continue to inherit the same number of methods. We then reason about the available methods. Naturally the `Weapon` class will have more methods available after the refactoring because the `getRange` method will be then declared in it. The classes that inherit from `Weapon` either directly or indirectly will have the same number of methods available after the refactoring since they will be able to access the `getRange` method from their parent classes.

```
1 imports "standard_library.refdsl"
2
3 in "model_original/model_java2kdm.xmi"
4 out "max_inheritance/model_java2kdm.xmi"
5
6 main() {
7   call safeRefactor()
8 }
9
10 /**
```

⁴A discussion covering how these are implemented, in particular ‘old’ variables, is reserved until subsection 4.2.5.

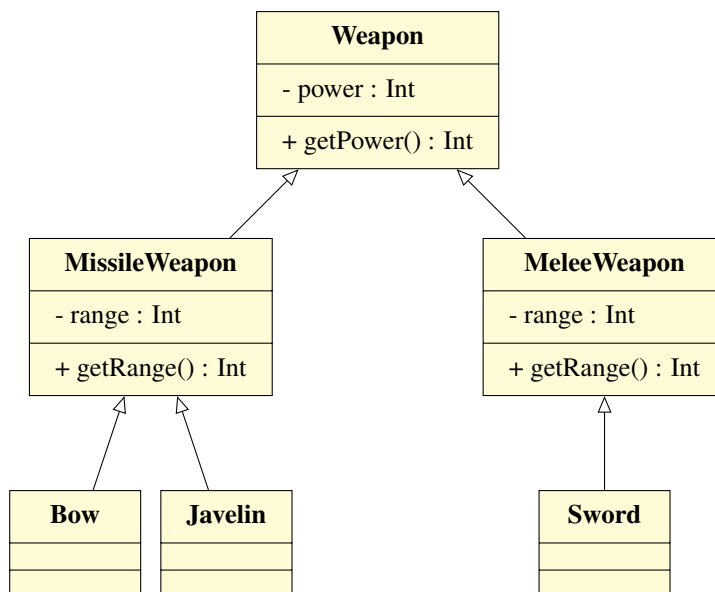


Figure 3.1: A UML class diagram demonstrating a poorly designed system to be refactored.

```

11 * Performs a pull up refactoring on the getRange procedure
12 * and range field.
13 *
14 * It also ensures that the metric values are changed appropriately.
15 */
16 proc safeRefactor()
17   requires $utils.classExists("Weapon")
18   requires $utils.classExists("MissileWeapon")
19   requires $utils.classExists("MeleeWeapon")
20   requires $utils.classExists("Bow")
21   requires $utils.classExists("Javelin")
22   requires $utils.classExists("Sword")
23   ensures $metrics.mi($utils.findClass("Weapon")) == $metrics.mi($utils.
24     findClass("~Weapon"))
25   ensures $metrics.mi($utils.findClass("MissileWeapon")) > $metrics.mi($utils.
26     findClass("~MissileWeapon"))
27   ensures $metrics.mi($utils.findClass("MeleeWeapon")) > $metrics.mi($utils.
28     findClass("~MeleeWeapon"))
29   ensures $metrics.mi($utils.findClass("Bow")) == $metrics.mi($utils.findClass("
30     ~Bow"))
31   ensures $metrics.mi($utils.findClass("Javelin")) == $metrics.mi($utils.
32     findClass("~Javelin"))
33   ensures $metrics.mi($utils.findClass("Sword")) == $metrics.mi($utils.findClass
34     ("~Sword"))
35   ensures $metrics.am($utils.findClass("Weapon")) > $metrics.am($utils.findClass
36     ("~Weapon"))
37   ensures $metrics.am($utils.findClass("MissileWeapon")) == $metrics.am($utils.
38     findClass("~MissileWeapon"))
39   ensures $metrics.am($utils.findClass("MeleeWeapon")) == $metrics.am($utils.
40     findClass("~MeleeWeapon"))
41   ensures $metrics.am($utils.findClass("Bow")) == $metrics.am($utils.findClass("
42     ~Bow"))

```

```

33  ensures $metrics.am($utils.findClass("Javelin")) == $metrics.am($utils.
      findClass("~Javelin"))
34  ensures $metrics.am($utils.findClass("Sword")) == $metrics.am($utils.findClass
      ("~Sword"))
35  {
36    call refactorField()
37    call refactorMethod()
38  }
39
40  /**
41   * This performs the pull up field refactoring on the range field
42   * range from classes MissileWeapon and MeleeWeapon.
43   */
44  proc refactorField() {
45    call moveFieldToSuperclass(range, MissileWeapon)
46    call deleteField(range, MeleeWeapon)
47  }
48
49  /**
50   * This pulls up the getRange method from MissileWeapon and MeleeWeapon.
51   */
52  proc refactorMethod() {
53    call moveMethodToSuperclass(getRange, MissileWeapon)
54    call deleteMethod(getRange, MeleeWeapon)
55  }
56

```

Listing 3.1: A RefDsl script that performs the pull up refactoring on the range field and getRange method for the MissileWeapon and MeleeWeapon classes presented in Figure 3.1.

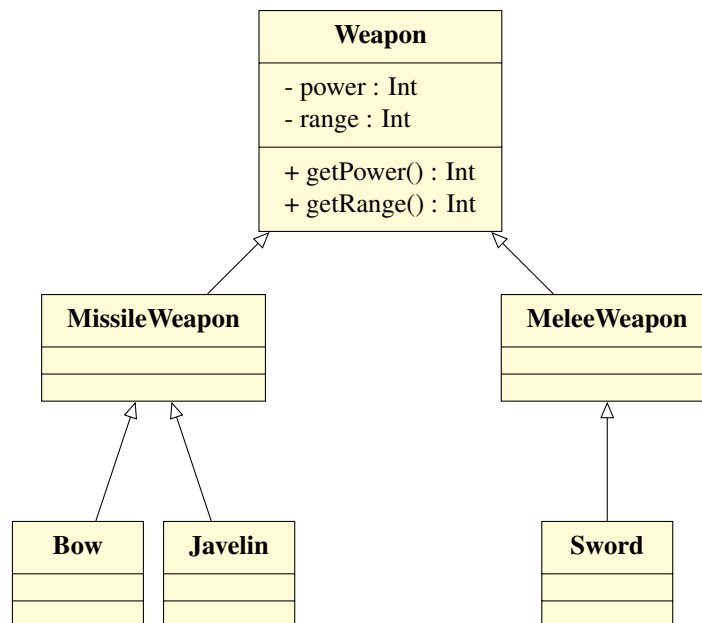


Figure 3.2: A UML class diagram that represents the system presented in Figure 3.1 refactored to obtain maximum inheritance.

This design is not much better. It means every class has a larger than necessary interface. Software engineers generally try to decrease the size of the interface to reduce the risk of afferent coupling. It provides them with leverage to be able to eliminate one or more of the methods without having to consider the ramifications for client code. To minimise the interface offered by the classes we aim for the design presented in Figure 3.3 and provide a RefDsl script for doing so in Listing 3.2. The structure of this script differs to Listing 3.1 in that we now use the `for` construct in the definition of the `buryMethod` procedure. It repeatedly applies the push down refactoring to send an entity to the lowest level of a class hierarchy.

```

1 imports "standard_library.refdsl"
2
3 in "model_original/model_java2kdm.xmi"
4 out "min_methods/model_java2kdm.xmi"
5
6 main() {
7   call safeRefactor()
8 }
9
10 /**
11  * Safely pulls up the range field into the Weapon class and pushes down
12  * the getPower and getRange methods.
13  *
14  * It also ensures that the metric measurements for the available methods and
15  * methods inherited
16  * measurements update appropriately.
17  */
18 proc safeRefactor()
19   requires $utils.classExists("Weapon")
20   requires $utils.classExists("MissileWeapon")
21   requires $utils.classExists("MeleeWeapon")
22   requires $utils.classExists("Bow")
23   requires $utils.classExists("Javelin")
24   requires $utils.classExists("Sword")
25   ensures $metrics.mi($utils.findClass("Weapon")) == $metrics.mi($utils.
26     findClass("~Weapon"))
27   ensures $metrics.mi($utils.findClass("MissileWeapon")) < $metrics.mi($utils.
28     findClass("~MissileWeappn"))
29   ensures $metrics.mi($utils.findClass("MeleeWeapon")) < $metrics.mi($utils.
30     findClass("~MeleeWeapon"))
31   ensures $metrics.mi($utils.findClass("Bow")) < $metrics.mi($utils.findClass("~
32     Bow"))
33   ensures $metrics.mi($utils.findClass("Javelin")) < $metrics.mi($utils.
34     findClass("~Javelin"))
35   ensures $metrics.mi($utils.findClass("Sword")) < $metrics.mi($utils.findClass(
36     "~Sword"))
37   ensures $metrics.am($utils.findClass("Weapon")) < $metrics.am($utils.findClass
38     ("~Weapon"))
39   ensures $metrics.am($utils.findClass("MissileWeapon")) < $metrics.am($utils.
40     findClass("~MissileWeappn"))
41   ensures $metrics.am($utils.findClass("MeleeWeapon")) < $metrics.am($utils.
42     findClass("~MeleeWeapon"))

```

```

33 ensures $metrics.am($utils.findClass("Bow")) == $metrics.am($utils.findClass("
    ~Bow"))
34 ensures $metrics.am($utils.findClass("Javelin")) == $metrics.am($utils.
    findClass("~Javelin"))
35 ensures $metrics.am($utils.findClass("Sword")) == $metrics.am($utils.findClass
    ("~Sword"))
36 {
37   call refactorField()
38   call refactorMethod()
39 }
40
41 /**
42  * This performs the pull up field refactoring on the field:
43  * range from classes MissileWeapon and MeleeWeapon.
44  */
45 proc refactorField()
46 {
47   call moveFieldToSuperclass(range, MissileWeapon)
48   call deleteField(range, MeleeWeapon)
49 }
50
51
52 /**
53  * The getPower method will be moved to the lowest levels in
54  * the inheritance hierarchy.
55  */
56 proc refactorMethod() {
57   call buryMethod(getPower, Weapon)
58   call buryMethod(getRange, MissileWeapon)
59   call buryMethod(getRange, MeleeWeapon)
60 }
61
62 /**
63  * The bury method refactoring sends a method to the
64  * lowest classes in the inheritance hierarchy.
65  */
66 proc buryMethod($methodName, $class) {
67   if($class.isParent()) {
68     call moveMethodToSubclass($methodName, $class)
69
70     for($cls in $class.children()) {
71       call buryMethod($methodName, $cls)
72     }
73   }
74 }
75

```

Listing 3.2: A RefDsl script that performs a pull up refactoring on the range field and a push down refactoring on the getPower method until it is at the lowest level of the inheritance hierarchy.

Although this is an improvement, a problem remains. In this design, Sword now has access

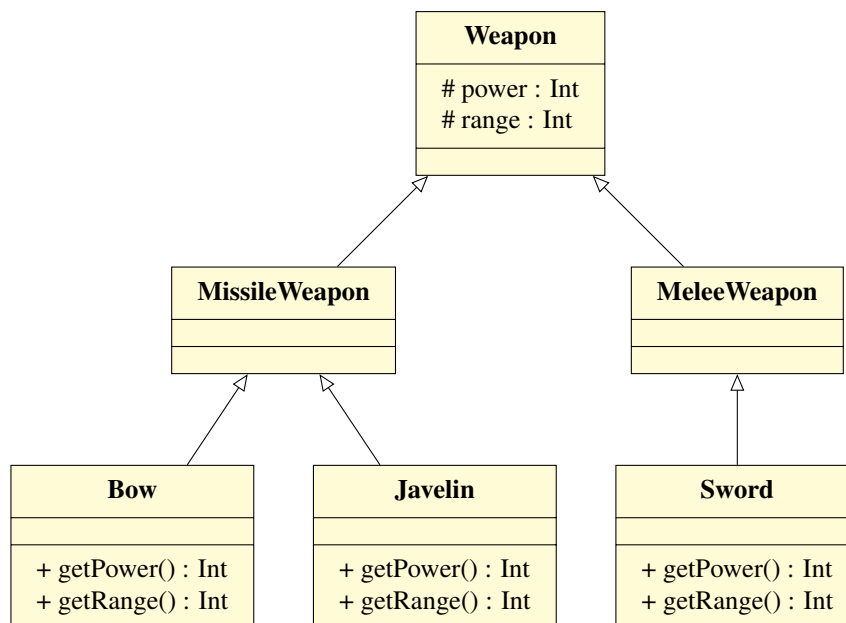


Figure 3.3: The UML class diagram from Figure 3.1 refactored for minimum methods per class.

to a method called `getRange`. A `MeleeWeapon` should not have a concept of range. We perform a final refactoring on this design to correct this and aim for the design shown in Figure 3.4, which is yielded by the `RefDsl` script in Listing 3.3. This is the optimised design according to search based approach to refactoring by O’Keeffe and Ó Cinnéide (2003).

```

1 imports "standard_library.refdsl"
2
3 in "model_original/model_java2kdm.xmi"
4 out "optimised/model_java2kdm.xmi"
5
6 main() {
7   call safeRefactor()
8 }
9
10 /**
11  * Safely pulls up the range field and removes the getRange method
12  * from the MeleeWeapon class.
13  *
14  * The metric measurements are also check to ensure they are updated
15  * appropriately.
16  */
17 proc safeRefactor()
18   requires $utils.classExists("Weapon")
19   requires $utils.classExists("MissileWeapon")
20   requires $utils.classExists("MeleeWeapon")
21   requires $utils.classExists("Bow")
22   requires $utils.classExists("Javelin")
23   requires $utils.classExists("Sword")
24   ensures $metrics.mi($utils.findClass("Weapon")) == $metrics.mi($utils.
    findClass("~Weapon"))
  
```

```

24 ensures $metrics.mi($utils.findClass("MissileWeapon")) == $metrics.mi($utils.
    findClass("~MissileWeappn"))
25 ensures $metrics.mi($utils.findClass("MeleeWeapon")) == $metrics.mi($utils.
    findClass("~MeleeWeapon"))
26 ensures $metrics.mi($utils.findClass("Bow")) == $metrics.mi($utils.findClass("
    ~Bow"))
27 ensures $metrics.mi($utils.findClass("Javelin")) == $metrics.mi($utils.
    findClass("~Javelin"))
28 ensures $metrics.mi($utils.findClass("Sword")) < $metrics.mi($utils.findClass(
    "~Sword"))
29 ensures $metrics.am($utils.findClass("Weapon")) == $metrics.am($utils.
    findClass("~Weapon"))
30 ensures $metrics.am($utils.findClass("MissileWeapon")) == $metrics.am($utils.
    findClass("~MissileWeappn"))
31 ensures $metrics.am($utils.findClass("MeleeWeapon")) < $metrics.am($utils.
    findClass("~MeleeWeapon"))
32 ensures $metrics.am($utils.findClass("Bow")) == $metrics.am($utils.findClass("
    ~Bow"))
33 ensures $metrics.am($utils.findClass("Javelin")) == $metrics.am($utils.
    findClass("~Javelin"))
34 ensures $metrics.am($utils.findClass("Sword")) < $metrics.am($utils.findClass(
    "~Sword"))
35 {
36   call refactorField()
37   call refactorMethod()
38 }
39
40 /**
41  * This performs the pull up field refactoring on the field:
42  * range from classes MissileWeapon and MeleeWeapopn.
43  */
44 proc refactorField() {
45   call moveFieldToSuperclass(range, MissileWeapon)
46   call deleteField(range, MeleeWeapon)
47 }
48
49 /**
50  * This will remove the getRange method from the MeleeWeapon
51  * class.
52  */
53 proc refactorMethod() {
54   call deleteMethod(getRange, MeleeWeapon)
55 }
56

```

Listing 3.3: A RefDsl script that performs a pull up refactoring on the range field and a push down refactoring on the getPower method until it is at the lowest level of the inheritance hierarchy.

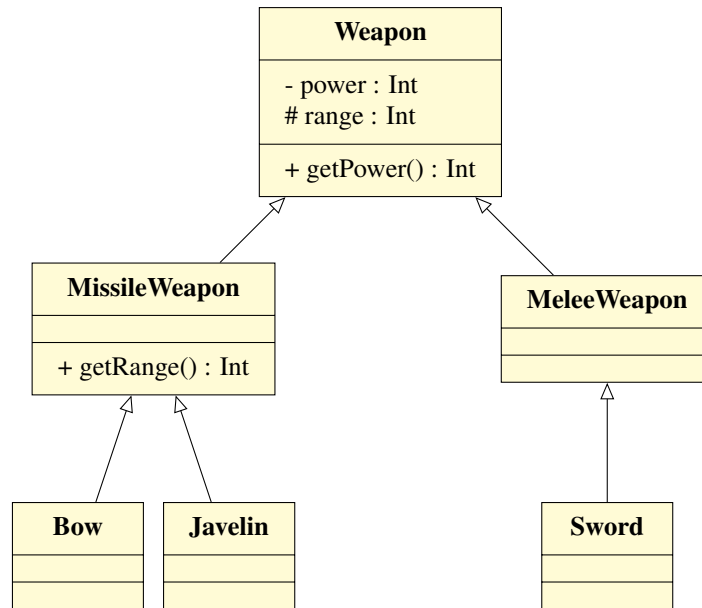


Figure 3.4: The UML class diagram from Figure 3.1 refactored using the optimisations calculated by the Dearthóir tool.

3.7 Summary

This chapter described a language to define and script software refactorings. The language was built using the Xtext language framework. It not only produced supporting tools for the language, such as a parser and Eclipse IDE integration, but it also provided a metamodel that can be instantiated by other researchers working in the area of refactoring. This metamodel allows for compatibility between tools. We also provided a standard library of primitive refactoring operations. These are used in our evaluation of RefDsl in Chapter 5. In the following chapter we discuss implementing an interpreter for the language in the form of an extensible framework that can refactor various kinds of software artefacts beyond the usual limits of source code.

Chapter 4

A Model Driven Framework to Interpret RefDstl Scripts

In the previous chapter, we discussed the syntax and implementation of the RefDstl language for specifying and scripting refactorings. We now focus on developing a framework for interpreting those scripts. We begin by describing its high level design and identify the various components that collectively form its architecture. In the latter part of this chapter, we describe how the framework can be extended with additional model transformations so that it is relevant for refactoring many heterogeneous artefacts rather than just source code. We also describe how the RefDstl language can be enriched with the notion of ‘script handles’.

4.1 High–Level Overview of the RefDstl Interpreter Framework

Figure 4.1 presents a UML component diagram illustrating the main components of the RefDstl framework. We have omitted the classes inside these components from the diagram for the sake of simplicity. We describe the general purpose of each of these components and how they interact before giving a more detailed discussion about each in turn.

We saw in Chapter 3 that the RefDstl language permits modularity by allowing scripts to be split into multiple files. However, before a RefDstl script can be executed those files must be ‘linked’ together so that procedures defined in external libraries, for example, can be resolved. The `ScriptManager` component facilitates this need. Once a script has been linked then execution is ready to commence. The `ScriptManager` relies on the `InstructionInterpreter` interface provided by the `Controllers` component to initiate execution.

The `Controllers` component contains multiple classes, which are each responsible for interpreting specific constructs and instructions provided by the RefDstl language. However, the controllers delegate most of their work. For example, the controller responsible for invoking procedures delegates the task of maintaining the environment (i.e. the mapping of values to variables) to the `Environment` component, which it accesses through the `Stack` interface. Furthermore, the task of interpreting the pre and post conditions on procedures is delegated to the `ConditionEvaluator` component through the `Interpreter` interface. The `ConditionEvaluator` component evaluates the truthness of boolean formulae. The ‘heavy lifting’ required to perform model transformations is delegated to to the `Transformations` component, which is accessed via the `CompositeTransformation` interface. The `Transformations` component contains multiple classes, each of which understand how to load, transform and persist

different (or heterogeneous) model types. The `ScriptHandles` component is also called upon by classes in the `Controllers` component. It provides utility methods that can be used in RefDsl scripts to enrich the language with features that are otherwise unavailable, for example, outputting messages to log files or the console.

We now elaborate on the details of each of these components.

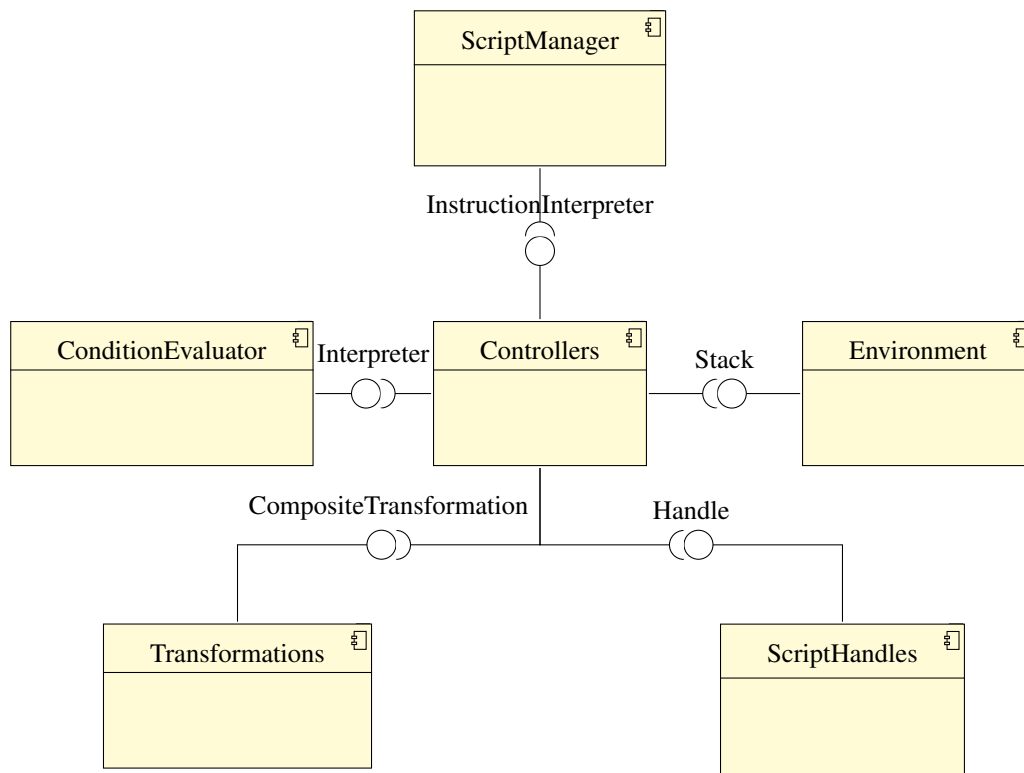


Figure 4.1: An overview of the components in the RefDsl interpreter framework.

Script Manager

The script manager has a similar purpose to a ‘linker’ used with traditional compiled software. It loads RefDsl scripts and inspects them for `import` statements. For example, the code snippet in Listing 4.1 shows a RefDsl script importing the code from a library called `my_refactorings.refdsl`. An `import` is equivalent to taking the procedures defined in the imported script and appending them to the bottom of the importing script. Of course an imported script, such as `my_refactorings.refdsl`, might also have its own `import` statements. These are recursively resolved until no more imports remain. Circular dependencies between imports should be avoided and the current implementation of RefDsl is unable to detect these. In a situation where one script and a subsequently imported script share a procedure with the same name then the procedure that was loaded the earliest is used. Parameter types are not used as part of the procedure’s signature; thus, overloaded procedures are unsupported and will result in a runtime error.

```

1 imports "my_refactorings.refdsl"
2
3 in "some_model.xmi"

```

```
4 out "some_model_refactored.xml"
5
6 /**
7  * Script begins execution in main.
8  */
9 main() {
10 ...
11 }
```

Listing 4.1: An example showing the syntax used to import one RefDsl script into another.

Controllers

Once a script has been fully loaded, it begins to execute. For each of the different refactoring instructions (create class, move field etc.) we implemented a controller. Controllers also exist that are specific to aspects of the RefDsl language, such as procedure calls.

The controllers responsible for interpreting refactoring instructions are mundane delegates. They merely extract the parameters for the refactoring instruction (such as the new name for a class in a ‘create class’ instruction) from the script and resolve any variables to their values using the `Environment` component. They then pass the parameters to the appropriate method in `CompositeTransformation` singleton object in the `Transformations` component. This repeats the message to all of the transformation classes registered in the framework, such as the transformation for refactoring the MoDisco model of the system or the transformation responsible for the metrics model of the system.

Other controllers that are specific to RefDsl language constructs have more interesting and involved activities. For example, the `MainInt` class in the `Controllers` component is responsible for beginning execution of the script. It attempts to locate the `main` procedure in the refactoring script and begins to execute the list of instructions in the body. Since a controller only knows how to interpret one type of instruction, it dispatches everything that it does not understand to the general instruction controller, named `InstructionInt`. This controller inspects the name of the actual type of the instruction and by using Java’s reflection capabilities it instantiates a new instance of the correct controller to perform the interpretation and further delegation occurs.

The `ConditionalStatementInt` is responsible for interpreting conditional `if` statements in the language. Guards on conditional statements use the same syntax as expressions in pre and post conditions so the interpretation of the guard is delegated to the conditions evaluator (discussed below). If the guard evaluates to true then `ConditionalStatementInt` evaluates the body of the condition through further use of delegation. Otherwise, it does nothing. For example, Listing 4.2 presents a procedure that will rename a class declaration to `Bar` if its name is currently `Foo`. The procedure takes as input a class declaration and the conditional statement inside the body of the procedure uses a guard to determine if the name of the passed class declaration is named `Foo` (in which case the guard evaluates to boolean true). If the guard is true then `$possibleFoo` will be renamed otherwise no action is taken.

```
1 /**
2  * This method will rename a class called Foo to Bar.
3  * Otherwise it does nothing.
4  */
```

```

5 renameFooToBar( $possibleFoo ) {
6     if( $possibleFoo.getName().equals(Foo)) {
7         call renameClass( $possibleFoo , Bar)
8     }
9 }

```

Listing 4.2: An example showing the syntax used to describe a conditional statement in RefDsl.

The `LoopingStatementInt` is responsible for performing a sequence of instructions over a collection. It extracts the object expression from the loop statement header. It delegates the task of interpreting this to the `ObjectExpressionInt`. The `ObjectExpressionInt` returns the result of the expression, which the `LoopingStatementInt` casts to an instance of a collection. Trying to loop over anything other than a collection is an obvious error. The `LoopingStatementInt` iterates over each member of the collection. During each iteration, it maps the current element to a variable name (determined by the script) in the environment (described later) and then executes each instruction in the loop body. In Listing 4.3, for example, we iterate over all of the classes in the project under refactoring and create an instance variable for logging. The `ObjectExpressionInt` evaluates `$utils.allClasses()` to a collection and the `LoopingStatementInt` binds each member of this collection to the variable named `$cls` during each iteration. The body of the loop is applied to the bound variable during each iteration, which in this case creates the variable for logging.

```

1 /**
2  * This loop will add a variable for logging to each class.
3  */
4 for( $cls in $utils.allClasses() ) {
5     call createField( LOG, public, $cls.getName(), "java.util.Logger" )
6 }

```

Listing 4.3: An example showing how to iterate in RefDsl.

The `ObjectExpressionInt` evaluates object expressions. These begin with a variable followed by a cascade of method calls. The parameters to these method calls can also be other object expressions. For example, in Listing 4.4 there are two object expressions that we use to compare two class names for equality. The first object expression to be executed is the nested expression `$clsTwo.getName()`. The `ObjectExpressionInt` locates the variable `$clsTwo` from the environment and then using Java's reflection mechanism invokes the `getName()` method. The resulting object is then cached. Next, `ObjectExpressionInt` locates the object assigned to the variable `$clsOne` and invokes the `getName()` method on it before, finally, evaluating the cascaded method `equals(...)`. The earlier cached result is passed as a parameter. The result of all object expressions is an object; the actual type such as `String` is unneeded.

```

1 $clsOne.getName().equals( $clsTwo.getName() )

```

Listing 4.4: An example object expression in RefDsl.

`ProcedureCallInt` handles RefDsl call instructions. In the case where no arguments are being passed with the procedure call this is simply a case of searching the script for the procedure with the same name and then proceeding to have `Controller` interpret each of the statements in

the procedure's body in turn. However, prior to executing the body `ProcedureCallInt` requests that the environment create a new frame on the stack, which is later destroyed when the procedure is finished executing. The purpose of the frame is to provide scope to the variables. Without frames the instructions in the body of a procedure would be able to read the parameters of the invoking procedure. This could lead to unusual behaviour. Once the procedure is finished executing the frame is destroyed for efficient memory usage. The `ProcedureCallInt` class is also responsible for placing the arguments on the stack frame. This is achieved by getting the names of the *parameters* that the procedure being called expects and the variable names or values of the *arguments* that are being provided. If the size of the argument list does not match the parameter list then an error is raised. Otherwise, on the frame created for the procedure call, each i^{th} argument is mapped to each i^{th} variable name. The procedure's body is then executed using, once again, more delegation.

Transformations

Model transformations are the primary motivation for this framework. There are no technical restrictions on the kind of models that can be transformed. However, the framework is intended to be used with transformations that are driven by refactoring instructions. The actual act of transforming the model is not required to be done via Java as we have done in our provided transformations. Developers extending the framework can choose to have the transformation conducted by calling out to an ATL script, for example. We believe this flexibility is a benefit of our framework.

We provide two transformations with the framework. The first is `MoDiscoTransformation`. If the framework is equipped with just this transformation then it behaves as any regular refactoring engine. Its responsibility is to transform `MoDisco` models that are representations of the Java system being refactored. It also provides a handle named `$project` in scripts. This can be used in the same way as variables in the `RefDstl` language. We do not go into detail about how this transformation actually refactors models. However, we point out a distinguishing feature that distinguishes the framework from other refactoring engines. The preconditions that most refactoring engines would check prior to the transformation are not embedded in the Java code that performs the transformations. Instead, the transformation assumes that the conditions have already been verified prior to the request to perform the transformation. In other words, the conditions must be expressed in the refactoring script in pre and post conditions on procedures. This is not to suggest that decoupling the conditions from the transformation make them easier to formulate but we do believe it makes it easier to correct erroneous conditions in the field. Also it is more convenient to add new transformations.

The second transformation we provide is the `MetricsTransformation`. The idea behind this is based on the work of McQuillan (2011) but the implementation provided here is our own. In this transformation, we maintain a representation of the system using McQuillan's *measurement* metamodel. The measurement metamodel contains a view of the measurable aspects of the system that are used by McQuillan's *metrics* metamodel. The metrics metamodel contains definitions of a variety of metrics as OCL expressions. When these expressions are evaluated they provide a numeric value. This transformation also provides a 'script handle' called `$metrics`. This is significant because it means that measurement values can be reasoned about in pre and postconditions written in scripts. In other words, `RefDstl` allows the script developer to define refactorings that will only be applied in the case where the quality of the software undergoing refactoring is

improved.

Condition Evaluator

From the RefDsl code examples previously shown, it is clear that there is a disparity between the syntax used to express pre and post conditions and the syntax used to describe refactoring operations. The disparity exists as a result of a language design choice to use a syntax for the conditions that would be familiar to programmers who have experience expressing pre and post conditions in other languages such as Spec#. However, this syntax is unsuitable to describe refactoring operations so we devised a declarative style reminiscent of SQL. It could be said that our language is really two languages in one and we held this in mind during the framework's design. For the sake of modularity, we placed the conditions evaluator into a separate component for this reason. The interpreter for evaluating conditions is called `ConditionInt`. Other than this class, the component contains one condition evaluator for each of the five levels of operator precedence in the RefDsl language. The `ConditionInt` delegates any condition to be evaluated to the interpreter class at the correct precedence.

Environment

`Environment` is responsible for maintaining state during script execution. It does this by providing a stack interface to clients. The stack is used much in the same way as in assembly programming to maintain local variables and for passing parameters to procedures. However, our implementation allows the client to explicitly create and destroy frames (which is done prior to a procedure call and immediately after). This prevents erroneous access of variables that are out of scope. Our environment does not need to provide a heap to store objects. Objects in RefDsl are just Java objects; hence they 'live' on the heap maintained by the JVM. When a reference to one of the objects is no longer on one of the stack frames maintained by `Environment` then that object can no longer be accessed so the object is garbage collected.

All variables that can be declared by RefDsl programmers have local scope. These can be accessed anywhere in the body of the procedure or in the conditions annotated on the procedure. Variables can also be declared as part of `forall` or `exists` constructs in conditions. These are only available in the body of the condition. The same applies to variables declared in the guard of `if` constructs. In loop constructs, variables can be declared in the header. These are available anywhere in the body of the loop construct including inside the bodies of any nested constructs. There are certain variables that are global. These are provided by 'script handles' (discussed later). For example, the `MoDiscoTransformation` provides a handle called `$project` that allows users to reason about the program that is being refactored. Although local and global variables are accessed and modified the same way, there is one major distinction in how they need to be treated in post conditions on procedures. In post conditions, it will be necessary to talk about the state of global variables prior to the procedure being executed so that conditions can be written to ensure that the model was refactored correctly. The environment supports features to enable this. It allows a deep copy to be made of a frame and of the objects/values in that frame. To access the old values, RefDsl programmers can use this syntax: `old($someGlobalVar)` to gain access to the value of `$someGlobalVar` prior to the procedure being executed. Local variables cannot be reassigned in the RefDsl language so they have no 'old' state.

Script Handles

The `ScriptHandles` component contains extension points for the RefDsl language. Other than when the RefDsl framework ‘starts up’, and registers each of the handles with the `Environment` component, the other components are unaware of its existence. Each script handle simply becomes another global variable. We discuss how script handles are developed in the next section.

4.2 Enriching the RefDsl Language via Script Handles

Script handles are one of the extension points on offer to developers extending the basic functionality of the RefDsl framework. They can be included to either provide additional features in the language or as a substitute for ‘macros’ so that common expressions can be expressed succinctly. For example, Listing 4.5 shows how a script handle we included can be used to reduce the number of lines of code required to reason over all of the classes in the system being refactored. Without this script handle, we would have to use just the methods provided by the `MoDisco` representation of the system under refactoring, as shown in Listing 4.6.

```
requires forall $cls in $utils.allClasses() : ( . . . )
```

Listing 4.5: Writing a precondition that uses the utility method in the `$utils` handle for convenience.

```
requires forall $cu in $project.getCompilationUnits() : forall $cls in $cu.
  getTypes() : $instance.Of($cls, ClassDeclaration) -> ( . . . )
```

Listing 4.6: Pure RefDsl code for universally quantifying over all the classes in a system

In this section, we discuss how developers can integrate their own script handles into the framework. We then briefly describe four of the script handles that we developed. Note that script handles are accessed via variables that are treated as global by the environment.

4.2.1 Embedding a New Handle

Embedding a new script handle into RefDsl involves following these steps:

1. Create a Singleton Class — The steps involved for creating a singleton class are described by Metsker and Wake (2006). Roughly speaking it involves providing a default constructor with a private access modifier so that instances of the class can only be created from inside the class. The class also contains a reference to an object (having the same type as the class) which is marked static and private and is initialised at declaration. This ensures that only one object with that class’s type can ever exist. The object can be accessed from outside the class by creating a static method (normally called `instance()`).
2. Implement the interface presented in Table 4.1 — The script handle must implement this interface so that the framework registers it correctly. The `getHandleName` method should return a `String` object. This will be the name of the variable that RefDsl script programmers use to access the script handle. `getHandleObject()` returns an instance of object. This should be the singleton object that is created as a part of enforcing the singleton design pattern. The class is free to contain any other methods. Methods that are declared public can be accessed in the script using the script handle.

Method Name	Purpose
<code>getHandleName</code>	This provides the name of the handle that can be accessed through the conditions. It is the responsibility of the person writing the script handle to make sure that the handle name is not already in use.
<code>getHandleObject</code>	This should return the instance of the handler class.

Table 4.1: The interface that must be implemented by a class to integrate it into the RefDsl framework as a script handle.

3. Register the handle in the framework — The script handle can be registered in the framework by modifying the constructor of the class `MachineStack` to include a line such as Listing 4.7. The `MachineStack` will then include a variable on every stack frame that points to the script handle's instance using its provided variable name. In the current implementation of RefDsl, there is no convenient way to determine what script handles are used in some script. In a future version, the interpreter could perform a static analysis of the script prior to execution to look for undeclared variables. Any undeclared variables could be suggested to be required script handles. However, this is not a guarantee as the script's author may have just misspelled a declared variable.

We discuss the script handles we implemented next.

```
scriptHandles.add(Utils.instance());
```

Listing 4.7: Java code that is used to register the `utils` script handle object with the RefDsl framework

4.2.2 Supporting Debugging in RefDsl Scripts

There are no debugging tools currently available for RefDsl scripts nor is there any native support for I/O in the language. For this reason, we introduced a crude form of debugging by developing a script handle that allows tagged messages to be written to the console and can also be processed by the Log4J system¹.

The plug-in can be accessed by using the `$debug` handle. It offers a method named `print` that takes a message to display or a tag followed by a message. A tag should be any of the Log4J levels such as `trace`, `debug`, `error`, `fatal`, `warn` or `info`. Depending on the tag, the debug message will be output at the most appropriate Log4J level and to the configured logger.

4.2.3 Checking Object Types

While writing the standard library of primitive refactoring operations we found that it was useful to be able to test the type of objects. For this purpose, we provided the `instance` script handler. It contains a method called `of` which accepts two parameters. The first argument should be an object and the second should be a string that is a type name. The method will test if the object is an instance of that type name and return the result as a boolean.

4.2.4 Accessing the Program Model

The model of the system can be accessed using the `$project` handle. This allows for full access to the project under refactoring's MoDisco representation. Any methods provided by

¹<http://logging.apache.org/log4j/2.x/>

the Java class implementation of MoDisco are available. For example, users can call the `getCompilationUnits` method to get the list of Java files in the project being refactored.

4.2.5 Accessing Old Variables

As discussed previously, global variables require two states to be maintained during the execution of a procedure in order to support postconditions: the state prior to execution and the post state. In a postcondition expression, the post value of the variable can be expressed by using the variable name — no special syntax is needed. However, the value of the variable prior to execution requires extra syntax. We provide this by using a script handle. The handle offers a single method called `valueOf` which takes a string object as a parameter. It reads the value for that variable from the frozen frames of the `MachineStack` and returns the value.

4.3 Custom Transformations into the RefDsl Framework

Users can integrate their own transformations into the framework by creating a singleton class that implements the `ITransformation` interface. The methods of this interface are documented in Table 4.2.

The user can create transformations for any type of representation that they choose but it is expected that all representations loaded during the execution of a script are just different views of the same model. However, there is no necessity for them to have been generated from some common core model.

We offer no guidance here regarding how to write the actual transformation code; this will depend very much on the representation the transformation is intended to refactor.

Method Name	Meaning
<code>nop</code>	In situations when a RefDsl programmer would like for a procedure, conditional or looping body to be empty they use a <code>nop</code> instruction. Transformations are not required to perform any action when a <code>nop</code> instruction is encountered.
<code>save</code>	A transformation is expected to have a reference to some model representation. The transformation should implement the <code>save</code> method to persist that model to a file or otherwise.
<code>deleteClass</code>	The script encountered a delete class instruction. The transformation should update the model representation by removing the class and any references to it.
<code>deleteField</code>	A delete field instruction was executed. The transformation should delete the field from its representation and any references to it.
<code>createClass</code>	A request to create a new class was encountered in the script. The model representation should create a new unreferenced class with the specified parameters.
<code>renameClass</code>	The script contained a rename class instruction. The transformation should modify the model by finding the class with the old name and update it with the new name.
<code>createField</code>	A new and unreferenced field should be created in a specified class.

createMethod	The transformation should create a new unreferenced method in the model representation.
deleteMethod	The transformation should search the model for a method with a specific signature and delete it as well as any associated references.
renameField	The field with a specific name should be found and its name updated to the name provided by the parameters.
changeMethodName	The model transformation should find a method with a particular name and signature and update its name.
changeFieldType	The type of an existing field should be changed by the transformation.
changeMethodType	The return type of an existing method should be changed in the representation held by the transformation.
changeMethodModifier	The access modifier of a method should be modified in the representation maintained by the transformation.
changeFieldModifier	An existing field should have its access modifier updated by the transformation.
addMethodArgument	An existing method should be modified so that a new argument appears in its signature.
deleteMethodArgument	An existing method should have its signature augmented so that a specified parameter is deleted.
reorderMethodArgument	The signature should be rewritten so that the arguments to a method are reordered.
changeSuperclass	An existing class should have its superclass changed from the current class to the specified one.
moveFieldToSuperclass	A field should be moved by the transformation from a subclass up the inheritance hierarchy into a superclass.
moveField	The transformation should relocate a field from the class that it is currently defined in to a new class.
moveFieldToSubclass	A copy of an existing field should be moved by the model transformation from the class it is currently located in to each of the subclasses.
moveMethod	A method should be moved from one class to another.
moveMethodToSuperclass	A method should be moved from a class to its superclass.
moveMethodToSubclass	A method should be moved from a class to its subclass.

Table 4.2: The methods that must be implemented by the transformation interface.

4.4 Disadvantage of Our Approach

The powerful extensibility of our framework is also its biggest disadvantage. RefDsl scripts that will execute on one person's installation might not work on another installation. We propose in the future to provide a mechanism to load script handles and transformations dynamically so that the dependencies of a script can be bundled with the script.

4.5 Summary

In this chapter, we presented the architecture of the RefDstl framework and identified the ‘key’ components. We emphasised how users may extend this framework using two approaches. The first involves creating script handles that enrich the scripting language with additional functionality, for example, interacting with the user via the console. The second involved incorporating additional model transformations, which are the primary reason the framework is capable of working on a variety of heterogeneous artefacts.

This chapter in conjunction with Chapter 3 concludes our discussion on the implementation of the RefDstl system². In Chapter 5, we concentrate on evaluating the efficacy of the system as well as the accuracy of the model transformation used to refactor MoDisco models representing Java systems.

²By ‘system’, we mean the language and the framework.

Chapter 5

A Case Study on Automatically Refactoring ‘God’ Classes

Up to this point, we have provided the motivation for a DSL to script refactorings and to define new composite refactorings. We presented the syntax of the RefDsl language and described how the system can be extended with custom model representations and transformations.

Now we focus our attention on evaluating the efficacy of our approach and our prototype implementation. We do this by performing a case study on a collection of ‘God’ classes selected from a corpus of Java software. At a low level, the study shows that: 1. RefDsl can be used to remove the ‘God’ class ‘smell’ using the extract class refactoring, and, 2. RefDsl can be used to measure coupling between Java classes. From a higher level, however, it is shown that RefDsl is suitable in practical environments to refactor ‘real world’ sized software and in academic environments as an extensible framework to evaluate research hypotheses.

5.1 Purpose of Experiment

In order to demonstrate the efficacy of RefDsl, we aim to show that it has the following qualities:

- Expressibility — We show that RefDsl is capable of expressing refactorings on a scale on par with what would be required by professionals and academics alike.
- Efficiency — Although our current implementation is only a prototype, it must be capable of performing real world refactorings in a reasonable amount of time. We follow the direction of Simon et al. (2001) who state that a refactoring engine should be able to perform the refactoring quicker than an experienced engineer performing it manually. Quantitative data is collected during the study and presented here that shows the time taken to perform the refactorings.

These objectives imply the following criteria should be considered when choosing the refactoring to be used in this study.

- The refactoring should resolve a ‘naturally’ occurring symptom of ‘code rot’. This criterion helps to serve as evidence that RefDsl is useful in ‘real world’ situations.
- The refactoring should affect multiple different entities in the system, i.e. classes, methods and fields. It should also involve multiple refactoring primitives. This ensures that the range of refactoring instructions provided by RefDsl are broad enough to handle reasonably

complex refactorings. It also allows us to consider the time taken by RefDstl to perform a reasonably large refactoring.

With this criteria in mind, we examined the literature and discovered that as classes age they become less cohesive; they effectively become ‘God’ classes. The literature discusses approaches to address the issue. For example Simon et al. (2001) use distance metrics to present visualisations to software engineers as a guide of where to apply refactorings. However, we believe RefDstl provides better opportunities, which are the focus of this study. We use RefDstl to automatically remove the ‘God’ class ‘smell’ from a selection of classes and explore the impact on the CBO metric (Chidamber and Kemerer, 1994) as well as its predictor by Chaparro et al. (2014). We also examine the time it takes to perform the refactoring. We discuss how the study was conducted next.

5.2 Summary of Approach

We use the following sequential activities identified by Mens and Tourwe (2004) as the basis for our approach. These include:

1. Identifying refactoring opportunities — Mens and Tourwe (2004) assume that the refactoring is being conducted by a software engineer working on their own project. However, in this case study we also have the additional step of selecting systems to be the subjects of the experiment. Once identified, we then identify classes that exhibit the ‘God’ class ‘smell’ using an analysis tool that is developed in section 5.5.
2. Choosing a refactoring to apply — We intentionally look for ‘God’ classes in this study. Consequently, the refactoring to apply is predetermined to be the extract class refactoring.
3. Ensuring the program behaviour is preserved — It is the responsibility of the programmer writing the RefDstl refactoring script to ensure that that pre and post conditions for the refactoring are sufficient to preserve program behaviour.
4. Apply the refactoring — The RefDstl refactoring engine applies the extract class refactoring script to MoDisco models that represent the system under refactoring. These models contain the ‘God’ classes. We do not do this manually.
5. Assess the effect of refactoring on quality characteristics of software — We investigate the impact of extract class on the CBO metric. We examine if the *post* CBO measurement of the class under refactoring can be predicted.
6. Maintain consistency with other software artefacts — This does not form a part of our study. However, it should be noted for all other artefacts that are representable by a model that these can be manipulated by the RefDstl engine *during* refactoring.

We discuss the first of these activities next.

5.3 Identifying a Corpus for the Experiment

Ideally this study would be carried out on commercial software systems. However we note that more often than not commercial code is proprietary and industry is unwilling to share their code

to maintain their competitive advantage. In experiments where commercial code has been used, the researchers are normally subjected to non-disclosure agreements that oftentimes require they do not name the product or further disseminate the code, for example in the study by Barker and Tempero (2007). This poses a problem for researchers where the subject of the experiments should be made available so that the experiment can be repeated and that fair comparisons can be made later with other studies. As an alternative we use a corpus of open source software. We also choose to restrict ourselves to using Java source code because we only have access to a MoDisco knowledge discoverer for transforming Java into MoDisco models.

We have identified the following potential sources to use in the experiment. Blackburn et al. (2006) proposed the DeCapo benchmark suite, which is described as ‘set of open source, real world applications with non-trivial memory loads’. Although it is primarily a suite of previously calculated and verified benchmarks and measurements (such as object-oriented metrics) the software systems that the measurements were taken from can be downloaded using the Ant¹ script in the source-code distribution. However, we choose not to use this suite since it is an older corpus than the Qualitas Corpus, which we discuss next.

Tempero et al. (2010) provide the curated Qualitas Corpus as three distributions: The *r* distribution provides the most recent version of a large collection of software systems. The *e* distribution is the evolutionary version. It provides fewer systems but each of the systems is provided with multiple versions. The *f* distribution completes the corpus with systems and versions that appear in neither of the other two distributions. We select the *e* distribution for this experiment as it is suitably large and fits the requirements for our evolutionary study. The following section provides a summary of the contents of the corpus and describes its size. We note that there exists multiple releases of the corpus that come with later releases of software with substitutions made for systems that are no longer actively maintained. From this point on, we refer to version 20130901e when discussing the Qualitas Corpus.

The Qualitas Corpus has gained acceptance among other researchers such as the following. Fontana et al. (2012) perform a study on the affects of individual code ‘smells’ on different design metrics. Griffith et al. (2014) examine the relationship between technical debt estimation models and quality models. Al-Mutawa et al. (2014) examine whether all circular dependencies should be considered as ‘undesirable’. De Roover et al. (2013) augment the Qualitas Corpus in the form of QuaAtlas so that it is more suitable for studies concerning API usage. Arcelli et al. (2015) use the corpus to discover what refactoring tools are easiest and most useful for removing code ‘smells’.

5.4 Content of the Qualitas Corpus

We briefly remark on the contents of the Qualitas Corpus in order to provide a sense of the size of this study. The corpus is disseminated as a split tape archive file, which is approximately 16.75GB in its unpacked state and 66.41GB in its decompressed installed state. It contains fifteen different systems from different domains. These are listed in Table 5.1 along with the number of versions of each system included. In total there are 579 versions of software and 3,320,594 Java classes (which are distributed across the various systems as shown by Figure 5.1). A variety of systems offers a degree of certainty that projects have disjoint sets of developers working on them. If the

¹<http://ant.apache.org>

projects were developed by the same developers then the corpus would be unsuitable for certain types of studies such as those sensitive to developer habits. This includes the study presented here.




Figure 5.1: The number of classes per system.

The corpus contains the binaries and source code for each system as well as a summary of the corpus content and a summary of each project (i.e. metadata). Each project summary includes:

- A listing of the classes in the project and the location of each of those classes.
- A count of the number of lines of code in each source file and non-commented lines of code.
- An indicator determining whether each class is public or private and top level or nested.

Clearly not all of the classes provided with the corpus have the ‘God’ class ‘smell’ and some processing is required to identify the relevant classes. We develop an analysis tool for this purpose in the next section.

5.5 An Analysis Tool for the Qualitas Corpus

We develop a tool for analysing the content of the Qualitas Corpus. Its purpose is to perform two activities.

- It produces a relational model of the Qualitas Corpus metadata. This allows us to query the contents of the corpus easily and efficiently using SQL, rather than writing several ad hoc programs to process the files which is the default format for the corpus metadata. It also

System Name	No. of Versions Included
Ant	23
Antlr	22
ArgoUML	16
Azureues	63
Eclipse SDK	52
Freecol	32
Freemind	16
Hibernate	115
JGraph	39
JMeter	24
JStock	31
Jung	23
JUnit	24
Lucene	36
Weka	63

Table 5.1: Systems and number of versions provided in the Qualitas Corpus version 20130901e.

allows us to succinctly describe what subjects were used for the case study as an SQL query, rather than providing a long list of class and system version names.

- It performs an analysis of the classes in the system. The results of this analysis are persisted to the relational database and stored close to the metadata for later querying.

5.5.1 Persisting the Metadata

The Qualitas Corpus metadata is mapped using a text processing program written in Java to the simplified relational model shown in the entity–relationship diagram presented in Figure 5.2. This is implemented using the PostgreSQL² database system.

5.5.2 Analysing the Classes

Analysing software source code involves one of the following analysis approaches:

- Static — This approach involves analysing the code without executing it and is suitable in situations where either the compiled binary or just the source code are available.
- Dynamic — Dynamic analysis involves executing the program. Some binary analyses can be conducted without needing access to the source code, for example, analysing the time taken by a program to finish executing. However, other more involved analyses require that the code is instrumented before being executed. For example, suppose we want to generate a list of the methods that are called during the execution of a program with unknown inputs. Dynamic analyses can be slower because of the time that it takes to execute the program undergoing analysis.

We chose to incorporate static analysis into the tool. At the core of the analysis tool is an iterator that goes through each of the classes provided in the corpus. For efficiency reasons, it

²<http://www.postgresql.org>

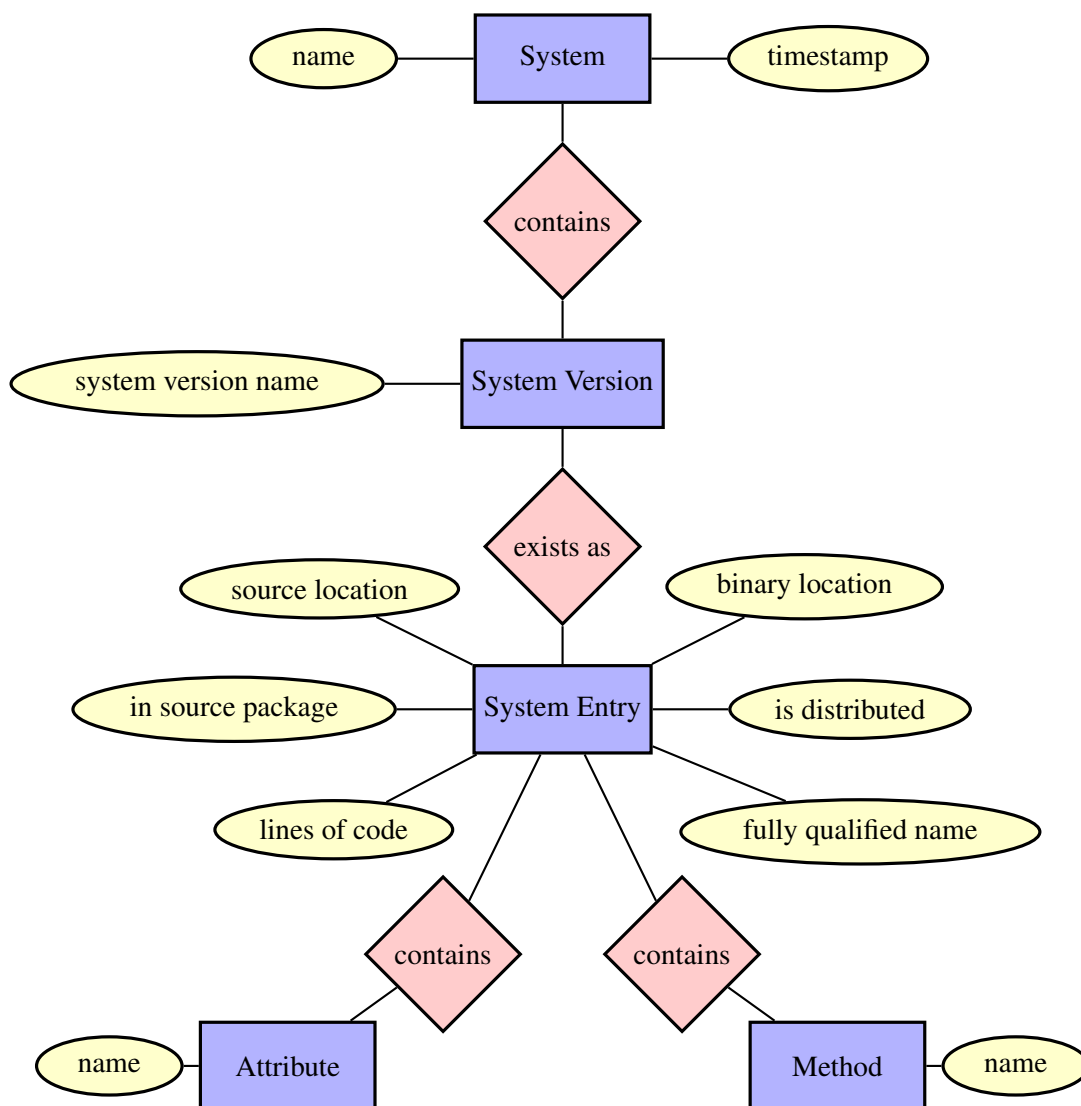


Figure 5.2: An entity relationship diagram illustrating the main entities and attributes persisted by the static analysis tool. Other entities, attributes and keys have been omitted for simplicity.

is designed to iterate over the classes stored in a Jar file sequentially. This has been done for efficiency because it means that a Java archive file only needs to be loaded and searched for Java class files once. While iterating over the classes, the actual analysis occurs. This requires using a library to analyse the classes. Two libraries were considered for this purpose: ASM (Bruneton et al., 2002) and BCEL³.

ASM works on the level of source code. It follows the visitor design pattern so developers define actions to occur when nodes of a particular type are encountered. This is similar to the SAX streaming API for parsing XML. However, it also offers a less memory efficient but more user friendly DOM style interface that allows the developer to ‘pull’ what they want from a tree structure of the program.

We opted however to use BCEL. It works at the lower byte-code level (although the byte-code does not need to be executed). BCEL of course requires that the program being analysed can be compiled. However, as mentioned earlier, the corpus provides the system binaries so this was not a problem. BCEL has an advantage over ASM in that some analyses are simpler. For example, with BCEL, we can search for patterns of bytecode so we can easily find all field accesses or mutations in a class.

Our static analysis computes class level metric measurements (TCC, WMC, ATFD) and saves the results to the database.

One might wonder why we developed a separate analysis tool for calculating these metrics rather than employing the RefDsl system. We could have used RefDsl for this purpose. However, initial experimentation showed that RefDsl would have been too slow to calculate the metrics for the three million classes in the corpus.

5.6 Selecting Refactoring Candidates

Our goal is to select the ‘God’ classes from the Qualitas corpus. However, we will choose only classes that appear in every version of a system distributed with the corpus and are ‘God’ classes throughout. This provides a degree of certainty that what we select are actually ‘God’ classes and not ‘borderline’ that could be reclassified due to a minor change occurring in some version.

To assess which of classes are ‘God’ classes, we use the criteria described by Lanza et al. (2005). They define a strategy for finding instances of the ‘God’ class ‘disharmony’ using these requirements:

- access to foreign data (ATFD) (Lanza et al., 2005) metric is greater than a few, and,
- weighted method per class (WMC) (Chidamber and Kemerer, 1994) metric is very high, and,
- tight class cohesion (TCC) (Bieman and Kang, 1995) metric is less than a third.

For the values ‘greater than a few’, we use the value 1, for very high we use 25%. This is in line with the values given in Marinescu (2004). The query to identify ‘God’ classes that appear in every version of a system is presented in Figure 5.3.

³<http://commons.apache.org/proper/commons-bcel/>

```

1 CREATE VIEW god_classes_all_versions AS
2 SELECT T.*, system_entry.*
3 FROM ( SELECT se.system_name, se.fully_qualified_name
4        FROM system_entry se
5        WHERE se.tcc < 0.33
6              AND se.atfd > 1
7              AND se.wmc > 25
8              AND se.system_name <> 'eclipse'
9        GROUP BY se.system_name, se.fully_qualified_name
10       HAVING count(se.fully_qualified_name) = ( SELECT count(sv.*)
11                                                FROM system_version sv
12                                                WHERE sv.system_name
13                                                  = se.system_name
14                                                GROUP BY sv.system_name )
15 ) T
16 JOIN system_entry USING (system_name, fully_qualified_name);

```

Figure 5.3: SQL query to discover ‘God’ classes in all versions.

```

1 SELECT *
2 FROM god_classes_all_versions
3 WHERE system_version_version IN ( SELECT *
4                                  FROM max_version )
5   AND system_name IN ( 'ant', 'argouml', 'freemind', 'jgraph', 'jmeter',
6                       'jstock', 'jung', 'junit', 'lucene', 'weka' )
7   AND system_version_version NOT IN ( 'argouml-0.16.1', 'argouml-0.18.1',
8                                       'freemind-0.6.7', 'weka-3.4',
9                                       'weka-3.4.12' )
10  AND level = 'Top Level Public'
11

```

Figure 5.4: SQL query to discover classes for the study.

This dataset had to be later narrowed down to the dataset given by the query in Figure 5.4. This was due to either the systems being too large to analyse (Eclipse) or being unable to generate models for those projects (which is discussed later).

5.6.1 Creating the MoDisco Models

A two stage process was involved to produce the MoDisco models for the projects in the Qualitas corpus:

1. We wrote a Java program that converted each system version in the corpus into a valid Eclipse project with the correct classpath settings. The classpath settings were created using the data pulled from the analysis tool described earlier.
2. With each system version in the form of an Eclipse project, we could produce MoDisco models. The MoDisco plugin for Eclipse comes with knowledge discoverers to do this. However, out of the box, these are designed to be interacted with manually. We developed a plugin for Eclipse that performed the knowledge discovery and saved the MoDisco models without human interaction. In total, it takes longer than six hours for the plugin to convert every project (ignoring the Eclipse SDK). The time spent developing the plugin paid for itself considering we had to generate the models three times before we were left with a




Figure 5.5: The percentage of ‘God’ classes for all versions of each of the software systems in the corpus.

reasonable number of models. We did not produce the models correctly the first two times because the classpath settings for the projects were initially incorrect. On our third attempt we managed to produce 458 valid models out of 579 systems. The large majority of the systems we could not produce models for were the hibernate system. We therefore excluded all hibernate systems from our study along with certain versions of the other systems. The dataset that was actually used for the study is given by the query in Figure 5.4.

In the two following sections, we discuss our protocol for deciding how to split an incohesive class as well as the RefDstl extract class script for performing the extraction.

5.7 An Extract Class Refactoring

Applying an extract class refactoring to an incohesive class involves moving elements of the incohesive class to a newly created class. This is accomplished using a series of move field and method refactorings as well as create class refactorings. However, deciding what fields and methods to move automatically involves using algorithms such as those used in social network analysis like the Girvan–Newman algorithm (Cassell et al., 2009; Girvan and Newman, 2002) or from the area of machine learning using algorithms such as k -means. The algorithm we used is as follows:

1. For every method in the class that is being automatically refactored, we create a cluster to contain that method.
2. We iteratively check each cluster against the other clusters to determine if the clusters should be merged. We say that two clusters should be merged if greater than half of the methods in

Method Name	Parameters	Description
cluster	Class name	This method will cluster the class with the fully qualified name specified in the input parameters. The method always returns true so it can be used safely within a pre or post condition expression.
fieldToClass	A field.	Returns the name of the class where the field should be moved.
methodToClass	A method.	Provides the name of the class to which the method should be moved.

Table 5.2: API for interacting with the clustering plugin.

the two clusters are similar. Two methods are considered similar if at least half of their field accesses are to the same fields.

3. After the clusters in the previous step converge, we then turn to allocating the fields to the clusters. Each field is assigned to the cluster that uses that field the most.

We implemented this algorithm in a script handle for the RefDstl framework since the RefDstl language has no way of expressing the algorithm otherwise. The API for using the handle is shown in Table 5.2.

A RefDstl script was created automatically for each of the classes to be refactored and identified in our dataset. An example of one of the scripts is shown in Figure 5.6. It shows the `main` procedure calling another to perform the extract class refactoring. This delegation is necessary since `main` does not support pre or post conditions.

The precondition is used to arrange the class elements into cohesive units: this is done using the call to the `clusterClass` method.

The first postcondition is used to calculate the CBO metric for the original class post refactoring. The formula for calculating the CBO value for a class is given in Equation 5.1. In words, it is a count of the number of classes a class c is coupled with. The last postcondition calculates the predicted value for CBO. This is done using the CBO predictor function by Chaparro et al. (2014), which we modified to take into account multiple extract class refactorings being performed at once. The formula is shown in Equation 5.2. It says that the CBO value for a class that has undergone an extract class refactoring should be equal to its original value plus n which is a count of the number of classes that were extracted less the number d which is a count of the classes coupled to by the extracted methods but not used by the original class. In the original formula proposed by Chaparro et al. (2014), n has the constant value of one.

$$CBO(c) = |c_{coupled}| \quad (5.1)$$

$$CBO_p(c_s) = CBO_b(c_s) + n - d \quad (5.2)$$

The results of applying the clustering algorithm to the classes that were selected for refactoring are discussed next, which includes comparing the actual CBO measurements to their predicted counterparts.


```

1 in 'model_java2kdm.xmi'
2 out 'model_out.xmi'
3
4 main()
5 {
6   call measureAndSubvert('org.apache.tools.ant.DirectoryScanner')
7 }
8
9 proc measureAndSubvert($originalGodClass)
10 requires $clusterer.clusterClass($utils.findClass($originalGodClass))
11 ensures $debug.printToFile('/Users/keith/Desktop/metrics.txt', 'ant', 'ant
    -1.8.4', 'org.apache.tools.ant.DirectoryScanner', $ripe.cbo($utils.
    findClass('org.apache.tools.ant.DirectoryScanner')))
12 ensures $debug.printToFile('/Users/keith/Desktop/metrics_pred.txt', 'org.
    apache.tools.ant.DirectoryScanner', 'ant', 'ant-1.8.4', $ripe.cboPredicted(
    $utils.oldClass('org.apache.tools.ant.DirectoryScanner'), $utils.findClass
    ('org.apache.tools.ant.DirectoryScanner')))
13 {
14   for($extractedClassName in $clusterer.classes()) {
15     if(~$extractedClassName.equals($originalGodClass) {
16       call createClass($extractedClassName)
17     }
18   }
19
20   for($method in $utils.findClass($originalGodClass).getMethods()) {
21     if(~$originalGodClass.equals($clusterer.methodToClass($method)) {
22       call makeMethodPublic($method.getName(), $method.getParamTypes(),
23         $originalGodClass)
24       call moveMethod($method.getName(), $method.getParamTypes(),
25         $originalGodClass, $clusterer.methodToClass($method))
26     }
27   }
28
29   for($field in $utils.findClass($originalGodClass).getFields()) {
30     if(~$originalGodClass.equals($clusterer.fieldToClass($field)) {
31       call makeFieldPublic($field.getName(), $originalGodClass)
32       call moveField($field.getName(), $originalGodClass, $clusterer.fieldToClass(
33         $field))
34     }
35   }
36
37   for($extractedClassName in $clusterer.classes()) {
38     if(~$extractedClassName.equals($originalGodClass) {
39       call createField($extractedClassName.toLowerCase(), 'private',
40         $originalGodClass, $extractedClassName)
41     }
42   }
43 }

```

Figure 5.6: RefDsl sourcecode used to perform the extract class refactoring on an identified 'God' class.

5.8 Results

We discuss the results of the case study from two angles: the correlation between the actual CBO of the measured class post refactoring and the correlation between the the time taken to perform the refactoring and the size of the system being refactored.

5.8.1 Predicting CBO

Figure 5.7 summarises the results of the case study. The x -axis represents the predicted value of the CBO metric while the actual value is shown along the y -axis. The diagonal blue line shows a linear model produced from the data. Ideally, there would exist a perfect correlation between the actual measurement and the predicted measurement. This would result in a slope of 1, i.e. when the x value increases by a single unit then the y value increases by an equal amount. The slope presented here is 0.6843. There are two reasons why we might have fallen short of 1. The first reason are potential flaws in our implementation. For example, the refactoring transformations that were implemented in the interpreter could have coding defects due to human error that is typical when developing any software system. Secondly, the predictor function is inaccurate. We note that in their evaluation, Chaparro et al. (2014) tested this particular CBO predictor function by having two postgraduate students perform between five and ten refactorings on two software systems. When the predicted and actual CBO measurements were compared for the two systems just 63% of the measurements were completely accurate.

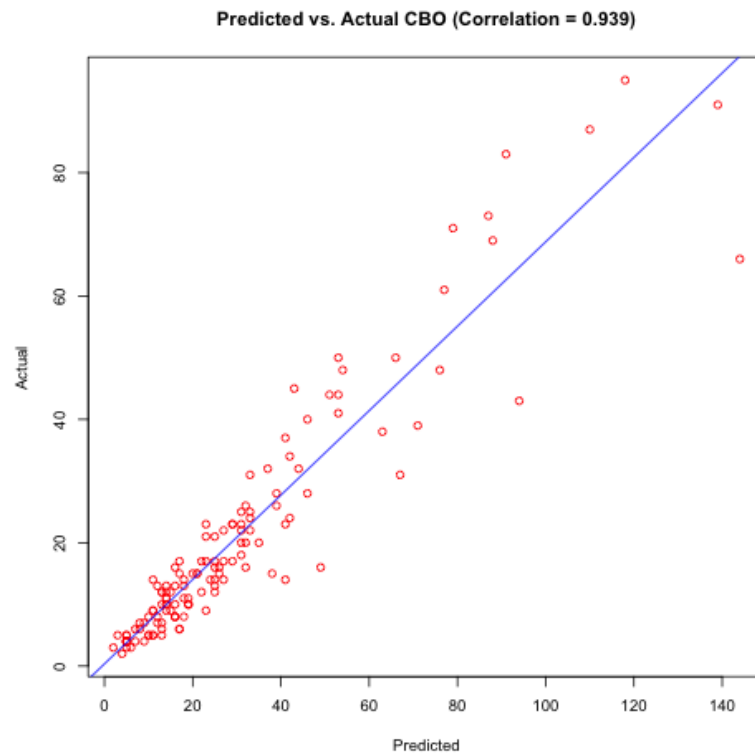


Figure 5.7: Pearson's correlation between the predicted CBO metric and the actual value .

We can further measure how close the predicted value correlates with the actual value using

Pearson's correlation coefficient, which is given by the `cor` function in the R language⁴ for statistical computing. In our results, the coefficient is 0.939. A perfect correlation would yield a value of 1 but nonetheless this value shows high correlation.

We note that there are some clear outliers in this data. The most apparent is due to the class `'org.yccheok.jstock.gui.MainFrame'` in `'jstock-1.0.7c'`. Its actual value after refactoring is 66 but its predicted value is 144. Further investigation, which involved reading the source code of that class, would suggest that we are not handling the case correctly where class attributes are accessed using the `this` Java keyword.

We remark on the distribution of the CBO measurements as shown by the three 'box-and-whisker' plots in Figure 5.8. In each plot, the 'whiskers' at the top and bottom of the plot represent the range of values (discarding the outliers shown as red coloured points in the figure). The thick black band inside each of the coloured boxes represents the median for that dataset while the beginning and end of each box represents the first and third quartiles for each dataset respectively. A perfectly distributed dataset would see the median located half way between the 'whiskers' while the first and third quartiles would be at the 25% and 75% marks. What's interesting, however, is how the plots compare to each other. We can see that the predicted values are generally higher than the original values. The outliers have higher values and the first quartile is at a higher value. We would expect this to be the case because for each class that undergoes an extract class refactoring it gains a newly created field that references the extracted class. In contrast, the actual values tend to be lower than the predicted values. The actual outliers have lower values and the quartiles are also lower. We have no rationale to explain this. However, it suggests the presence of a 'bug' in our implementation: most likely we are not considering some class usages. These would increase the value of d in Equation 5.2 and hence reduce the predicted values.

We do not compare our results with those from the study by Chaparro et al. (2014). Doing this would be meaningless for the following reasons:

- We did not use the same dataset. Those authors restrict themselves to ArgoUML⁵ and aTunes⁶ (although they do not specify what version of each). On the other hand we use classes from a wider range of projects, including ArgoUML.
- We did not use the same measurement tool.
- We measured different representations of the systems (for example, the transformation for turning Java code into MoDisco models may have some bugs.)

5.8.2 Time Taken

As shown in Figure 5.9, there also exists a strong correlation (0.981) between the time taken to perform the refactoring and the size of the model. This supports the idea that most of the time spent by the refactoring engine goes toward loading the model to be refactored and writing it back to disk after refactoring. There is some small variation in the time taken to refactor models of different sizes. However, the greatest variation is around seven seconds. We consider this to be

⁴<http://www.r-project.org>

⁵<http://argouml.tigris.org>

⁶<http://www.atunes.org>

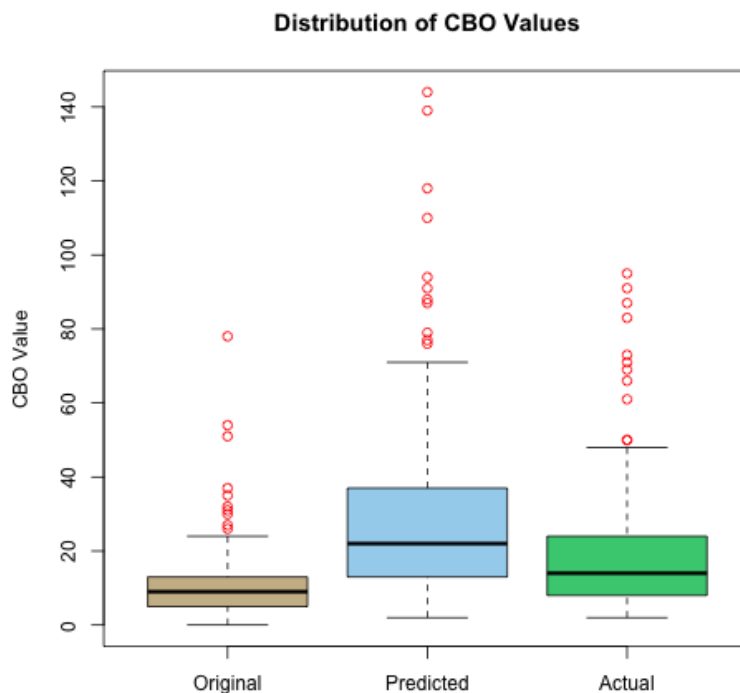


Figure 5.8: The distribution of the original, predicted post refactoring and actual post refactoring CBO measurements for the software systems used in the study.

negligible and bear in mind that these timings reflect ‘wall time’ and not ‘CPU time’ so slight variations might be caused by the operating system switching to background processes during the refactoring etc.

Overall the time taken to perform the refactorings was sixteen minutes. We do not believe that a developer could perform these 134 refactorings in this amount of time: deciding how to refactor just one ‘God’ class would probably take a developer as long.

We note that in this study we used Java code to calculate the CBO metric and its predicted counterpart. We had tried to use the metrics transformation also provided with RefDstl but we found that this was too slow (approximately six minutes to do one refactoring). The slowdown is due to having to transform every class to the metrics metamodel rather than just the needed classes. Future versions of RefDstl will address this weakness.

5.9 Discussion

We believe our case study shows that RefDstl is a suitable system for targeting the ‘four Is’ that were identified in Chapter 1:

1. The correlation between CBO and the predicted value for CBO suggests that our implementation has some degree of accuracy. If there were no correlation then this would suggest that either the prediction model is wrong or the implementation is flawed. However, given the results it would require a coincidence to get this correlation. A refined implementation should yield better results still.

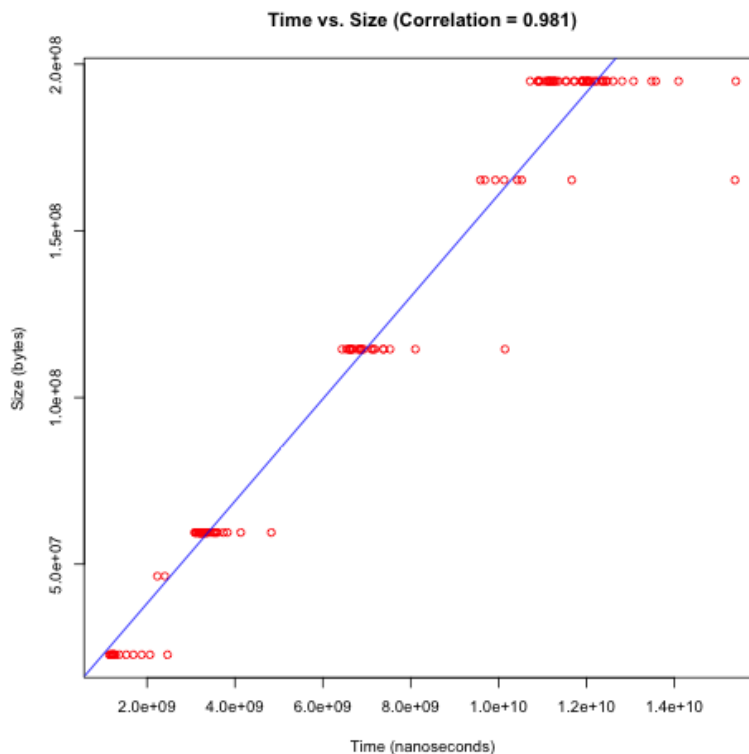


Figure 5.9: Pearson’s correlation between the size of the model being refactored and the time taken to refactor the model in nanoseconds.

2. We used RefDstl to define a new composite refactoring that can be reused and shared in a library. This demonstrates that our approach affords flexibility to refactoring unlike many mainstream refactoring engines.
3. Our framework is extensible since it considered CBO metrics during the refactoring and can be adapted to consider other metrics or representations.
4. RefDstl had the foresight to know which fields and methods to move. This is dissimilar to other ‘blind’ refactoring engines that merely follow the (mis–)guidance of developers who could make incorrect decisions.

5.10 Summary

In this chapter, we have presented a case study of refactoring 134 ‘God’ classes from the Qualitas Corpus. We refactored the classes automatically using an algorithm that we embedded into the RefDstl system. We evaluated our approach by comparing the CBO measurements post refactoring with predicted values. The high correspondence between the values provided some guarantee (although not absolute) about the accuracy of our refactoring system and the time measurements show the feasibility of our approach when considering time.

In the final chapter that follows we conclude this dissertation with a reflection of what has been achieved and how. We also remark on how this work might be expanded in the future.

Chapter 6

Conclusion

In this chapter, we highlight the contributions that have been made in this dissertation and we discuss topics that merit exploration in future work.

6.1 Summary of Contributions

Refactoring is the methodical approach to restructuring a software system such that its observable behaviour remains unchanged (Opdyke, 1992). In Chapter 1 and Chapter 2, we recognised the importance of refactoring engines in the refactoring process. However, we highlighted that these are susceptible to four recurrent weaknesses that we named the ‘four Is’.

I1 referred to inaccuracy, which is when the conditions required for a refactoring to be applied are defined too strongly or weakly by the refactoring engine. In Chapter 2, we saw that attempts have been made to *formally* demonstrate that refactorings preserve behaviour for OO *specification* languages and for a small set of Java refactorings. However, the approaches taken by those researchers can only demonstrate that conditions are not weak. They cannot demonstrate that conditions are overly strong. We proposed that, until this issue is resolved rigorously, that it would be better if refactoring engines placed their conditions *away* from the transformation code so that conditions could be altered in ‘the field’. For this reason, we developed the RefDstl language in Chapter 3. It can be used to specify new composite refactorings (using a DBC style) from provided primitive operations. The language also addresses I2 because the defined refactorings can be saved to scripts for later application.

With regard to I3, we implemented a framework for interpreting RefDstl scripts. It has been carefully designed so that additional transformations can be easily integrated. Each additional transformation is capable of refactoring a different type of artefact (although all artefacts should be different representations of the same model) and we provided two transformations for demonstration purposes: one to actually refactor models of Java systems and one for transforming the system under refactoring into a measurement metamodel developed by McQuillan (2011). By permitting additional transformations and allowing extension points called ‘script handles’, we allow RefDstl programmers to reason about different aspects or views of the system under refactoring in their refactoring pre and post conditions. This means that the RefDstl framework is not improvement-like traditional refactoring engines (I4). It allows users to reason about the quality of the software during refactoring (which we demonstrated at the end of Chapter 3).

We evaluated the efficacy and, to a degree, the accuracy of our prototype implementation in Chapter 5. We showed how RefDstl could be used to automatically improve the structure of

‘God’ classes using a script handle that calculates the CBO value for the class under refactoring. The accuracy of our transformations was shown by checking the correlation between the actual CBO values post refactoring and the predicted values using a predictor function given by Chaparro et al. (2014). We explained that this was not a ‘water tight’ approach but a sensible heuristic. The predictor function could be incorrect and thus we would have achieved a low correlation anyway. However, it would require coincidence for either the predictor function or our transformations to be incorrect and to still achieve the high correlation that we did (0.939 using Pearson’s correlation coefficient).

Also, as a starting point for those eager to use RefDstl, we have contributed a standard library of primitive refactoring operations in Appendix B. It was used in the examples throughout this dissertation.

Finally, we also want to highlight that our approach to developing the RefDstl language equips us with a reusable metamodel of the primitive refactorings described by Opdyke (1992). This can be used by researchers in areas beyond what has been described in this dissertation.

Now, having summarised our model driven approach for refactoring heterogeneous software artefacts, we briefly remark on avenues available for future work.

6.2 Future Work

In future work, we will take a different approach to tackle the problem of showing that refactoring preserves a program’s behaviour. For Java programs equipped with JML specifications, we will integrate a transformation for JML annotations into the RefDstl framework. This will mean that when JML annotated programs are refactored that their specifications will also be refactored in tandem. By using theorem provers, we could then show that the specification of a program before refactoring is equivalent to the specification after refactoring.

For programs that are not annotated with specifications, we will offer an approach that allows users to reason about the structural properties of the software. We will develop the MoDisco metamodel in the Gallina language used by the Coq theorem prover¹ as a collection of structures². We will integrate a transformation into the RefDstl framework that converts the system being refactored into instances of those structures (effectively a MoDisco to Gallina transformation). This would allow users to reason about the structure of their systems; for example, they could prove that all references to a renamed entity are preserved after refactoring.

Finally, we will enrich the interpreter with the capability to calculate the weakest precondition for any RefDstl procedure. This can be done using Dijkstra’s Weakest Precondition Calculus but it would require that loop invariants are added to the RefDstl language.

¹<https://coq.inria.fr>

²Similar to structs in ‘C’ or classes in Java.

Acronyms

ASCII American Standard Code for Information Interchange.

AST Abstract Syntax Tree.

ATFD Access to Foreign Data.

ATL Atlas Transformation Language.

BNF Backus Naur Form.

CASE Computer Aided Software Engineering.

CBO Coupling Between Objects.

DBC Design by Contract.

DSL Domain Specific Language.

DSTL Domain Specific Transformation Language.

EMF Eclipse Modelling Framework.

IDE Integrated Development Environment.

JML Java Modelling Language.

LCC Loose Class Cohesion.

LCOM Lack of Cohesion of Methods.

MDE Model Driven Engineering.

MOF Meta Object Facility.

MPS Meta Programming System.

OMG Object Management Group.

OO Object-Oriented.

QVTO Query View Transformation Operational.

QVTR Query View Transformation Relational.

SQL Structured Query Language.

TCC Tight Class Cohesion.

TIKZ Tikz ist kein Zeichenprogramm.

UML Unified Modelling Language.

WMC Weighted Methods Per Class.

XML Extensible Markup Language.

XSD XML Schema Definition.

Appendix A

RefDsl Language Grammar for Xtext

```
1 grammar ie.pop.refdsl.RefDsl
2
3 hidden(WS, ML_COMMENT, SL_COMMENT)
4
5 import "http://www.eclipse.org/emf/2002/Ecore" as ecore
6
7 generate refDsl "http://www.pop.ie/refdsl/RefDsl"
8
9 RefactoringScript:
10   requires += Import*
11   ('in' project = STRING)?
12   ('out' output = STRING)?
13   (main = Main)?
14   refactorings += Procedure*;
15
16 Import:
17   'imports' file = STRING;
18
19 /* Procedures. */
20 Main:
21   'main' '(' ')'
22   '{'
23     instructions += (RefactoringInstruction | ConditionalStatement |
24     ProcedureCall | LoopingStatement)+
25   '}' ;
26 Procedure:
27   'proc' name=ID_PART '(' (arguments=ParameterList)? ')'
28   ('requires' preconditions+=Condition1)*
29   ('ensures' postconditions+=Condition1)*
30   '{'
31     instructions += (RefactoringInstruction | ConditionalStatement |
32     ProcedureCall | LoopingStatement)+
33   '}' ;
34 ProcedureCall:
35   'call' name=ID_PART '(' args = ArgumentList? ')';
36
37 ConditionalStatement:
```

```

38  'if' '(' condition=Condition1 ')' '{'
39  instructions +=(RefactoringInstruction | ConditionalStatement | ProcedureCall
40  | LoopingStatement)+
41  '}' ;
42 LoopingStatement :
43  'for' '(' id=VID 'in' (collection=ObjectExpression) ')' '{'
44  instructions += (RefactoringInstruction | ConditionalStatement |
45  ProcedureCall | LoopingStatement)+ '}' ;
46 RefactoringInstruction : { RefactoringInstruction }
47  action=(CreateClass
48  | CreateField
49  | CreateMethod
50  | DeleteClass
51  | DeleteField
52  | DeleteMethod
53  | ChangeClassName
54  | ChangeFieldName
55  | ChangeMethodName
56  | ChangeFieldType
57  | ChangeMethodType
58  | ChangeMethodModifier
59  | ChangeFieldModifier
60  | AddMethodArgument
61  | DeleteMethodArgument
62  | ReorderMethodArguments
63  | ChangeSuperclass
64  | MoveFieldToSuperclass
65  | MoveFieldToSubclass
66  | MoveMethodToSuperclass
67  | MoveMethodToSubclass
68  | MoveField
69  | MoveMethod
70  | NoOp);
71
72 /* Create syntax. */
73 CreateClass :
74  'create' (accessor=Accessor)? 'class' name=(ID_PART | FULL_ID | VID) ('with'
75  'superclass' superName=(ID_PART | FULL_ID | VID))?;
76 CreateField :
77  'create' (accessor=Accessor)? (static='static')? 'field' name=(ID_PART | VID)
78  'in' 'class' unit=(ID_PART | FULL_ID | VID) 'of' 'type' type=(FULL_ID |
79  ID_PART | VID);
80 CreateMethod :
81  'create' (accessor=Accessor)? (static='static')? 'method' name=(ID_PART | VID)
82  'in' 'class' unit=(ID_PART | FULL_ID | VID) ('taking' '[' args=
  JavaParameterList ']')? ('returning' return=(FULL_ID | ID_PART | VID))?;

```

```

83 DeleteClass :
84   'delete' 'class' name=(FULL_ID | ID_PART | VID);
85
86 DeleteField :
87   'delete' 'field' name=(ID_PART | VID) 'in' 'class' unit=(FULL_ID | VID |
88     ID_PART);
89
90 DeleteMethod :
91   'delete' 'method' name=(ID_PART | VID) 'in' 'class' unit=(FULL_ID | VID |
92     ID_PART);
93
94 /* Change syntax. */
95 ChangeClassName :
96   'rename' 'class' oldName=(FULL_ID | VID | ID_PART) 'to' newName=(FULL_ID |
97     VID | ID_PART);
98
99 ChangeFieldName :
100  'rename' 'field' oldName=(ID_PART | VID) 'to' newName=(ID_PART | VID) 'in' '
101    class' unit=(FULL_ID | VID | ID_PART);
102
103 ChangeMethodName :
104  'rename' 'method' oldName=(ID_PART | VID) 'to' newName=(ID_PART | VID) 'in' '
105    class' unit=(FULL_ID | ID_PART | VID);
106
107 ChangeFieldType :
108  'change' 'type' 'of' 'field' name=(FULL_ID | ID_PART | VID) 'in' 'class' unit
109    =(FULL_ID | ID_PART | VID) 'to' newType=(FULL_ID | VID | ID_PART);
110
111 ChangeMethodType :
112  'change' 'type' 'of' 'method' name=(ID_PART | VID) 'in' 'class' unit=(FULL_ID
113    | ID_PART | VID) 'to' newType=(FULL_ID | VID | ID_PART);
114
115 ChangeMethodModifier :
116  'make' 'method' name=(ID_PART | VID) ('with' 'args' '[' args=
117    JavaParameterList ']')? 'in' 'class' unit=(FULL_ID | VID | ID_PART)
118    newModifier=Accessor;
119
120 ChangeFieldModifier :
121  'make' 'field' name=(ID_PART | VID) 'in' 'class' unit=(FULL_ID | ID_PART |
122    VID) newModifier=Accessor;
123
124 AddMethodArgument :
125  'add' 'argument' arg=JavaParameter 'to' 'method' name=(ID_PART | VID) ('with'
126    'args' '[' args=JavaParameterList ']')?
127    'in' 'class' unit=(FULL_ID | VID | ID_PART);
128
129 DeleteMethodArgument :
130  'delete' 'argument' arg=(ID_PART | VID) 'from' 'method' name=(ID_PART | VID)
131    ('with' 'args' '[' args=JavaParameterList ']')?
132    'in' 'class' unit=(FULL_ID | VID | ID_PART);
133
134 ReorderMethodArguments :

```

```

124 'place' 'argument' arg=(ID_PART | VID) 'in' 'method' name=(ID_PART | VID) '
    with' 'args' '[' args=JavaParameterList ']' 'before' successor=(ID_PART |
    VID) 'in' 'class' unit=(FULL_ID | VID | ID_PART);
125
126 ChangeSuperclass:
127 'class' unit=(FULL_ID | ID_PART | VID) 'extends' newSuperclass=(FULL_ID | VID
    | ID_PART);
128
129 /* Move syntax. */
130 MoveFieldToSuperclass:
131 'move' 'field' name=(ID_PART | VID) 'from' unit=(FULL_ID | VID | ID_PART) 'up
    ';
132
133 MoveFieldToSubclass:
134 'move' 'field' name=(ID_PART | VID) 'from' unit=(FULL_ID | VID | ID_PART) '
    down';
135
136 MoveMethodToSuperclass:
137 'move' 'method' name=(ID_PART | VID) ('with' 'args' '[' args=
    JavaParameterList '])'? 'from' unit=(FULL_ID | VID | ID_PART) 'up';
138
139 MoveMethodToSubclass:
140 'move' 'method' name=(ID_PART | VID) ('with' 'args' '[' args=
    JavaParameterList '])'? 'from' unit=(FULL_ID | VID | ID_PART) 'down';
141
142 MoveField:
143 'move' 'field' name=(ID_PART | VID) 'from' unit=(FULL_ID | VID | ID_PART) 'to
    ' new_unit=(FULL_ID | VID | ID_PART)
144 ;
145
146 MoveMethod:
147 'move' 'method' name=(ID_PART | VID) 'from' unit=(FULL_ID | VID | ID_PART) '
    to' new_unit=(FULL_ID | VID | ID_PART)
148 ;
149
150 NoOp:
151 name = 'nop'
152 ;
153
154 /* Helpers. */
155 Accessor:
156 value=('public' | 'private' | 'protected' | 'default');
157
158 /*
159 * Contracts.
160 */
161 Condition1:
162 left=Condition2 operator='->' right=Condition1
163 | left=Condition2 operator='==' right=Condition1
164 | left=Condition2 operator='~=' right=Condition1
165 | left=Condition2 operator='>' right=Condition1
166 | left=Condition2 operator='<' right=Condition1

```

```

167 | left=Condition2 operator='>=' right=Condition1
168 | left=Condition2 operator='<=' right=Condition1
169 | higher=Condition2;
170
171 Condition2:
172     operator='forall' id=VID 'in' collection=ObjectExpression ':' left=
    Condition1
173 | operator='exists' id=VID 'in' collection=ObjectExpression ':' left=
    Condition1
174 | higher=Condition3;
175
176 Condition3:
177     left=Condition4 operator='/\\"' right=Condition3
178 | left=Condition4 operator='\\/\"' right=Condition3
179 | higher=Condition4;
180
181 Condition4:
182     operator='~' left=Condition4
183 | higher=Condition5;
184
185 Condition5:
186     '(' bracketed=Condition1 ')'
187 | arg=ObjectExpression
188 | constant=LogicalConstant
189 | scalar=STRING;
190
191 /*
192  * Parameters.
193  */
194
195 JavaParameter:
196     type = (FULL_ID | ID_PART | VID) name = (FULL_ID | ID_PART | VID)
197 ;
198
199 JavaParameterList:
200     parameters += JavaParameter (',' parameters += JavaParameter)*
201 ;
202
203 Parameter:
204     name = VID
205 ;
206
207 ParameterList:
208     parameters += Parameter (',' parameters += Parameter)*
209 ;
210
211 /*
212  * Arguments.
213  */
214
215 Argument:

```

```

216   constant = ID_PART | variable = VID | logical = LogicalConstant | expression
      = ObjectExpression | string = STRING
217 ;
218
219 ArgumentList:
220   arguments += Argument ( ',' arguments += Argument)*
221 ;
222
223 ObjectExpression:
224   base = VID (calls += ObjectExpressionMethodInvocation)+
225 ;
226
227 ObjectExpressionMethodInvocation:
228   '.' methodName = ID_PART '(' (arguments = ArgumentList)? ')'
229 ;
230
231 LogicalConstant:
232   value=("TRUE" | "FALSE");
233
234 /* Terminal definitions. */
235
236 terminal VID returns ecore::EString : '$' ID_PART;
237
238 //terminal ID returns ecore::EString : '~?(?a..'z'|'A..'Z'|'_') (?a..'z'|'
      A..'Z'|'_'|'0'..'9')*((('.') (?a..'z'|'A..'Z'|'_') (?a..'z'|'A..'Z'|'_
      '|'0'..'9')*)*);
239
240 terminal FULL_ID returns ecore::EString : ID_PART+ ('.' ID_PART*)+;
241
242 terminal ID_PART returns ecore::EString : '^?(?a..'z'|'A..'Z'|'_') (?a..'z'
      |'A..'Z'|'_'|'0'..'9')*);
243
244 //terminal INT returns ecore::EInt: ('0'..'9')+;
245
246 terminal STRING :
247   '"' ( '\\ ' . /* 'b'|'t'|'n'|'f'|'r'|'u'|'"'|'\"'|'\\' */ | !('\\'|'"') )*
      '"' |
248   "'" ( '\\ ' . /* 'b'|'t'|'n'|'f'|'r'|'u'|'"'|'\"'|'\\' */ | !('\\'|'"') )*
      "'";
249 ;
250
251 terminal ML_COMMENT : '/*' -> '*/';
252
253 terminal SL_COMMENT : '// ' !('\n'|'\r')* ('\r'? '\n')?;
254
255 terminal WS : (' '|'\t'|'\r'|'\n')+;
256
257 terminal ANY_OTHER: .;

```

Listing A.1: Grammar definition for RefDsl

Appendix B

RefDstl Standard Library

```
1 /**
2  * This file is the standard library of refactorings for the RefDstl system
3  * when being used to refactor Java programs.
4  *
5  * Essentially these are CRUD operations over a model of Java programs.
6  */
7
8 /*
9  * Create operations.
10 */
11
12 /**
13  * Creates a new class with no members and no superclass.
14  *
15  * $className - The name of the class to create.
16  */
17 proc createClass($className)
18   requires $utils.classDoesNotExist($className)
19 {
20   create class $className
21 }
22
23 /**
24  * Creates a new class with no members but with a given superclass.
25  *
26  * $className - The name of the class to create.
27  * $superClassName - The name of the class' superclass.
28  */
29 proc createClassWithSuper($className, $superClassName)
30   requires forall $cls in $utils.allClasses()
31     : ~$cls.getFullyQualifiedName().equals($className)
32   requires exists $cls in $utils.allClasses()
33     : $cls.getFullyQualifiedName().equals($superClassName)
34 {
35   create class $className with superclass $superClassName
36 }
37
38 /**
39  * Creates a new field in a class.
```



```

40 *
41 * $fieldName - The name of the field to create.
42 * $className - The fully qualified name of the class where the field is to be
    created.
43 * $typeName - The type of the field.
44 */
45 proc createField($fieldName, $accessor, $className, $typeName)
46   requires exists $cls in $utils.allClasses()
47     : $cls.getFullyQualifiedName().equals($className)
48   requires exists $cls in $utils.allClasses()
49     : $cls.getFullyQualifiedName().equals($className)
50     -> forall $fld in $cls.getFields()
51       : ~$fld.getName().equals($fieldName)
52 {
53   create field $fieldName in class $className of type $typeName
54 }
55
56 /**
57 * Creates a new static field in a class.
58 *
59 * $fieldName - The name of the field to create.
60 * $className - The fully qualified name of the class where the field is to be
    created.
61 * $typeName - The type of the field.
62 */
63 proc createStaticField($fieldName, $accessor, $className, $typeName)
64   requires exists $cls in $utils.allClasses()
65     : $cls.getFullyQualifiedName().equals($className)
66   requires exists $cls in $utils.allClasses()
67     : $cls.getFullyQualifiedName().equals($className)
68     -> forall $fld in $cls.getFields()
69       : ~$fld.getName().equals($fieldName)
70 {
71   create static field $fieldName in class $className of type $typeName
72 }
73
74 /**
75 * Creates a void method in a class with the given argument list.
76 *
77 * $methodName - The name of the method to create.
78 * $className - The fully qualified name of the class where the method is to be
    created.
79 * $arguments - The arguments to the method.
80 * $returnType - The type of the value returned by the method.
81 */
82 proc createMethod($methodName, $className, $arguments, $returnType)
83   requires exists $cls in $utils.allClasses()
84     : $cls.getFullyQualifiedName().equals($className)
85   requires forall $cls in $utils.allClasses()
86     : $cls.getFullyQualifiedName().equals($className)
87     -> (forall $mtd in $cls.getMethods()
88       : ~$utils.signatureMatches($mtd, $arguments))

```

```

89 {
90   create public method $methodName in class $className /*taking [ $arguments ]
91     */ returning $returnType
92 }
93 /**
94   * Creates a void method in a class with the given argument list.
95   *
96   * $methodName - The name of the method to create.
97   * $className - The fully qualified name of the class where the method is to be
98     created.
99   * $arguments - The arguments to the method.
100  */
101 proc createVoidMethod($methodName, $className, $arguments)
102   requires exists $cls in $utils.allClasses()
103     : $cls.getFullyQualifiedName().equals($className)
104   requires forall $cls in $utils.allClasses()
105     : $cls.getFullyQualifiedName().equals($className)
106     → (forall $mtd in $cls.getMethods()
107       : ~$utils.signatureMatches($mtd, $arguments))
108 {
109   create public method $methodName in class $className /*taking [ $arguments ]
110     */
111 }
112 /**
113   * Update operations.
114   */
115 /**
116   * Renames a class.
117   *
118   * $oldName - The current name of the class.
119   * $newName - The new name of the class.
120   */
121 proc renameClass($oldName, $newName)
122   requires forall $cls in $utils.allClasses()
123     : ~$cls.getFullyQualifiedName().equals($newName)
124 {
125   rename class $oldName to $newName
126 }
127 /**
128   * Renames a field.
129   *
130   * $oldName - The current name of the field.
131   * $newName - The new name for the field.
132   * $className - The name of the class where the field is declared.
133   */
134 proc renameField($oldName, $newName, $className)
135   requires forall $cls in $utils.allClasses()
136     : $cls.getFullyQualifiedName().equals($className)
137

```

```

138         -> (forall $fld in $cls.getFields()
139             : ~$fld.getIdentifier().equals($newName))
140     {
141     rename field $oldName to $newName in class $className
142     }
143
144 /**
145  * Renames a method.
146  *
147  * $oldName - The current name for the method.
148  * $newName - The desired new name for the method.
149  * $className - The name of the class where the method is currently located.
150  */
151 proc renameMethod($oldName, $newName, $className)
152 requires forall $cls in $utils.allClasses()
153         : $cls.getFullyQualifiedName().equals($className)
154         -> (forall $mtd in $cls.getMethods()
155             : ~$mtd.getName().equals(newName))
156     {
157     rename method $oldName to $newName in class $className
158     }
159
160 /**
161  * Changes the type of a field.
162  *
163  * $fieldName - The name of the field whose type is to change.
164  * $newType - The new type for the field.
165  * $className - The name of the class where the field is declared.
166  */
167 proc changeFieldType($fieldName, $newType, $className)
168 requires forall $cls in $utils.allClasses()
169         : $cls.getFullyQualifiedName().equals($className)
170         -> (forall $fld in $cls.getFields()
171             : $fld.getIdentifier().equals($fieldName)
172             -> $fld.eCrossReferences().isEmpty())
173     {
174     change type of field $fieldName in class $className to $newType
175     }
176
177 /**
178  * Changes the return type of a method.
179  *
180  * $methodName - The name of the method to change.
181  * $newType - The new return type for the method.
182  * $className - The name of the class where the method is defined.
183  */
184 proc changeMethodType($methodName, $newType, $className)
185 requires forall $cls in $utils.allClasses()
186         : $cls.getFullyQualifiedName().equals($className)
187         -> forall $mtd in $cls.getMethods()
188             : $mtd.getName().equals($methodName)
189             -> $mtd.eCrossReferences().isEmpty()

```

```

190 {
191   change type of method $methodName in class $className to $newType
192 }
193
194 /**
195  * Changes the access modifier on a method to public.
196  *
197  * $methodName - The name of the method to be given the public access modifier.
198  * $args - The arguments that form the method's signature.
199  * $className - The name of the class where the method is defined.
200  */
201 proc makeMethodPublic($methodName, $args, $className)
202   requires TRUE
203 {
204   make method $methodName in class $className public
205 }
206
207 /**
208  * Changes the access modifier on a method to private.
209  *
210  * $methodName - The name of the method to be given the private access modifier
211  *
212  * $args - The arguments that form the method's signature.
213  * $className - The name of the class where the method is defined.
214  */
215 proc makeMethodPrivate($methodName, $args, $className)
216   requires forall $cls in $utils.allClasses()
217     : $cls.getFullyQualifiedName().equals($className)
218     → forall $mtd in $cls.getMethods() :
219       $mtd.getName().equals($methodName) → $mtd.
220     eCrossReferences().isEmpty()
221 {
222   make method $methodName in class $className private
223 }
224
225 /**
226  * Changes the access modifier on a method to protected.
227  *
228  * $methodName - The name of the method to be given the protected access
229  * modifier.
230  * $args - The arguments that form the method's signature.
231  * $className - The name of the class where the method is defined.
232  */
233 proc makeMethodProtected($methodName, $args, $className)
234   requires forall $cls in $utils.allClasses()
235     : $cls.getFullyQualifiedName().equals($className)
236     → forall $mtd in $cls.getMethods() :
237       $mtd.getName().equals($methodName) → $mtd.
238     eCrossReferences().isEmpty()
239 {
240   make method $methodName in class $className protected
241 }

```

```

238
239 /**
240  * Changes the access modifier on a method to default.
241  *
242  * $methodName - The name of the method to be given the default access modifier
243  *
244  * $args - The arguments that form the method's signature.
245  * $className - The name of the class where the method is defined.
246  */
247 proc makeMethodDefault($methodName, $args, $className)
248   requires forall $cls in $utils.allClasses()
249     : $cls.getFullyQualifiedName().equals($className)
250     → forall $mtd in $cls.getMethods() :
251       $mtd.getName().equals($methodName) → $mtd.
252       eCrossReferences().isEmpty()
253 {
254   make method $methodName in class $className default
255 }
256
257 /**
258  * Changes the access modifier on a field to public.
259  *
260  * $fieldName - The name of the field to become public.
261  * $className - The name of the class where the field is located.
262  */
263 proc makeFieldPublic($fieldName, $className)
264   requires TRUE
265 {
266   make field $fieldName in class $className public
267 }
268
269 /**
270  * Changes the access modifier on a field to private.
271  *
272  * $fieldName - The name of the field to be marked as private.
273  * $className - The name of the class where the field is declared.
274  */
275 proc makeFieldPrivate($fieldName, $className)
276   requires forall $cls in $utils.allClasses()
277     : $cls.getFullyQualifiedName().equals($className)
278     → forall $fld in $cls.getFields() :
279       $fld.getName().equals($fieldName) → $fld.
280       eCrossReferences().isEmpty()
281 {
282   make field $fieldName in class $className private
283 }
284
285 /**
286  * Changes the access modifier on a field to protected.
287  *
288  * $fieldName - The name of the field which should have its access modifier
289  * changed.

```

```

286 * $className - The name of the class where the field is located.
287 */
288 proc makeFieldProtected($fieldName, $className)
289   requires forall $cls in $utils.allClasses()
290     : $cls.getFullyQualifiedName().equals($className)
291     -> forall $fld in $cls.getFields() :
292       $fld.getName().equals($fieldName) -> $fld.
293     eCrossReferences().isEmpty()
294 {
295   make field $fieldName in class $className protected
296 }
297 /**
298 * Changes the modifier on a field to default access.
299 *
300 * $fieldName - The field to make default.
301 * $className - The name of the class where the field is declared.
302 */
303 proc makeFieldDefault($fieldName, $className)
304   requires forall $cls in $utils.allClasses()
305     : $cls.getFullyQualifiedName().equals($className)
306     -> forall $fld in $cls.getFields() :
307       $fld.getName().equals($fieldName) -> $fld.
308     eCrossReferences().isEmpty()
309 {
310   make field $fieldName in class $className default
311 }
312 /**
313 * Adds an argument to a method.
314 *
315 * $argName - The name for the new argument.
316 * $argType - The type of the new argument.
317 * $methodName - The name of the method where the new argument is to be added.
318 * $className - The name of the class where the method is located.
319 */
320 proc addMethodArgument($argName, $argType, $methodName, $className)
321   requires forall $cls in $utils.allClasses()
322     : $cls.getFullyQualifiedName().equals($className)
323     -> forall $mtd in $cls.getMethods()
324       : $mtd.getName().equals($methodName)
325       -> forall $p in $mtd.getParameters()
326         : ~$p.getName().equals($argName)
327 {
328   add argument $argType $argName to method $methodName in class $className
329 }
330
331 /**
332 * Removes an argument from a method.
333 *
334 * $argName - The name of the argument to delete.
335 * $methodName - The name of the method to delete.

```

```

336 * $className - The name of the class where the method is currently located.
337 */
338 proc deleteMethodArgument($argName, $methodName, $className)
339     requires forall $cls in $utils.allClasses()
340         : $cls.getFullyQualifiedName().equals($className)
341         -> forall $mtd in $cls.getMethods()
342             : $mtd.getName().equals($methodName)
343             -> forall $p in $mtd.getParameters()
344                 : $p.getName().equals($argName)
345                 -> $p.eCrossReferences().isEmpty()
346 {
347     delete argument $argName from method $methodName in class $className
348 }
349
350 /**
351 * Reorders the arguments to a method by moving one argument in front of the
352 * other.
353 *
354 * $methodName - The name of the method that should have its arguments
355 * rearranged.
356 * $predecessor - The argument that should be closest to the head in the list
357 * of arguments.
358 * $successor - The argument that should be closest to the tail in the list of
359 * arguments.
360 * $className - The class where the method is currently located.
361 */
362 proc reorderMethodArguments($methodName, $predecessor, $successor, $className)
363     requires TRUE
364 {
365     place argument $predecessor in method $methodName with args [ ARTYPE ARGNAME
366     ] before $successor in class $className
367 }
368
369 /**
370 * Changes the superclass that a subclass currently extends.
371 *
372 * $subclass - The class whose parent class is to be changed.
373 * $superclass - The new class that $subclass extends.
374 */
375 proc changeSuperclass($subclass, $superclass)
376     requires forall $cls in $utils.allClasses()
377         : $cls.getFullyQualifiedName().equals($className)
378         -> $cls.eCrossReferences().isEmpty()
379 {
380     class $subclass extends $superclass
381 }
382
383 /**
384 * Moves a field one step up the inheritance hierarchy.
385 *
386 * $fieldName - The name of the field to move up the inheritance hierarchy.
387 * $className - The name of the class where the field is currently located.

```

```

383 */
384 proc moveFieldToSuperclass($fieldName , $className)
385   requires forall $fld in $utils.findClass($className).getFields()
386     : $fld.getName().equals($fieldName) → ($fld.isUnused() ∨ ~$fld.
      isPublic())
387 {
388   move field $fieldName from $className up
389 }
390
391 /**
392  * Moves a field one step down the inheritance hierarchy.
393  *
394  * $fieldName - The name of the field to move.
395  * $className - The name of the class where the field is currently located.
396  */
397 proc moveFieldToSubclass($fieldName , $className)
398   requires forall $cls in $utils.allClasses()
399     : $cls.getFullyQualifiedName().equals($className)
400     → forall $fld in $cls.getFields() :
401       $fld.getName().equals($fieldName) → $fld.
      eCrossReferences().isEmpty()
402 {
403   move field $fieldName from $className down
404 }
405
406 /**
407  * Moves a method from a class to its superclass.
408  *
409  * $methodName - The name of the method to move.
410  * $className - The name of the class where the method is currently located.
411  */
412 proc moveMethodToSuperclass($methodName , $className)
413   requires $utils.findMethod($className , $methodName).isUnused() ∨ $utils.
      findMethod($className , $methodName).isPublic()
414 {
415   move method $methodName from $className up
416 }
417
418 /**
419  * Moves a method from a class to its subclass.
420  *
421  * $methodName - The name of the method to move.
422  * $className - The class where that method is currently located.
423  */
424 proc moveMethodToSubclass($methodName , $className)
425   requires forall $cls in $utils.allClasses()
426     : $cls.getFullyQualifiedName().equals($className)
427     → forall $mtd in $cls.getMethods() :
428       $mtd.getName().equals($methodName) → $mtd.
      eCrossReferences().isEmpty()
429 {
430   move method $methodName from $className down

```



```

431 }
432
433 /**
434  * Moves a field from one class to another.
435  *
436  * $fieldName - The name of the field to move.
437  * $fromClass - The class where the field is currently located.
438  * $toClass - The class where the field is to be moved.
439  *
440  * requires $utils.findClass($fromClass).getField($fieldName).isUnused();
441  * requires ~$utils.fieldExists($toClass, $fieldName)
442  */
443 proc moveField($fieldName, $fromClass, $toClass)
444     requires ~$utils.fieldExists($toClass, $fieldName)
445 {
446     move field $fieldName from $fromClass to $toClass
447 }
448
449 /**
450  * Moves a method by name from one class to another.
451  *
452  * Note: ALL methods with the name will be moved. This does not take the
453  * remainder of the method signature into account.
454  *
455  * $methodName - The name of the method to move.
456  * $fromClass - The class where the method is currently located.
457  * $toClass - The class where the method is currently located.
458  */
459 proc moveMethod($methodName, $fromClass, $toClass)
460     requires $utils.findClass($fromClass).getMethod($methodName).isUnused()
461     requires ~$utils.methodExists($toClass, $methodName)
462 {
463     move method $methodName from $fromClass to $toClass
464 }
465
466 /*
467  * Delete operations.
468  */
469
470 /**
471  * Deletes a class by name.
472  *
473  * $className - The name of the class to delete from the project.
474  */
475 proc deleteClass($className)
476     requires forall $cls in $utils.allClasses() :
477         (forall $mtd in $cls.getMethods() : $mtd.eCrossReferences().isEmpty())
478         /\ (forall $fld in $cls.getFields() : $fld.eCrossReferences().isEmpty())
479 {
480     delete class $className
481 }
482

```

```
483 /**
484  * Deletes an unused field from a class.
485  *
486  * $fieldName - The name of the field to delete.
487  * $className - The name of the class where the field is declared.
488  */
489 proc deleteField($fieldName, $className)
490   requires $utils.findField($className, $fieldName).isUnused()
491   {
492     delete field $fieldName in class $className
493   }
494
495 /**
496  * Deletes an unused method from a class.
497  *
498  * $methodName - The name of the method to delete.
499  * $className - The name of the class to delete.
500  */
501 proc deleteMethod($methodName, $className)
502   requires $utils.findMethod($className, $methodName).isUnused()
503   {
504     delete method $methodName in class $className
505   }
```

Listing B.1: Standard Library for RefDsl

Bibliography

- Aho, A. V., Lam, M. S., Sethi, R., and Ullman, J. D. (2006). *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, 2nd edition.
- Al-Mutawa, H. A., Dietrich, J., Marsland, S., and McCartin, C. (2014). On the Shape of Circular Dependencies in Java Programs. In *Australasian Software Engineering Conference*, pages 48–57, Sydney. IEEE.
- Arcelli, F., Mangiacavalli, M., Pochiero, D., and Zanoni, M. (2015). On Experimenting Refactoring Tools to Remove Code Smells. In *Scientific Workshop Proceedings of the International Conference on Agile Software Development*, Helsinki. ACM.
- Barker, R. and Tempero, E. (2007). A Large-Scale Empirical Comparison of Object-Oriented Cohesion Metrics. In *Asia-Pacific Software Engineering Conference*, pages 414–421, Aichi. IEEE.
- Beck, K. (1999). *Extreme Programming Explained: Embrace Change*. Addison-Wesley, Boston, 1st edition.
- Bézivin, J. (2005). On the Unification Power of Models. *Software & Systems Modeling*, 4(2):171–188.
- Bieman, J. and Kang, B. K. (1995). Cohesion and Reuse in an Object-Oriented System. In *Symposium on Software Reusability*, pages 259–262, New York. ACM.
- Blackburn, S. M., Garner, R., Hoffmann, C., Khan, A. M., McKinley, K. S., Bentzur, R., Diwan, A., Feinberg, D., Frampton, D., Guyer, S. Z., Hirzel, M., Hosking, A., Jump, M., Lee, H., Moss, J. E. B., Phansalkar, A., and Stefanovi, D. (2006). The DaCapo Benchmarks : Java Benchmarking Development and Analysis. In *Object-Oriented Programming Systems, Languages, and Applications*, pages 169–190, Portland. ACM.
- Boehm, B. W. (1988). A Spiral Model of Software Development and Enhancement. *IEEE Computer*, 21(5):61–72.
- Brett, D., Digg, D., Garcia, K., and Marinov, D. (2007). Automated Testing of Refactoring Engines. In *European Software Engineering Conference and the Symposium on the Foundations of Software Engineering*, pages 9–10, Cavtat. ACM.
- Bruneton, E., Lenglet, R., and Coupaye, T. (2002). ASM: A Code Manipulation Tool to Implement Adaptable Systems. In *Adaptable and Extensible Component Systems*.
- Carvalho Júnior, A., Silva, L., and Cornélio, M. (2007). Using CafeOBJ to Mechanise Refactoring Proofs and Application. *Electronic Notes in Theoretical Computer Science*, 184:39–61.
- Cassell, K., Andreae, P., Groves, L., and Noble, J. (2009). Towards Automating Class-Splitting Using Betweenness Clustering. In *International Conference on Automated Software Engineering*, pages 595–599, Auckland. ACM/IEEE.

- Chaparro, O., Bavota, G., Marcus, A., and Penta, M. D. (2014). On the Impact of Refactoring Operations on Code Quality Metrics. In *International Conference on Software Maintenance and Evolution*, pages 456–460, Victoria. IEEE.
- Chidamber, S. R. and Kemerer, C. F. (1991). Towards a Metrics Suite for Object-Oriented Design. In *Object Oriented Programming Systems Languages and Applications*, pages 197–211, Phoenix. ACM.
- Chidamber, S. R. and Kemerer, C. F. (1994). A Metrics Suite for Object Oriented Design. *IEEE Transactions on Software Engineering*, 20(6):476–493.
- Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., and Talcott, C. (2003). The Maude 2.0 System. *Rewriting Techniques and Applications*, LNCS 2706:76–87.
- De Roover, C., Lammel, R., and Pek, E. (2013). Multi-Dimensional Exploration of API Usage. In *International Conference on Program Comprehension*, pages 152–161, San Francisco. IEEE.
- Dig, D., Comertoglu, C., Marinov, D., and Johnson, R. (2006). Automated Detection of Refactorings in Evolving Components. In *European Conference on Object-Oriented Programming*, chapter 24, pages 404–428. Springer, Berlin.
- Fokaefs, M. and Tsantalis, N. (2011). JDeodorant: Identification and Application of Extract Class Refactorings. In *International Conference on Software Engineering*, pages 1037–1039, Honolulu. ACM.
- Fontana, F. A., Ferme, V., and Spinelli, S. (2012). Investigating the Impact of Code Smells Debt on Quality Code Evaluation. In *International Workshop on Managing Technical Debt*, pages 15–22, Zurich. IEEE.
- Fowler, M. (1999). *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, Boston.
- Garrido, A. and Meseguer, J. (2006). Formal Specification and Verification of Java Refactorings. In *International Workshop on Source Code Analysis and Manipulation*, pages 165–174, Philadelphia. IEEE.
- Ge, X. and Murphy-Hill, E. (2014). Manual Refactoring Changes with Automated Refactoring Validation. In *International Conference on Software Engineering*, pages 1095–1105, Hyderabad. ACM.
- Girvan, M. and Newman, M. E. J. (2002). Community Structure in Social and Biological Networks. *National Academy of Sciences of the United States of America*, 99(12):7821–7826.
- Glass, R. (2001). Frequently Forgotten Fundamental Facts about Software Engineering. *IEEE Software*, 18(3):112–111.
- Griffith, I., Reimanis, D., Izurieta, C., Codabux, Z., Deo, A., and Williams, B. (2014). The Correspondence Between Software Quality Models and Technical Debt Estimation Approaches. In *International Workshop on Managing Technical Debt*, pages 19–26, Victoria. IEEE.
- Harman, M. and Clark, J. (2004). Metrics are Fitness Functions Too. In *International Software Metrics Symposium*, pages 58–69, Chicago. IEEE.
- Henkel, J. and Diwan, A. (2005). CatchUp! Capturing and Replaying Refactorings to Support API Evolution. *Proceedings. 27th International Conference on Software Engineering, 2005. ICSE 2005.*, pages 0–9.
- Hitz, M. and Montazeri, B. (1995). Measuring Coupling and Cohesion In Object-Oriented Systems. *Angewandte Informatik*, 50:1–10.

- Hoare, C. A. R. (1969). An Axiomatic Basis for Computer Programming. *Communications of the ACM*, 12(10):576–580.
- Jastram, M. (2014). *Rodin: User's Handbook*. CreateSpace Independent Publishing Platform.
- Jouault, F., Allilaire, F., Bézivin, J., and Kurtev, I. (2008). ATL: A Model Transformation Tool. *Science of Computer Programming*, 72(1-2):31–39.
- Karsai, G., Krahn, H., Pinkernell, C., Rumpe, B., Schindler, M., and Völkel, S. (2009). Design Guidelines for Domain Specific Languages. In *Workshop on Domain-Specific Modeling at Object Oriented Systems and Languages*, Orlando. ACM.
- Lanza, M., Marinescu, R., and Ducasse, S. (2005). *Object-Oriented Metrics in Practice*. Springer.
- Leavens, G. T., Baker, A. L., and Ruby, C. (1999). JML: A Notation for Detailed Design. In Kilov, H., Rumpe, B., and Harvey, W., editors, *Behavioural Specifications for Business and Systems*, chapter 12, pages 175–188. Kluwer.
- Lehman, M. and Fernández-Ramil, J. C. (2006). Software Evolution. In Madhavji, N. H., Fernández-Ramil, J., and Perry, D. E., editors, *Software Evolution and Feedback: Theory and Practice*, chapter 1, pages 7–37. John Wiley & Sons.
- Lincke, R., Lundberg, J., and Löwe, W. (2008). Comparing Software Metrics Tools. In *International Symposium on Software Testing and Analysis*, pages 131–142, Seattle. ACM.
- Lorenz, M. and Kidd, J. (1994). *Object-Oriented Software Metrics: A Practical Guide*. Prentice-Hall.
- Marinescu, R. (2004). Detection Strategies: Metrics-Based Rules for Detecting Design Flaws. In *International Conference on Software Maintenance*, pages 350–359, Chicago. IEEE.
- Martin, R. C. (2003). *Agile Software Development: Principles, Patterns, and Practices*. Prentice Hall.
- McCabe, T. J. (1976). A Complexity Measure. *IEEE Transactions on Software Engineering*, 2(4):308–320.
- McQuillan, J. A. (2011). *Using Model Driven Engineering to Reliably Automate the Measurement of Object-Oriented Software*, PhD Dissertation. Maynooth University.
- Meneely, A., Smith, B., and Williams, L. (2012). Validating Software Metrics: A Spectrum of Philosophies. *ACM Transactions on Software Engineering and Methodology*, 21(4):1–28.
- Mens, T. and Tourwe, T. (2004). A Survey of Software Refactoring. *IEEE Transactions on Software Engineering*, 30(2):126–139.
- Mens, T., Van Eetvelde, N., Demeyer, S., and Janssens, D. (2005). Formalizing Refactorings with Graph Transformations. *Journal of Software Maintenance and Evolution: Research and Practice*, 17(4):247–276.
- Metsker, S. J. and Wake, W. C. (2006). *Design Patterns in Java*. Addison-Wesley, 2nd edition.
- Meyer, B. (1997). *Object-Oriented Software Construction*. Prentice Hall, 2nd edition.
- Milne, I. and Rowe, G. (2002). Difficulties in Learning and Teaching Programming - Views of Students and Tutors. *Education and Information Technologies*, 7(1):55–66.
- Moghadam, I. H. and Ó Cinnéide, M. (2011). Code-Imp: A Tool for Automated Search-Based Refactoring. In *Workshop on Refactoring Tools*, pages 41–44, Honolulu. ACM.
- Muller, P.-A., Fondement, F., Baudry, B., and Combemale, B. (2012). Modeling modeling modeling. *Software & Systems Modeling*, 11(3):347–359.

- Murphy-Hill, E. and Black, A. P. (2008). Seven Habits of a Highly Effective Smell Detector. In *International Workshop on Recommendation Systems for Software Engineering*, pages 36–40, Atlanta. ACM.
- Murphy-Hill, E., Parnin, C., and Black, A. P. (2012). How We Refactor, and How We Know It. *IEEE Transactions on Software Engineering*, 38(1):5–18.
- Myers, G. J., Sandler, C., and Badgett, T. (2011). *The Art of Software Testing*. Wiley Publishing, 3rd edition.
- Nastov, B. (2013). *Grammar and Graphical Concrete Syntaxes Generator Assistant for Domain Specific Modelling Languages*, M.Sc. Dissertation. Montpellier II.
- O’Keeffe, M. and Ó Cinnéide, M. (2003). A Stochastic Approach to Automated Design Improvement. In *Principles and Practice of Programming in Java*, pages 16–18, Kilkenny. ACM.
- O’Keeffe, M. and Ó Cinnéide, M. (2006). Search-Based Software Maintenance. In *European Conference on Software Maintenance and Reengineering*, Bari. IEEE.
- O’Keeffe, M. and Ó Cinnéide, M. (2008a). Search-Based Refactoring : An Empirical Study. *Journal of Software Maintenance*, 20(5):345–364.
- O’Keeffe, M. and Ó Cinnéide, M. (2008b). Search-based Refactoring for Software Maintenance. *Journal of Systems and Software*, 81(4):502–516.
- OMG (2015). Meta Object Facility (MOF) 2.0 Query/View/Transformation Specification. Technical report.
- Opdyke, W. F. (1992). *Refactoring Object-Oriented Frameworks*, Ph.D. Dissertation. University of Illinois at Urbana-Champaign.
- Overbey, J. L. and Johnson, R. E. (2011). Differential Precondition Checking: A Lightweight, Reusable Analysis for Refactoring Tools. In *Automated Software Engineering*, pages 303–312, Oread. IEEE/ACM.
- Overbey, J. L., Johnson, R. E., and Munawar, H. (2016). Differential Precondition Checking: A Lightweight, Reusable Analysis for Refactoring Tools. *Automated Software Engineering*, 23(1):77–104.
- Schaefer, I., Bettini, L., Bono, V., Damiani, F., and Tanzarella, N. (2010). Delta-oriented Programming of Software Product Lines. *Software Product Lines: Going Beyond*, 6287 LNCS:77–91.
- Schneider, J.-G., Vasa, R., and Hoon, L. (2010). Do Metrics Help to Identify Refactoring? In *Workshop on Software Evolution and International Workshop on Principles of Software Evolution*, pages 3–7, Antwerp. ACM.
- Simon, F., Steinbruckner, F., and Lewerentz, C. (2001). Metrics Based Refactoring. In *Conference on Software Maintenance and Reengineering*, pages 30–38, Lisbon. IEEE.
- Soares, G., Gheyi, R., Serey, D., and Massoni, T. (2010). Making Program Refactoring Safer. *IEEE Software*, 27(4):52–57.
- Stachowiak, H. (1973). *Allgemeine Modelltheorie*. Springer.
- Steinberg, D., Budinsky, F., Paternostro, M., and Merks, E. (2009). *EMF: Eclipse Modeling Framework 2.0*. Addison-Wesley, 2nd edition.
- Tempero, E., Anslow, C., Dietrich, J., Han, T., Li, J., Lumpe, M., Melton, H., and Noble, J. (2010). The Qualitas Corpus: A Curated Collection of Java Code for Empirical Studies. In *Asia Pacific Software Engineering Conference*, pages 336–345, Sydney. IEEE.

-
- Thompson, S. (2011). *Haskell: The Craft of Functional Programming*. Addison-Wesley, 3rd edition.
- Tokuda, L. (2001). Evolving Object-Oriented Designs with Refactorings. *Automated Software Engineering*, 8:89–120.
- Völter, M., Benz, S., Dietrich, C., Engelmann, B., Helander, M., Kats, L., Visser, E., and Wachsmuth, G. (2013). *DSL Engineering Designing, Implementing and Using Domain-Specific Languages*. CreateSpace Independent Publishing Platform.