

DISS. ETH No. 21293

Enhancing Separation Logic for Object-Oriented

A dissertation submitted to
ETH ZURICH

for the degree of
Doctor of Sciences

presented by
Stephanus Johannes van Staden
Master of Science, International University in Germany

born
February 22nd, 1983

citizen of
South Africa

accepted on the recommendation of

Prof. Dr. Bertrand Meyer, examiner
Prof. Dr. Peter O'Hearn, co-examiner
Dr. Matthew Parkinson, co-examiner
Dr. Cristiano Calcagno, co-examiner

2013

ACKNOWLEDGEMENTS

I would like to thank Bertrand Meyer, my doctoral advisor, for his patience and support. The stimulating atmosphere at the Chair of Software Engineering at ETH Zurich made the ideas possible that eventually culminated in this thesis. Many thanks to the current and former group members who gave me feedback on draft papers: Jason (Yi) Wei, Sebastian Nanz, Martin Nordio, Carlo Furia and Scott West. My thanks also goes to Andreas Leitner, Ilinca Moser, Manuel Oriol, Michela Pedroni, Marco Piccioni, Chris Poskitt, Benjamin Morandi, Christian Estler, Julian Tschannen, Marco Trudel, Mischael Schill, Nadia Polikarpova and Claudia Günthart. My research was supported by ETH Research Grant ETH-15 10-1, for which I am very grateful. The organisation and infrastructure of ETH Zurich made it a real pleasure to do the work.

I am greatly indebted to Cristiano Calcagno for his advice and help. He spent many hours reading, correcting, advising and guiding my work, and just as many in explaining things and making helpful suggestions. Apart from our research collaboration, I am also thankful for his constant encouragement and support.

My thanks also goes to Peter O’Hearn and Matthew Parkinson, who made time in their busy schedules to co-examine this work. Matthew kindly explained details about jStar to me that I needed for implementing MultiStar. My internship at Microsoft Research in Cambridge and several discussions at conferences and workshops helped to shape my work.

Finally, I would like to thank my family and friends for their encouragement and support. Apart from my parents, brothers and sister, my thanks goes to Wha Suck, Dean, Alexandru, Antoine and Riaan. Last but not least, a very special thanks goes to Teresa, my wife, the sunshine in my life.

CONTENTS

1	Introduction	1
1.1	History and motivation	1
1.2	Overview of contributions	5
2	Background material	9
2.1	Language syntax	9
2.2	Logic syntax and semantics	11
2.3	Specification refinement	12
2.4	The specification environment	13
2.5	Statement verification	13
2.6	Method verification	15
2.7	Class and program verification	16
2.8	Useful lemmas	17
3	Reasoning about related abstractions	19
3.1	Examples	23
3.1.1	Intertwining ancestor abstractions	23
3.1.2	Access control and call protocols	28
3.1.3	Diamond inheritance	30
3.2	MultiStar	35
3.2.1	Front-end	35
3.2.2	Back-end	36
3.3	Case study	38
3.4	Formalisation	41
3.4.1	Language syntax	41
3.4.2	Operational semantics	41
3.4.3	Logic syntax and semantics	43
3.4.4	Specification refinement	43
3.4.5	The specification environment	43
3.4.6	Export information verification	43

3.4.7	Axiom verification	44
3.4.8	Statement verification	44
3.4.9	Method verification	44
3.4.10	Class and program verification	46
3.5	Conclusions and related work	47
4	Verifying executable contracts	51
4.1	Background	53
4.1.1	The abstract setting: triples and footprints	54
4.1.2	The concrete setting	54
4.2	Precondition verification	55
4.3	Postcondition verification	58
4.4	Class invariant verification	62
4.5	Model-based specifications	67
4.6	Relative purity and predicate extraction	71
4.6.1	Relative purity	71
4.6.2	Predicate extraction	73
4.7	Implementation	74
4.8	Conclusions and related work	75
5	Correctness by construction	77
5.1	Freefinement	79
5.1.1	The Inputs	79
5.1.2	The Extended Language and Formal System	81
5.1.3	System V_2 and Refinement	87
5.1.4	The Refinement of Refinement Systems	88
5.1.5	Discussion	95
5.2	Applications	99
5.2.1	Lambda Calculus	99
5.2.2	Hoare Logic	102
5.2.3	Discussion	110
5.3	Constructing correct OO programs	111
5.3.1	Language and proof rules	111
5.3.2	Development calculus	114
5.4	Related Work	118
6	Conclusion	121
A	Semantics of the proof system	123
B	Antitone Galois connections	127

ABSTRACT

Reasoning about programs can be tricky and error-prone. Formal verification facilitates the formulation and demonstration of rigorous arguments about program correctness. To be usable and useful, a program logic, i.e. a proof system for program verification, should accommodate the principal mechanisms of the programming language and crystallise the informal reasoning of programmers.

Recently, separation logic emerged as a promising tool for reasoning about shared mutable state, which pervades mainstream programming languages such as C, Java and Eiffel. Parkinson and others proposed a proof system for Object-Oriented (OO) programs that combines separation logic with mechanisms to reason about inheritance and dynamic dispatch. This system can verify a wide range of OO programs and design patterns in a concise way. The flexibility and simplicity of the system made it an attractive target for further improvement and broader application.

The contributions of this thesis address the following problems:

1. Specifying, verifying and using relationships between state abstractions. This is especially important when reasoning about OO programs that use multiple inheritance.
2. Reasoning about executable contracts. Executable contracts are often weak and may perform side-effects. Yet they capture useful design information, are programmer-friendly and assist in debugging.
3. Refinement and correctness by construction. Instead of first writing the code and then proving it correct, these techniques make it possible to write a correct program in the first place. The program and its correctness proof grow and evolve together.

State abstraction mechanisms, such as abstract predicates, are useful for reasoning about programs with modules that encapsulate state and hide information. Parkinson adapted abstract predicates to the OO setting, and

they play a central role in his proof system. Since OO code frequently relies on relationships between state abstractions of a class or a class hierarchy, this thesis enhances Parkinson’s system with mechanisms for specifying, verifying and exploiting such relationships. The extension also makes it possible to establish the logical consistency of a class hierarchy without considering implementation details, and it facilitates reasoning about multiple inheritance.

Existing OO code often contains contracts in the form of executable preconditions, postconditions and class invariants. These specifications are typically weaker than separation logic assertions, but they are more lightweight and perhaps more likely to be written by programmers. Contracts also record valuable information about program design and are useful for testing and debugging. This thesis contributes a new technique for using the separation logic assertions to verify that executable contracts will always hold at runtime and that they will not perform unwanted side-effects. As a result, verified contracts need not be monitored at runtime, and they add confidence in the correctness of the code and the separation logic specification.

Correctness by construction is an important feature of a mature engineering discipline. In the context of software engineering, it is realised by a calculus for top-down program development that features refinement as a central technique. A refinement calculus helps to construct correct code from a given specification in a series of steps. This thesis proposes freefinement – an algorithm for obtaining a sound refinement calculus from a modular program logic. The resulting refinement calculus can interoperate closely with the program logic, and it is even possible to reuse and translate proofs between them.

Many aspects of the work also apply to other settings. None of the contributions rely on a specific flavour of separation logic. The work on multiple inheritance subsumes interface inheritance, and the reasoning techniques for executable contracts generalise to non-OO languages that use explicit memory management. Finally, freefinement applies to a great variety of formal systems, including program logics for other languages and type systems.

ZUSAMMENFASSUNG

Beweisführung über Programme kann knifflig und fehleranfällig sein. Formale Verifikation unterstützt das Formulieren und Beweisen von präzisen Behauptungen über die Korrektheit von Programmen. Um brauchbar und nützlich zu sein, sollte eine Programmlogik, d.h. ein Beweissystem für Programmverifikation, die Hauptmechanismen der Programmiersprache unterstützen und die informelle Beweisführung der Programmierer präzisieren.

In letzter Zeit hat sich Separation Logic als nützliches Werkzeug erwiesen, um Beweise über einen gemeinsam genutzten, veränderlichen Zustandsraum zu führen, der in populären Programmiersprachen wie z.B. C, Java und Eiffel allgegenwärtig ist. Parkinson et al. schlugen ein Beweissystem für objektorientierte (OO) Programme vor, das Separation Logic mit Mechanismen für die Beweisführung über Vererbung und dynamischer Bindung kombiniert. Dieses System kann ein breites Spektrum von OO-Programmen und Entwurfsmustern knapp und präzise verifizieren. Die Flexibilität und Einfachheit des Systems haben es zu einem attraktiven Zielobjekt für weitere Verbesserungen und Einsatzmöglichkeiten gemacht.

Die Beiträge dieser Dissertation befassen sich mit den folgenden Problemen:

1. Spezifikation, Verifikation und Verwendung von Beziehungen zwischen Zustandsabstraktionen. Dies ist speziell dann wichtig, wenn Beweise über OO-Programme geführt werden, die Mehrfachvererbung verwenden.
2. Beweisführung über ausführbare Spezifikationen. Ausführbare Spezifikationen sind oft schwach und können mit Nebeneffekten behaftet sein. Dennoch erfassen sie nützliche Designinformationen, sind programmierfreundlich und helfen bei der Fehlersuche.
3. Verfeinerung und konstruktionsbedingte Korrektheit. Anstatt zuerst den Programmtext zu schreiben und dann zu verifizieren, erlauben diese Techniken es, das Programm von Anfang an korrekt zu schreiben. Das

Programm und der Beweis seiner Korrektheit wachsen und entwickeln sich gemeinsam.

Mechanismen für die Zustandsabstraktion, wie z.B. abstrakte Prädikate, sind nützlich für die Beweisführung über Programme mit Modulen, die Zustand und Informationen kapseln. Parkinson adaptierte abstrakte Prädikate an das OO-Umfeld und sie spielen eine zentrale Rolle in seinem Beweissystem. Da OO-Programme sich oft auf Beziehungen zwischen Zustandsabstraktionen einer Klasse oder Klassenhierarchie stützen, erweitert diese Dissertation Parkinsons System mit Mechanismen für die Spezifikation, Verifikation und Instrumentalisierung dieser Beziehungen. Die Erweiterung ermöglicht es auch, die logischen Konsistenz einer Klassenhierarchie zu begründen ohne Implementationsdetails zu berücksichtigen, und sie unterstützt die Beweisführung über Mehrfachvererbung.

Bereits existierende OO-Programme enthalten oft Spezifikationen in der Form von ausführbaren Vor-, Nachbedingungen und Klasseninvarianten. Diese Spezifikationen sind typischerweise schwächer als Zusicherungen in Separation Logic, aber sie sind auch schlanker und werden vielleicht eher von Programmierern geschrieben. Auch enthalten sie wertvolle Informationen über den Programmaufbau und sind nützlich für Tests und die Fehlerbehebung. Ein Beitrag dieser Dissertation ist eine neue Technik für die Verwendung von Zusicherungen in Separation Logic, die sicherstellt, dass ausführbare Spezifikationen zur Laufzeit immer eingehalten werden und keine ungewollten Nebeneffekte haben. Dadurch müssen ausführbare Spezifikationen zur Laufzeit nicht mehr überprüft werden und sie erhöhen das Vertrauen in die Korrektheit des Programms und der Spezifikationen in Separation Logic.

Konstruktionsbedingte Korrektheit (correctness-by-construction) ist eine wichtige Eigenschaft ausgereifter Ingenieursdisziplinen. Im Kontext von Software Engineering wird es durch einen Kalkül für Top-Down-Programmentwicklung realisiert, der Verfeinerung als zentrale Technik aufweist. Ein Verfeinerungskalkül hilft bei der Konstruktion korrekter Programme in einer Serie von Schritten, ausgehend von einer gegebenen Spezifikation. Diese Dissertation schlägt Freefinement vor – ein Algorithmus um einen korrekten Verfeinerungskalkül aus einer modularen Programmlogik zu erhalten. Der resultierende Verfeinerungskalkül ist in der Lage, mit der Programmlogik eng zu interagieren, und es ist sogar möglich, Beweise zwischen diesen beiden wiederzuverwenden und zu übersetzen.

Viele Aspekte dieser Arbeit sind auch auf andere Umgebungen anwendbar, keiner der Beiträge hängt von spezifischen Ausprägungen von Separation Logic ab. Die Arbeit an Mehrfachvererbungen umfasst auch Schnittstellenvererbung, und die Beweisführungstechnik für ausführbare Spezifikationen kann

auf Nicht-OO-Sprachen mit explizitem Speichermanagement verallgemeinert werden. Überdies betrifft Freefinement eine grosse Vielfalt von formalen Systemen, inklusive Programmlogiken für andere Sprachen und Typsysteme.

CHAPTER 1

INTRODUCTION

1.1 History and motivation

Object-oriented (OO) programming languages are widely used in software engineering. They follow the imperative paradigm, i.e. programs are commands that transform state. Perhaps the most distinguishing feature of OO languages is that they promote a data-centered approach to programming: objects encapsulate data and operations associated with that data. Since humans often think about a system or a problem domain in terms of its constituent objects, the object abstraction is a natural fit for modelling them and hence reflecting this understanding in software.

Mainstream OO languages have many features, such as classes, inheritance, dynamic binding, dynamic allocation of objects and sharing of references to objects. Although most OO programs use these features, their effects can be subtle and hard to understand. Especially in more complex programs where several features interact, it is often difficult to see that the code is correct, i.e. that it will do what it is supposed to do. What is needed is a formalism that facilitates simple reasoning about the correctness of OO programs.

A system for verifying the correctness of imperative programs that manipulate simple variables has long been known: Hoare logic [32]. The basic judgement of Hoare logic is the triple $\{P\}c\{Q\}$, where P is the precondition, c the command and Q the postcondition. Informally, it means: if c is executed in an initial state satisfying P , and it terminates, then the final state will satisfy Q . An example of a valid triple is $\{x = 7\}x := x + 1\{x = 8\}$.

Hoare logic is compositional, or modular, which is important for larger verification efforts. Both programs and specifications should compose nicely. For example, the rule for reasoning about a sequential composition involves

only the specifications of its two operands:

$$\frac{\{P\}c\{Q\} \quad \{Q\}c'\{R\}}{\{P\}c; c'\{R\}}$$

The rule of constancy is a good example of how specifications compose. Provided c does not modify any of the free variables in R , we have:

$$\frac{\{P\}c\{Q\}}{\{P \wedge R\}c\{Q \wedge R\}}$$

The rule states that a command preserves all properties of variables it did not modify. Continuing the earlier example, we would be able to conclude $\{x = 7 \wedge y = 6\}x := x + 1\{x = 8 \wedge y = 6\}$. A simple syntactic check suffices for determining the variables that a command might modify. Everything else must stay the same.

The situation is rather straightforward for imperative languages that manipulate simple variables: the state is just a mapping from variable names to values, and the only way to manipulate a variable is to mention it explicitly in the program text. OO languages are more complicated. In addition to simple variables (in the stack), the state also has shared objects (in the heap) that are accessed indirectly through references. Because of reference aliasing, there is no simple syntactic check for determining whether an object remains unmodified by a piece of code. OO programs regularly use many objects and much sharing, so it is important to reason about them in a convenient way.

It turns out that there are several ways to extend Hoare logic to programs that manipulate shared mutable state such as objects. One family of approaches keep the rule of constancy. Doing so requires more complicated techniques to determine 1) the set of objects or parts of objects that a command may modify, and 2) whether the predicate R describes any of them. Because 1 is not a simple syntactic check, and 2 tends to dissect R and not limit the inspection to its free variables, we shall not consider such approaches in this thesis.

Separation logic [54, 60] is another way of generalising Hoare logic to programs with shared mutable state. It adapts the rule of constancy by replacing the $\wedge$ with $*$, a new logical connective called separating conjunction. A state satisfies the predicate $P * Q$ when 1) its simple variables satisfy the constraints imposed on them by P and Q , and 2) it is possible to partition the objects into two disjoint sets of storage locations (i.e. heap locations or object-field pairs) such that the one set satisfies the constraints of P and the other those of Q . The reasoning about simple variables is exactly the same as in Hoare logic. Separation logic recognises that the concept of separation

is important for reasoning about the new kind of storage: mutable locations that may be aliased by references.

The adapted rule is known as the *frame rule*:

$$\frac{\{P\}c\{Q\}}{\{P * R\}c\{Q * R\}}$$

It is applicable whenever c does not modify any of the free variables in R . This simple syntactic check needs no information about the definitions of predicates that appear in R , which is good for abstraction. The frame rule says that the separate objects described by R will not be modified. In order for this to work, it must be impossible for c to change the storage locations that R describes. The separation logic triple $\{P\}c\{Q\}$ therefore has a generalised meaning¹: if c is executed in an initial state satisfying P , then it will not access storage locations except those described in P or allocated by c , and if it terminates, then the final state will satisfy Q . The frame rule is very useful for knowing how a command behaves in a larger context, i.e. when more objects are present. It is heavily used in many proofs of correctness, and it tends to make them simple and short.

Separation logic simplifies reasoning about shared mutable state, but OO programs also use other features that are historically difficult to reason about. Inheritance and the associated dynamic dispatch (or dynamic binding) of routine calls enable a single piece of code to have an incredible variety of runtime behaviours. Static reasoning therefore needs to impose some discipline on inheritance and the behaviour of (overridden) routines. Moreover, the OO paradigm encourages programmers to think about objects as black boxes with hidden implementation details. Since the purpose of formalism is to make intuitive arguments rigorous, it is desirable to express and exploit such abstraction patterns in formal reasoning.

In his PhD thesis [57] and work with others [56, 58], Parkinson proposed an elegant proof system that addresses these issues. The main ingredients of the system are as follows (Chapter 2 contains the details):

- It uses separation logic to reason about the fields of objects.
- Abstract predicate families (apfs) provide encapsulated data abstractions in the logic. An apf is a predicate with a designated ‘target’ argument. Each class C can define the apf predicate for the case when the target argument has dynamic type C . It is only possible to use this

¹That c does not access (i.e. read or modify) these locations is a stronger requirement than what we actually need here, but it is widely used and useful for reasoning about concurrency.

definition within class C when the dynamic type of the target argument is C – the predicate must be treated abstractly in all other cases. Apfs can thus express the idea that objects (or parts of objects or object aggregates) are black boxes for outsiders, and that subclasses may change the representation. A specification of a routine typically says how the code transforms one or more of these black boxes into others.

- Each routine has a static and a dynamic specification, each consisting of a precondition and a postcondition. The static specification states in detail what the routine body does. It is used to reason about calls that use static dispatch (e.g. *precursor* or *super* calls). The dynamic specification is typically more abstract – it describes the idea behind the routine that must be respected in all subclasses. It imposes discipline on inheritance, and is used to reason about dynamically dispatched calls.

The organisation of specifications is modular: a routine body is verified only against the static specification, which must in turn be consistent with the dynamic specification. Static specifications make it unnecessary to re-verify code in subclasses, so each routine body is verified exactly once.

These ingredients combine to yield a compact system that can reason about many uses and abuses of inheritance, including behaviour extension, restriction and modification, and change of representation in subclasses.

Because of all its desirable features, this system is an ideal starting point for further enhancements. This thesis considers the improvement of the system and its broader application in reasoning. It tackles the following problems:

1. The basic apf mechanism is not always flexible enough. A class might want to share some properties of its definitions with others. Also, some properties of definitions, or relationships between them, may be invariant in a whole class hierarchy, and client code might rely on this. For example, it could be a design decision that all subclasses must use the same definition and hence the same implementation of a logical abstraction. Or it might be intended that an implication between two different apf predicates will always hold. We need ways to express, verify and use such information.
2. OO languages like Eiffel use multiple inheritance and deferred (abstract) classes and routines. Reasoning about these constructs requires extensions to the original system.

3. Much existing OO code has been specified with executable contracts, such as preconditions, postconditions and class invariants. They capture expectations and intentions in a programmer-friendly way. There has been several attempts to use them as a basis for verification, but it tends to be complicated because they 1) are often too weak, 2) do not have a separating conjunction, and hence struggle with framing and aliasing, and 3) often contain side-effects. The separation logic system ignores these assertions, but it would be better to reason about them and exploit them in verification.
4. Instead of hacking away at code and verifying its correctness afterwards, it is often easier to develop code that is correct by construction. However, the separation logic system is geared for bottom-up verification instead of top-down development. It does not directly support techniques for stepwise transformation such as refinement. A major goal of this thesis is therefore to support refinement and correctness by construction in a way that complements the separation logic system.

1.2 Overview of contributions

Chapter 3 addresses the first problem by adding features to Parkinson's system. It shows with examples how a single class may employ multiple apfs to achieve its purpose, and that these are often related and combined in disciplined ways. Since the correctness of OO code frequently depends on the relationships between these abstractions, we need mechanisms to express, verify and use such information in a rigorous fashion. In particular, the work introduces two general specification mechanisms: *export clauses* for relating abstractions in individual classes, and *axiom clauses* for relating abstractions in a class and all its descendants, i.e. for an entire (sub-)hierarchy of classes. The information expressed in export and axiom clauses is first verified before it is assumed for the verification of code. Proof rules that avoid logical inconsistencies formalise the manipulation.

Chapter 3 also extends Parkinson's system to support multiple inheritance and deferred (abstract) classes and routines. This subsumes reasoning about interface inheritance, a form of multiple inheritance used in languages like Java. Multiple inheritance usually intertwines several logical abstractions from parent classes, and it turns out that export and axiom clauses are useful for specifying such information. Knowing more about ancestor classes through specifications means that fewer of their implementation details have to be considered, and verification is consequently more modular.

Chapter 4 attacks the third problem: reasoning about executable contracts and exploiting them in verification. The work shows how separation logic can be used to verify that executable contracts are satisfied, i.e. that they will always hold at runtime and not perform undesirable side-effects. Both the program and its executable assertions are verified with respect to separation logic specifications. A novel notion called *relative purity* embraces historically problematic side-effects in executable specifications, and verification boils down to proving *connecting implications*. Even model-based specifications, i.e. expressive executable contracts that use mathematical abstractions, can be verified. A failed verification attempt may indicate a discrepancy between the two kinds of specification. For example, the separation logic precondition of a routine may be too liberal for the particular application, and the executable precondition can help to uncover this problem. Or a failed verification may point to situations in which a runtime violation of an executable postcondition would occur. In any case, the separation logic and executable specifications complement each other in verification. Several examples illustrate the use and utility of the approach.

The above mechanisms are well-suited to automation. We implemented them in MultiStar, a fully automatic verification tool for verifying Eiffel programs. MultiStar uses separation logic specifications that are embedded in the source program. The verification proceeds in two steps. Firstly, the specifications and code are translated into an intermediate representation that is simpler to process. Secondly, a reasoning engine based on jStar [23] performs the actual verification. Section 3.2 describes the architecture of MultiStar and its treatment of export/axiom clauses and multiple inheritance. Section 4.7 explains how it supports reasoning about executable contracts. MultiStar successfully verified the examples in Chapters 3 and 4. Moreover, Section 3.3 describes a case study in which MultiStar was used to verify the core iterator hierarchy of a popular data structure library. MultiStar is freely available as part of the EVE (Eiffel Verification Environment) download [28].

Chapter 5 tackles refinement and correctness by construction. The approach is to obtain a refinement calculus for OO commands, and then to use it as part of a larger calculus that handles other OO features. This reflects the fact that developing a routine body involves different reasoning than, say, adding an attribute to a class or a dynamic specification to a routine.

It is moreover a cost-effective way to structure a solution, because the refinement calculus comes for free! Section 5.1 describes *free refinement*, an algorithm that constructs a sound refinement calculus from a verification system under certain conditions. In this work, a verification system is any formal system for establishing whether an inductively defined term, such as a command, satisfies a specification. Examples of verification systems include

Hoare logics and type systems. Freefinement first extends the term language to include specification terms, and builds a verification system for the extended language that is a sound and conservative extension of the original system. The extended system is then transformed into a sound refinement calculus. The resulting refinement calculus can interoperate closely with the verification system – it is even possible to reuse and translate proofs between them. Freefinement gives a semantics to refinement at an abstract level: it associates each term of the extended language with a set of terms from the original language, and refinement simply reduces this set. For illustration purposes, Section 5.2 applies freefinement to a simple type system for the lambda calculus and also to a Hoare logic.

It is straightforward to apply freefinement to the part of the separation logic proof system that reasons about OO commands. The resulting refinement calculus is plugged into a larger system that deals with other OO constructs, and this larger calculus is compatible with (and complements) the full separation logic system. In particular, all its transformation steps preserve correctness with respect to the system, and every verifiable OO program can be constructed from scratch in a stepwise manner.

Many of the results in this thesis are quite general and also apply to different settings. None of the contributions rely on a particular flavour of separation logic. The work on executable contracts can be used for non-OO languages with explicit memory management such as C. Finally, freefinement can give refinement calculi for a great variety of formal systems, including modular type systems such as System F.

CHAPTER 2

BACKGROUND MATERIAL

This chapter describes the OO programming language and its separation logic proof system that are used throughout the thesis. It summarises the published work of Parkinson [57] and his joint work with Bierman [58] with minor changes in the presentation.

2.1 Language syntax

The grammar of our kernel OO language with specifications is shown in Figure 2.1. It is deliberately minimal to keep the presentation simple, but it contains the essential language constructs of popular OO languages. The concrete notation is based on Eiffel [26].

A class is divided into top-level sections. The **inherit** section lists its parent class (or ‘ANY’), the **define** section contains predicate definitions, and methods and fields are written in the **feature** section. Empty sections and ‘**inherit** ANY’ will simply be omitted in examples.

A sequence of *c*’s is denoted by $\bar{c}$. The letters G and H are used for class names, m for method names, and f for field names. Variables are denoted by u, x, y and z. **Void** corresponds to ‘null’ in other languages. Two reserved program variables **Current** and **Result** denote the current object (‘this’) and the result of a function call respectively. **Current** is never **Void**. The term *feature* describes methods and fields, sometimes called the ‘members’ of a class. Feature overloading, including method overloading, is not allowed.

Separate namespaces exist for class names, p, m and f. We assume the absence of clashes when names are introduced. This precludes method overloading and field shadowing. Later, when we extend the language with multiple inheritance, it will guarantee that methods or fields with the same name in parent classes stem from common ancestors.

$L ::= \text{class } G \text{ inherit } H \text{ define } \overline{D} \text{ feature } \overline{M} \overline{F} \text{ end}$	
$D ::= x.\text{p}_G(\overline{t}:\overline{y}) \text{ as } P$	<i>Define clause</i>
$M ::= \text{introduce } m(\text{Args}) \text{ Rt Sd Ss B}$	<i>Method declaration</i>
override $m(\text{Args}) \text{ Rt Sd Ss B}$	
inherit $m(\text{Args}) \text{ Rt Sd Ss}$	
$F ::= f: \text{Type}$	<i>Field declaration</i>
$Sd ::= \text{dynamic Spec}$	<i>Dynamic specification</i>
$Ss ::= \text{static Spec}$	<i>Static specification</i>
$\text{Spec} ::= \{P\}\text{-}\{Q\} \mid \{P\}\text{-}\{Q\} \text{ also Spec}$	<i>Specification</i>
$B ::= \text{do } s \text{ end}$	<i>Method body</i>
$s ::= \text{skip}$	<i>Statement</i>
$x := e$	<i>Assignment</i>
$s ; s$	<i>Sequential composition</i>
if e then s else s end	<i>Conditional</i>
while e do s end	<i>Loop</i>
$x: \text{Type}. s$	<i>Local variable block</i>
$x := \text{new } G$	<i>Object allocation</i>
$x := y.f$	<i>Field lookup</i>
$x.f := e$	<i>Field update</i>
$x := y.m(\overline{e}) \mid y.m(\overline{e})$	<i>Dynamically dispatched call</i>
$x := y.G::m(\overline{e}) \mid y.G::m(\overline{e})$	<i>Direct method call</i>
$e ::= x \mid e + e \mid e = e \mid \text{Void} \mid 0 \mid 1 \mid 2 \dots$	<i>Expression</i>
$\text{Type} ::= \text{INT} \mid \text{BOOL} \mid G$	
$\text{Args} ::= \overline{x: \text{Type}}$	<i>Formal arguments</i>
$\text{Rt} ::= \epsilon \mid : \text{Type}$	<i>Return type</i>

Figure 2.1: The kernel language grammar.

A constructor is simply an introduced method m where m is a class name. Except for the restriction that subclasses cannot inherit or override constructors, no special treatment is needed otherwise. In the examples, we will omit the target of a method call or field assignment if it is **Current**.

Methods have both **static** and **dynamic** specifications. A specification is written in pre-post form $\{P\}\text{-}\{Q\}$, or $\{P_1\}\text{-}\{Q_1\}$ **also** $\{P_2\}\text{-}\{Q_2\}$ to indicate that both are satisfied (Section 2.3 shows how the **also** form abbreviates a single pre-post specification). A method's dynamic specification must be satisfied by all subclasses, and is used to verify dynamically-dispatched calls. A static specification describes properties about the particular method body, and is used to verify statically-dispatched calls, including **Precursor** ('super' or 'base') calls and direct calls $x.C::m(\bar{e})$ in C++ style.

To provide subclasses with the opportunity to respecify a method and to simplify the proof rules that follow later, we require a subclass to inherit or override explicitly all non-constructor methods present in its parent (in the examples of later chapters and the MultiStar tool, specification shorthands are employed to achieve this).

We assume the formal argument names of methods stay the same in subclasses. This simplifies the proof rules that follow, which would otherwise need additional substitutions.

Finally, to improve the readability of examples in later chapters, we frequently write statements in a list and omit the interspersed sequential composition operators.

2.2 Logic syntax and semantics

The predicates used in specifications and proofs have the following grammar.

$$\begin{aligned} P, Q, S, T, \Delta ::= & \forall x.P \mid P \Rightarrow Q \mid \text{false} \mid e = e' \mid x : G \mid x <: G \mid x.f \hookrightarrow e \mid P * Q \\ & \mid x.\textcolor{blue}{p}(\overline{t:e}) \quad \textit{Apf predicate} \\ & \mid x.\textcolor{red}{p}_G(\overline{t:e}) \quad \textit{Apf entry} \end{aligned}$$

The predicate $x : G$ means x references an object whose dynamic type is exactly G , and $x <: G$ means x references an object whose dynamic type is a subtype of G . In both cases $x \neq \mathbf{Void}$, and $x : G \Rightarrow x <: G$ holds. Within a context, if x is declared of type G then $x <: G$ whenever $x \neq \mathbf{Void}$.

Abstract predicate families [56, 58] provide logical abstractions of state for the OO setting. An abstract predicate family (abbreviated *apf*) provides an abstract predicate p for which each class C can define an entry p_C . The first argument of an apf predicate or entry is called the *root*. The root followed by a dot is written before the apf predicate or entry name. The whole prefix

is omitted when the root is **Current**. An apf predicate's root can never be **Void**. Since the meaning of an apf predicate depends on the dynamic type of the root object, it can be seen as mirroring dynamic dispatch of object-orientation in the logic. The other arguments of an apf predicate or entry are a set (i.e. order-independent collection) of tagged arguments, where tag names t provide useful hints about the purpose of tagged values. Apfs offer high levels of abstraction in specifications and proofs: apf predicates and entries are treated abstractly by the clients of a class.

Other predicates have the usual intuitionistic separation logic semantics. Informally, the predicate $x.f \hookrightarrow e$ means that the f field of object x has value e ($f \hookrightarrow e$ abbreviates **Current**. $f \hookrightarrow e$), and $P * Q$ means that P and Q hold for disjoint portions of the heap. Readers are referred to [60, 54, 57] for a formal treatment of separation logic. Symbols such as $\Leftrightarrow$, $\neg$, **true**, $\vee$, $\wedge$ and $\exists$ are encoded in the standard way. Every occurrence of $_$ in a predicate denotes a fresh existentially quantified variable, where the quantifier is placed in the innermost position. $FV(P)$ denotes the free variables of P ; every method precondition P must satisfy **Result** $\notin FV(P)$.

In the rest of the formalisation, the symbols P , Q , S and T are used for assertions and predicates, and Δ for assumptions.

2.3 Specification refinement

A specification $\{P_1\}\text{-}\{Q_1\}$ is refined by $\{P_2\}\text{-}\{Q_2\}$ when any statement s that satisfies $\{P_1\}\text{-}\{Q_1\}$ also satisfies $\{P_2\}\text{-}\{Q_2\}$. Under the assumptions Δ , the specification $\{P_1\}\text{-}\{Q_1\}$ is refined by $\{P_2\}\text{-}\{Q_2\}$ if we can prove $\Delta \vdash \{P_1\}\text{-}\{Q_1\} \implies \{P_2\}\text{-}\{Q_2\}$, i.e. provide a proof tree with leaves $\Delta \vdash \{P_1\}\text{-}\{Q_1\}$ and root $\Delta \vdash \{P_2\}\text{-}\{Q_2\}$ built with the structural rules of separation logic (Consequence, Frame, Auxiliary Variable Elimination, Disjunction, and others). In the context of method specification refinement, Δ contains the apf assumptions of the class, and the Consequence and Frame rules are given by:

$$\frac{\Delta \Rightarrow (P' \Rightarrow P) \quad \Delta \vdash \{P\}\text{-}\{Q\} \quad \Delta \Rightarrow (Q \Rightarrow Q')}{\Delta \vdash \{P'\}\text{-}\{Q'\}} \text{Consequence}$$

$$\frac{\Delta \vdash \{P\}\text{-}\{Q\}}{\Delta \vdash \{P * T\}\text{-}\{Q * T\}} \text{Frame}$$

The Frame rule is applicable whenever **Result** $\notin FV(T)$, and expresses that disjoint portions of the heap stay unchanged.

Method specifications can be combined with **also** (Definition 1 in [58]):

$$\{P_1\}_{-}\{Q_1\} \textbf{ also } \{P_2\}_{-}\{Q_2\} \stackrel{\text{def}}{=} \{(P_1 \wedge x = 1) \vee (P_2 \wedge x \neq 1)\}_{-}\{(Q_1 \wedge x = 1) \vee (Q_2 \wedge x \neq 1)\}$$

where x denotes a fresh auxiliary variable. The specifications $\{P_1\}_{-}\{Q_1\}$ and $\{P_2\}_{-}\{Q_2\}$ are *equivalent w.r.t. Δ* iff both $\Delta \vdash \{P_1\}_{-}\{Q_1\} \implies \{P_2\}_{-}\{Q_2\}$ and $\Delta \vdash \{P_2\}_{-}\{Q_2\} \implies \{P_1\}_{-}\{Q_1\}$. Two specifications are *equivalent* iff they are equivalent w.r.t. all Δ . It can be shown that **also** is commutative, associative and idempotent modulo equivalence with identity $\{\text{false}\}_{-}\{\text{true}\}$. The notation **also** _{$i \in I$} $\{P_i\}_{-}\{Q_i\}$ denotes the specification $\{P_{e_1}\}_{-}\{Q_{e_1}\} \textbf{ also } \dots \textbf{ also } \{P_{e_m}\}_{-}\{Q_{e_m}\}$, where $e_1 \dots e_m$ are the elements of the finite index set I . Furthermore, when I is the empty set:

$$\textbf{also}_{i \in \emptyset} \{P_i\}_{-}\{Q_i\} \stackrel{\text{def}}{=} \{\text{false}\}_{-}\{\text{true}\}$$

It always holds that $\Delta \vdash \{P\}_{-}\{Q\} \implies \{\text{false}\}_{-}\{\text{true}\}$. Other useful lemmas involving **also** are given in Section 2.8. Finally, we use the abbreviation

$$\begin{aligned} \Delta \vdash \{P_1\}_{-}\{Q_1\} &\stackrel{\text{Current} : G}{\implies} \{P_2\}_{-}\{Q_2\} \\ &\stackrel{\text{def}}{=} \Delta \vdash \{P_1\}_{-}\{Q_1\} \implies \{P_2 * \text{Current} : G\}_{-}\{Q_2\} \end{aligned}$$

in the formalisation of the Dynamic dispatch proof obligation for methods in Section 2.6.

2.4 The specification environment

Most of the proof rules that follow use an environment Γ , which maps method names to their specifications for all classes in a program:

$$\begin{aligned} \Gamma &::= G.m \mapsto (\bar{x}, \{P\}_{-}\{Q\}) && \text{Method dynamic specification} \\ &| G::m \mapsto (\bar{x}, \{S\}_{-}\{T\}) && \text{Method static specification} \\ &| \bar{\Gamma} \end{aligned}$$

The $\bar{x}$ in a specification of m denote its formal argument names. Γ is guaranteed to be a partial function for well-formed programs, and we write $\Gamma(G.m) = (\bar{x}, \{P\}_{-}\{Q\})$ for $G.m \mapsto (\bar{x}, \{P\}_{-}\{Q\}) \in \Gamma$, and $\Gamma(G::m) = (\bar{x}, \{S\}_{-}\{T\})$ for $G::m \mapsto (\bar{x}, \{S\}_{-}\{T\}) \in \Gamma$.

2.5 Statement verification

The assumptions Δ used to verify statements contain apf information about the enclosing class. The rules for most statements are standard:

$$\begin{array}{c}
\frac{}{\Delta; \Gamma \vdash_s \{P\} \mathbf{skip} \{P\}} \\
\frac{}{\Delta; \Gamma \vdash_s \{P[e/x]\} x := e \{P\}} \\
\frac{\Delta; \Gamma \vdash_s \{P\} s \{Q\} \quad \Delta; \Gamma \vdash_s \{Q\} s' \{R\}}{\Delta; \Gamma \vdash_s \{P\} s; s' \{R\}} \\
\frac{\Delta; \Gamma \vdash_s \{P * e\} s \{Q\} \quad \Delta; \Gamma \vdash_s \{P * \neg e\} s' \{Q\}}{\Delta; \Gamma \vdash_s \{P\} \mathbf{if} \ e \ \mathbf{then} \ s \ \mathbf{else} \ s' \ \mathbf{end} \{Q\}} \\
\frac{\Delta; \Gamma \vdash_s \{P * e\} s \{P\}}{\Delta; \Gamma \vdash_s \{P\} \mathbf{while} \ e \ \mathbf{do} \ s \ \mathbf{end} \{P * \neg e\}}
\end{array}$$

The rule for local variables has the proviso that $x \notin FV(P) \cup FV(Q)$:

$$\frac{\Delta; \Gamma \vdash_s \{P\} s \{Q\}}{\Delta; \Gamma \vdash_s \{P\} x : \text{Type.} \ s \{Q\}}$$

In the rule for object allocation, $allfields(G)$ denotes the set of field names listed in G and all its ancestors:

$$\frac{allfields(G) = \{f_1, f_2, \dots, f_n\}}{\Delta; \Gamma \vdash_s \{\mathbf{true}\} \\ x := \mathbf{new} \ G \\ \{x.f_1 \hookrightarrow _ * x.f_2 \hookrightarrow _ * \dots * x.f_n \hookrightarrow _ * x : G\}}$$

For field lookup, when x is not free in e and not the same as y , we have:

$$\frac{}{\Delta; \Gamma \vdash_s \{y.f \hookrightarrow e\} x := y.f \{y.f \hookrightarrow e * x = e\}}$$

Field update is simple:

$$\frac{}{\Delta; \Gamma \vdash_s \{x.f \hookrightarrow _ \} x.f := e \{x.f \hookrightarrow e\}}$$

Dynamically dispatched calls use the dynamic specs of methods in Γ , while direct calls use the static ones. Provided x is not y and x is not free in $\bar{e}$, the rules for result-returning calls are:

$$\frac{\Gamma(G.m) = (\bar{u}, \{P\} _ \{Q\})}{\Delta; \Gamma \vdash_s \{P[y, \bar{e}/\mathbf{Current}, \bar{u}] * y <: G\} \\ x := y.m(\bar{e}) \\ \{Q[y, \bar{e}, x/\mathbf{Current}, \bar{u}, \mathbf{Result}]\}}$$

$$\frac{\Gamma(G::m) = (\bar{u}, \{S\} \cdot \{T\})}{\frac{\Delta; \Gamma \vdash_s \{S[y, \bar{e}/\mathbf{Current}, \bar{u}] * y \neq \mathbf{Void}\} \quad x := y.G::m(\bar{e}) \quad \{T[y, \bar{e}, x/\mathbf{Current}, \bar{u}, \mathbf{Result}]\}}{\Delta; \Gamma \vdash_s \{S[y, \bar{e}/\mathbf{Current}, \bar{u}] * y \neq \mathbf{Void}\} \quad x := y.G::m(\bar{e}) \quad \{T[y, \bar{e}, x/\mathbf{Current}, \bar{u}, \mathbf{Result}]\}}}$$

The important structural rules for this thesis are Frame, Consequence, Auxiliary Variable Elimination, and Disjunction. The Frame rule is the key to local reasoning. Provided s modifies no variable in $FV(T)$:

$$\frac{\Delta; \Gamma \vdash_s \{P\}s\{Q\}}{\Delta; \Gamma \vdash_s \{P * T\}s\{Q * T\}} \text{ Frame}$$

The rule of Consequence allows the use of assumptions Δ :

$$\frac{\Delta \Rightarrow (P' \Rightarrow P) \quad \Delta; \Gamma \vdash_s \{P\}s\{Q\} \quad \Delta \Rightarrow (Q \Rightarrow Q')}{\Delta; \Gamma \vdash_s \{P'\}s\{Q'\}} \text{ Consequence}$$

In Auxiliary Variable Elimination, the variable x to be eliminated may not appear free in s .

$$\frac{\Delta; \Gamma \vdash_s \{P\}s\{Q\}}{\Delta; \Gamma \vdash_s \{\exists x.P\}s\{\exists x.Q\}} \text{ AuxVarElim}$$

Finally, the Disjunction rule is standard:

$$\frac{\Delta; \Gamma \vdash_s \{P\}s\{Q\} \quad \Delta; \Gamma \vdash_s \{P'\}s\{Q'\}}{\Delta; \Gamma \vdash_s \{P \vee P'\}s\{Q \vee Q'\}} \text{ Disjunction}$$

2.6 Method verification

As for statement verification, the assumptions Δ used to verify method definitions contain apf information about the method's enclosing class.

A newly introduced method's static and dynamic specifications must be consistent¹, and its body must satisfy the static specification. These two requirements are captured by the Dynamic dispatch [D.d.] and Body verification [B.v.] proof obligations respectively.

¹We establish that the dynamic specification follows from the static one when the dynamic type of **Current** is G . If the dynamic type of x is G , then the body of m in G will be executed if $x.m$ is called. The static and dynamic specifications must be consistent with each other in this case, since the dynamic specification is used for reasoning about the call statement, whereas the body was verified only w.r.t. the static specification.

$$\begin{array}{l}
B = \mathbf{do\ s\ end} \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
\Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\Delta; \Gamma \vdash_s \{S_G\}s\{T_G\} \quad [\text{B.v.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{introduce}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ B\ \text{in}\ G\ \text{parent}\ H
\end{array}$$

The next rule is used whenever a method body gets redefined. Consistency must be proven between the new dynamic specification and those in the parent class; this is embodied in the Behavioural subtyping [B.s.] proof obligation. The other proof obligations are identical to those for method introduction above.

$$\begin{array}{l}
\Gamma(H.m) = (\bar{x}, \{P_H\}\text{-}\{Q_H\}) \\
B = \mathbf{do\ s\ end} \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
\Delta \vdash \{P_G\}\text{-}\{Q_G\} \Rightarrow \{P_H\}\text{-}\{Q_H\} \quad [\text{B.s.}] \\
\Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\Delta; \Gamma \vdash_s \{S_G\}s\{T_G\} \quad [\text{B.v.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{override}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ B\ \text{in}\ G\ \text{parent}\ H
\end{array}$$

When a method is inherited, its static specification must follow from the one in the parent class. The Inheritance [Inh.] obligation ensures that this will be the case. The Behavioural subtyping and Dynamic dispatch obligations serve the same purposes as before.

$$\begin{array}{l}
\Gamma(H.m) = (\bar{x}, \{P_H\}\text{-}\{Q_H\}) \\
\Gamma(H::m) = (\bar{x}, \{S_H\}\text{-}\{T_H\}) \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
\Delta \vdash \{P_G\}\text{-}\{Q_G\} \Rightarrow \{P_H\}\text{-}\{Q_H\} \quad [\text{B.s.}] \\
\Delta \vdash \{S_H\}\text{-}\{T_H\} \Rightarrow \{S_G\}\text{-}\{T_G\} \quad [\text{Inh.}] \\
\Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{inherit}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ \text{in}\ G\ \text{parent}\ H
\end{array}$$

2.7 Class and program verification

A class is verified by verifying all of its methods. Once again, the formula Δ contains apf information that is specific to the class.

$$\frac{\forall M_i \in \overline{M} \cdot \Delta; \Gamma \vdash_m M_i \text{ in } G \text{ parent } H}{\Delta; \Gamma \vdash_c \text{class } G \text{ inherit } H \text{ define } \overline{D} \text{ feature } \overline{M} \overline{F} \text{ end}}$$

Finally, here is the rule for program verification:

$$\frac{\begin{array}{l} \forall i \in 1..n. L_i = \text{class } G_i \dots \text{end} \\ \Gamma = \text{specs}(L_1 \dots L_n) \\ \forall i \in 1..n. \text{apf}(L_i); \Gamma \vdash_c L_i \\ \text{true}; \Gamma \vdash_s \{\text{true}\} s \{\text{true}\} \end{array}}{\vdash_p L_1 \dots L_n s}$$

The function *apf* translates the abstract predicate family definitions of a class into a formula – its apf assumptions. It is adapted from [58] for tagged arguments:

$$\begin{aligned} \text{apf}(\text{class } G \dots \text{define } D_1 D_2 \dots D_n \text{ feature } \dots \text{end}) &\stackrel{\text{def}}{=} \text{apf}_G(D_1) \wedge \dots \wedge \text{apf}_G(D_n) \\ \text{apf}_G(x.\text{p}_G(Y) \text{ as } P) &\stackrel{\text{def}}{=} \text{FtoE}(p, G, Y) \wedge \text{EtoD}(x.\text{p}_G(Y) \text{ as } P) \wedge (\forall x <: G \cdot \text{TR}(p, x, Y)) \\ \text{FtoE}(p, G, \overline{t:y}) &\stackrel{\text{def}}{=} \forall x, \overline{y}. x : G \Rightarrow [x.\text{p}(\overline{t:y}) \Leftrightarrow x.\text{p}_G(\overline{t:y})] \\ \text{EtoD}(x.\text{p}_G(\overline{t:y}) \text{ as } P) &\stackrel{\text{def}}{=} \forall x, \overline{y}. x.\text{p}_G(\overline{t:y}) \Leftrightarrow P \\ \text{TR}(p, x, \overline{t:y}) &\stackrel{\text{def}}{=} \bigwedge_{\overline{t':y'} + \overline{t'' : y''} \equiv \overline{t:y}} \forall \overline{y'}. x.\text{p}(\overline{t':y'}) \Leftrightarrow x.\text{p}(\overline{t':y'} + \overline{t'' : _}) \end{aligned}$$

2.8 Useful lemmas

Lemmas 2.1 and 2.2 are frequently used in proofs of Behavioural Subtyping and Inheritance:

Lemma 2.1. $\Delta \vdash (\text{also}_{i \in I} \{P_i\} - \{Q_i\}) \Longrightarrow \{P_k\} - \{Q_k\}$ for all $k \in I$.

Lemma 2.2. If $\Delta \vdash \{P\} - \{Q\} \Longrightarrow \{S_i\} - \{T_i\}$ for all $i \in I$, then $\Delta \vdash \{P\} - \{Q\} \Longrightarrow (\text{also}_{i \in I} \{S_i\} - \{T_i\})$.

For Body Verification:

Lemma 2.3. If $\Delta; \Gamma \vdash_s \{S_i\} s \{T_i\}$ for all $i \in I$, then under assumptions Δ and Γ , statement s satisfies $(\text{also}_{i \in I} \{S_i\} - \{T_i\})$.

CHAPTER 3

REASONING ABOUT RELATED ABSTRACTIONS

The use of data abstractions is a hallmark of OO programming. A class is a typical example of such an abstraction. In interface or general multiple inheritance hierarchies, a class can combine and maintain several abstractions offered by its parents. Although most examples of this chapter involve abstractions in connection with inheritance, not all data abstractions are directly coupled with language constructs. Classes use them for various purposes: to simplify how clients manipulate a class, to separate various concepts that are combined in a class, or to encourage or enforce particular call protocols. For example, a complex object with a long initialisation phase can use the abstractions ‘initialising’ and ‘ready’, with methods applicable to an ‘initialising’ object, others to a ‘ready’ object, and some to both. Or algorithms can manipulate a mutable data structure under an ‘immutable’ abstraction, even if there is no interface making this explicit.

Relationships between data abstractions are important when reasoning about OO code. This chapter explores the problem of relating abstractions in an information hiding setting, where implementation details of abstractions are hidden from clients. Suppliers must therefore express and fulfill relationships between abstractions. If a class offers ‘student’ and ‘person’ abstractions (by using inheritance, or other means), for example, it might allow clients to convert a ‘student’ abstraction into a ‘person’ one. Clients can then manipulate the ‘person’ abstraction by calling e.g. library routines. After the manipulation, they might be allowed to convert back and assume that the number of exams the ‘student’ has taken is still the same. Specification mechanisms are needed to express and enforce the programmer’s intentions about such relationships. This is especially important in multiple inheritance

hierarchies where classes combine multiple abstractions in complicated ways.

Abstract predicate families are a flexible mechanism for data abstraction in OO specifications. An apf P provides a predicate name for an abstraction; each class C can define an *entry* predicate P_C . The definition of P_C describes how class C implements the apf P , and is hidden from other classes. For example, apfs S and P can be used to provide an abstraction of students and persons in a program respectively. Class `STUDENT` can define the entries $S_{STUDENT}$ and $P_{STUDENT}$, while other classes can define their apf entries differently. The predicate $x.S(\text{age}:a, \text{exm}:e)$ describes object x under the ‘student’ abstraction: its age is ‘ a ’ and the number of exams taken is ‘ e ’. If the dynamic type of x is `STUDENT`, then class `STUDENT` can use the fact that

$$x.S(\text{age}:a, \text{exm}:e) \Leftrightarrow x.S_{STUDENT}(\text{age}:a, \text{exm}:e)$$

In other words, the dynamic type of the first argument of an apf predicate (in this case the dynamic type of x) determines which apf entry applies. Apf predicates can therefore be seen to mirror dynamic dispatch of OO programs in the logic. The apf mechanism is modular and exercises information hiding: only the class defining an apf entry knows the definition and can relate its entry to the apf predicate.

The relationship described before, namely that a ‘student’ abstraction can be converted into a ‘person’ one, can be expressed as follows:

$$x.S(\text{age}:a, \text{exm}:e) \Rightarrow x.P(\text{age}:a)$$

Allowing the back conversion without affecting the number of exams requires a stronger property that uses separation logic’s $*$ -connective:

$$x.S(\text{age}:a, \text{exm}:e) \Leftrightarrow [x.P(\text{age}:a) * x.\text{RestStoP}(\text{exm}:e)] \quad (A)$$

where RestStoP abstracts the parts of a ‘student’ abstraction that are independent from and not included in a ‘person’ one. With this property, a client can now reason as the following proof outline shows:

```

{x.S(age: a, exm: e)}
{x.P(age: a) * x.RestStoP(exm: e)}
  {x.P(age: a)}
    // Manipulation of ‘person’ abstraction by library routines.
    {x.P(age: a + 1)}
  {x.P(age: a + 1) * x.RestStoP(exm: e)}
{x.S(age: a + 1, exm: e)}
```

The Frame rule of separation logic guarantees that the disjoint $x.\text{RestStoP}(\text{exm}:e)$ remains unchanged. In essence, the client uses property (A) in combination with the Frame rule to infer an S -based specification for the library manipulation. It is not necessary to re-specify or re-verify the library – knowledge of the relationship saves specification overhead and keeps reasoning modular.

To which objects the property (A) applies is a design choice: a programmer might express that *selected* classes in a heterogeneous hierarchy fulfill the relationship, or that *all* classes in a homogeneous hierarchy fulfill it. We introduce two general specification mechanisms for the two cases: *export clauses* to express properties that hold for individual classes, and *axiom clauses* to describe properties of entire hierarchies. If class STUDENT specifies

export

$$\forall x,a,e. x : \text{STUDENT} \Rightarrow [x.\text{S}(\text{age}:a, \text{exm}:e) \Leftrightarrow [x.\text{P}(\text{age}:a) * x.\text{RestStoP}(\text{exm}:e)]]$$

where $\{\}$

then a client must know that an object's dynamic type is exactly STUDENT before using the information in reasoning. On the other hand, if class STUDENT specifies

axiom

$$\text{S_P}: \forall a,e. \text{S}(\text{age}:a, \text{exm}:e) \Leftrightarrow [\text{P}(\text{age}:a) * \text{RestStoP}(\text{exm}:e)]$$

then clients can use the stronger implication

$$\forall x,a,e. x <: \text{STUDENT} \Rightarrow [x.\text{S}(\text{age}:a, \text{exm}:e) \Leftrightarrow [x.\text{P}(\text{age}:a) * x.\text{RestStoP}(\text{exm}:e)]]$$

Knowledge that the object's dynamic type is a subtype of STUDENT (including STUDENT) suffices to use the relationship. This is much more convenient for clients: a sound OO type system will guarantee that if a variable has static type STUDENT and references an object, then the object's dynamic type will always be a subtype of STUDENT.

Axiom clauses offer a general facility to constrain the implementation of abstractions in subclasses. For example, class STUDENT can express that the number of exams a 'student' has taken is always non-negative, and that all subclasses should use its implementation of the 'student' abstraction:

axiom

$$\text{exm_non_neg}: \forall a,e. \text{S}(\text{age}:a, \text{exm}:e) \Rightarrow 0 \leq e$$

$$\text{S_constraint}: \forall a,e. \text{S}(\text{age}:a, \text{exm}:e) \Leftrightarrow \text{S}_{\text{STUDENT}}(\text{age}:a, \text{exm}:e)$$

Axiom clause 'exm_non_neg' guarantees clients that $0 \leq e$ whenever they know $x.\text{S}(\text{age}:a, \text{exm}:e)$ and $x <: \text{STUDENT}$. Subclass representation constraints such as the one expressed in the 'S_constraint' clause are useful for

ensuring safe interaction between statically and dynamically dispatched calls on the same object – a pervasive pattern in OO programs (see Section 5.5 of [58]).

The claims made in export and axiom specifications must be checked to obtain sound reasoning. Whether or not a class fulfills axiom clauses often depends on properties of particular other classes, such as its parents. For this reason our proof system has a layered assumption structure: axiom verification can use export information of all classes in a program, and method verification can additionally use axiom information. Several examples will show how export and axiom clauses are verified and applied in verification.

This chapter extends the proof system of Parkinson and Bierman (see chapter 2) with export/axiom clauses, abstract classes, abstract methods and shared multiple inheritance¹ where fields and methods of common ancestor classes are not replicated in the descendant [26]. Apart from the use of apfs to support abstraction and information hiding, Parkinson and Bierman’s system has the attractive property that it can verify a wide range of inheritance uses and abuses. Flexible handling of inheritance is vital in a proof system for multiple inheritance, since classes often interrelate methods and data from parents in complicated ways: methods can intertwine ancestor abstractions as the example in Section 3.1.1 shows, or abstractions can share parts in the case of diamond inheritance. Export and axiom clauses provide flexible ways to relate abstractions, thereby giving clients an integrated view of how methods affect abstractions from different inheritance paths. Very few proof systems exist for multiple inheritance, and no proof system we know of can facilitate reasoning about multiple related abstractions at the same level of abstraction as ours.

We implemented our proof system in MultiStar – a fully automatic verification tool. MultiStar has a two-tier architecture: a GUI front-end that translates Eiffel code and specifications into a simpler form for verification, and a language-independent back-end based on jStar [23] for reasoning. The front-end uses specifications written inside classes. Future front-ends for e.g. Java and C# can reuse the MultiStar back-end: with its support for interface inheritance, export/axiom clauses, abstract classes and abstract methods, a wide range of programs can be verified. As we shall see, export and axiom clauses can also be useful in single-inheritance situations.

All the examples presented in this chapter have been verified with MultiStar. To demonstrate the flexibility of our approach, we also used MultiStar to verify the iterator hierarchy of the Gobo data structure library [29]. The complete code and specifications of the examples and Gobo case study are

¹Inheritance with *virtual base classes* in C++ terminology [27].

available online [28].

Chapter outline Several examples illustrating the new specification mechanisms and proof system follow in Section 3.1. Section 3.2 presents the MultiStar tool, and Section 3.3 reports on the case study with Gobo iterator classes. A formal exposition of our proof system appears in Section 3.4. Section 3.5 mentions related work and concludes. Appendix A contains an overview of the formal semantics of our proof system and a proof of soundness.

3.1 Examples

A class may inherit from several parent classes. This is indicated by listing the names of all parents in its **inherit** section. The new **export** and **axiom** sections respectively contain the export and axiom clauses of a class. As usual, empty sections will simply be omitted.

We use the method specification shorthands of [58]: if only a static specification is listed, the dynamic specification is assumed to be exactly the same, and if only a dynamic specification is listed, then a static specification is derived by replacing each apf predicate whose first argument is **Current** with the entry predicate of the class. In other words, if the shorthand is used in class C , then $p(\overline{t:e})$ is replaced with $p_C(\overline{t:e})$. Specifications of non-constructor methods are furthermore propagated down the hierarchy: if a class does not explicitly list an inherited method, then it is assumed to have the same static and dynamic specifications as determined for the parent class. To avoid ambiguity, we require that if the method is available in multiple parents, then they must all have identical specifications for it.

The examples do not discuss details that are uninteresting from the perspective of this thesis, such as proofs of correctness of simple ancestor classes. The paper of Parkinson and Bierman [58] contains several examples that illustrate the basics.

3.1.1 Intertwining ancestor abstractions

Classes CELL and COUNTER are shown in Figure 3.1. CELL models mutable integer-valued cells and uses apf **Cell**, while COUNTER uses apf **Cn**. The apfs provide logical abstractions of mutable cells and counters respectively. Class CCELL in Figure 3.2, the focus of this example, inherits from CELL and COUNTER. It intertwines the functionality of its parents by overriding *set_value* to store the value and increment the count. It uses apf **Cc** to

```

class CELL
define x.CellCELL(val: v) as x.value  $\hookrightarrow$  v
feature
  introduce CELL(v: INT)
  dynamic {value  $\hookrightarrow$  _} - {Cell(val: v)}
  do value := v end

  introduce value(): INT
  dynamic {Cell(val: v)} - {Cell(val: v) * Result = v}
  do Result := value end

  introduce set_value(v: INT)
  dynamic {Cell(val: _)} - {Cell(val: v)}
  do value := v end

  value: INT
end

class COUNTER
define x.CnCOUNTER(cnt: c) as x.count  $\hookrightarrow$  c
feature
  introduce COUNTER()
  dynamic {count  $\hookrightarrow$  _} - {Cn(cnt: 0)}
  do count := 0 end

  introduce count(): INT
  dynamic {Cn(cnt: c)} - {Cn(cnt: c) * Result = c}
  do Result := count end

  introduce increment()
  dynamic {Cn(cnt: c)} - {Cn(cnt: c + 1)}
  do tmp: INT. tmp := count; count := tmp + 1 end

  count: INT
end

```

Figure 3.1: The CELL and COUNTER classes.

```

class CCELL inherit CELL COUNTER
define
  x.CellCCELL(val: v, cnt: c) as x.CcCCELL(val: v, cnt: c)
  x.CnCCELL(cnt: c) as x.CnCOUNTER(cnt: c)
  x.CcCCELL(val: v, cnt: c) as x.CellCELL(val: v) * x.CnCOUNTER(cnt: c)
export
   $\forall x. x : \text{CCELL} \Rightarrow [\forall c, v. x.\text{Cc}(\text{val}: v, \text{cnt}: c) \Leftrightarrow x.\text{Cell}(\text{val}: v, \text{cnt}: c) \Leftrightarrow$ 
     $(x.\text{Cn}(\text{cnt}: c) * \text{Rest}(x, v))]$  where { Rest(x, v) = x.CellCELL(val: v) }
feature
  introduce CCELL(v: INT)
  dynamic {value  $\hookrightarrow$  _ * count  $\hookrightarrow$  _} - {Cc(val: v, cnt: 0)}
  do Precursor{CELL}(v); Precursor{COUNTER}() end

  inherit value(): INT
  dynamic {Cc(val: v, cnt: c)} - {Cc(val: v, cnt: c) * Result = v}
    also {Cell(val: v, cnt: c)} - {Cell(val: v, cnt: c) * Result = v}

  override set_value(v: INT)
  dynamic {Cc(val: _, cnt: c)} - {Cc(val: v, cnt: c + 1)}
    also {Cell(val: _, cnt: c)} - {Cell(val: v, cnt: c + 1)}
  do CCELL::increment(); CELL::set_value(v) end

  inherit count(): INT
  dynamic {Cc(val: v, cnt: c)} - {Cc(val: v, cnt: c) * Result = c}
    also {Cn(cnt: c)} - {Cn(cnt: c) * Result = c}

  inherit increment()
  dynamic {Cc(val: v, cnt: c)} - {Cc(val: v, cnt: c + 1)}
    also {Cn(cnt: c)} - {Cn(cnt: c + 1)}
end

// In an arbitrary class or library:
use_counter(c: COUNTER)
  dynamic {c.Cn(cnt: v)} - {c.Cn(cnt: v + 10)}

use_cell(c: CELL, v: INT)
  dynamic {c.Cell(val: _)} - {c.Cell(val: v)}

```

Figure 3.2: The CCELL class and two library methods.

provide an abstraction of such objects in the logic, and ‘grows’ $\text{Cell}_{\text{CCELL}}$ to accommodate method *set_value*, as we shall see.

The single export clause of CCELL relates the Cc , Cell and Cn abstractions. Only the predicate in front of **where** is exported for reasoning. Predicate definitions following **where** are used only to verify the clause and allow a class to hide implementation/representation details in its interface without introducing new predicate families. For example, if we modify the representation of class CCELL, then the definition of $\text{Rest}(x,v)$ can be changed without invalidating correctness proofs of client code.

To verify an export clause, we must prove that the exported predicate follows from the standard apf assumptions of the class and the predicate definitions after the **where** keyword. The proof for CCELL’s export clause is trivial (Section 2.7 describes the standard apf assumptions of a class). Note that export clauses are not verified for a particular dynamic type, since the standard apf assumptions of a class do not assume a particular one. It is therefore sound to use exported information to verify axiom clauses and methods of other classes, as we will do in later examples. However, nothing prevents the user from writing exported predicates as implications where the antecedent is of the form $x : \text{Type}$ (see the export clause of CCELL). In this case information in the consequent can be applied only to objects satisfying this type constraint.

For the constructor we have to prove that its body satisfies the static specification (note that a **Precursor** call is syntactic sugar for a direct call):

```
{value  $\leftrightarrow$  _ * count  $\leftrightarrow$  _}
  CELL::CELL(v)
{CellCCELL(val:v) * count  $\leftrightarrow$  _}
  Precursor{COUNTER}()
{CellCCELL(val:v) * CnCOUNTER(cnt:0)}
{CcCCELL(val:v, cnt:0)}
```

The constructor body simply passes the needed fields to parent constructors and treats their internal representations abstractly thereafter.²

Method *value* is respecified in CCELL with Cell and Cc specifications. Since it inherits the body from CELL, we must prove that the new static specification is satisfied assuming the body’s static specification of CELL. This method proof obligation is called Inheritance in the formalisation. By applying the Frame rule (with $\text{Cn}_{\text{COUNTER}}(\text{cnt}:c)$) and then the rule of Con-

²To simplify the formal presentation of proofs and make them more transparent, we mention fields explicitly in constructor preconditions. MultiStar injects them automatically – see Section 3.2.1 for more discussion.

sequence, we can derive each **also**-ed static specification, which is sufficient to conclude the proof by Lemma 2.2. Another proof obligation for *value* is Behavioural subtyping, where we must show that the dynamic specification listed in `CELL` follows from the new one, i.e. that `CCELL` maintains the old specification. For the proof, we first ‘choose’ the `Cell` dynamic spec with Lemma 2.1, and then remove the `cnt` tag by applying the Auxiliary Variable Elimination and Consequence rules. The application of Auxiliary Variable Elimination quantifies the variable `c` existentially in the pre- and postcondition, and the application of Consequence uses tag reduction information which is part of `CCELL`’s standard apf assumptions.

Behavioural subtyping of *set_value* is similar. For its Body verification obligation, we must prove that both `CellCCELL` and `CcCCELL` static specifications are satisfied. The proof proceeds as follows:

```

{CellCCELL(val: _, cnt: c)}
  CCELL::increment()
{CellCCELL(val: _, cnt: c + 1)}
{CellCCELL(val: _) * CnCOUNTER(cnt: c + 1)}
  CELL::set_value(v)
{CellCCELL(val: v) * CnCOUNTER(cnt: c + 1)}
{CellCCELL(val: v, cnt: c + 1)}

```

An application of Consequence proves the other **also**-ed static spec and Lemma 2.3 completes the proof. As the body operates on state described by `CellCELL` and `CnCOUNTER`, the proof obligations and separation logic’s faulting semantics demand that we ‘grow’ `CellCCELL` to include both state parcels.

Now consider the two library routines at the bottom of Figure 3.2. The export clause contains the necessary information to prove the two triples:

```

{true} cc := new CCELL(5); use_counter(cc) {cc.Cc(val: 5, cnt: 10)}
{true} cc := new CCELL(5); use_cell(cc, 20) {cc.Cc(val: 20, cnt: _)}

```

The proof of the second triple reduces and expands tags according to standard apf rules:

```

{true}
  cc := new CCELL(5)
{cc : CCELL * cc.Cc(val: 5, cnt: 0)}
{cc : CCELL * cc.Cell(val: 5, cnt: 0)}
{cc : CCELL * cc.Cell(val: 5, cnt: _)}
{cc : CCELL * cc.Cell(val: 5)}
  use_cell(cc, 20)
{cc : CCELL * cc.Cell(val: 20)}

```

```
{cc : CCELL * cc.Cell(val: 20, cnt: _)}
{cc.Cc(val: 20, cnt: _)}
```

Information about `cnt` is lost in the postcondition, which is unavoidable because *use_cell* could call *set_value* more than once. In a version of CCELL where `CnCCELL` is defined to include the `CellCCELL` state and the equivalence of `Cc`, `Cn` and `Cell` is exported, information about `val` will likewise be lost in the first triple. Also note that dynamic type information is required to use the exported relationships, since the subclasses of CCELL are not obliged to implement them.

3.1.2 Access control and call protocols

Our proof system can enforce interesting access control patterns in verified programs. Consider class CCEL2 in Figure 3.3 which has the same executable code as CCELL but different specifications. Its export clause relates the `Cell` and `Cn` abstractions in a one-directional way. The constructor produces a `Cell` apf predicate with which methods *value*, *set_value* and *count* can be called. Verified clients cannot call *increment* with the `Cell` predicate. They must use exported information to get a `Cn` predicate, yet they lack information to change back after the call: no export or axiom clause is available to do this, and every method producing a `Cell` predicate requires one. The following proof attempt where `cc2 : CCEL2` shows the problem:

```
{cc2.Cell(val: v, cnt: c)}
{cc2.Cn(cnt: c)}
  cc2.increment()
{cc2.Cn(cnt: c + 1)}
{???}
{cc2.Cell(val: _, cnt: _)} // The weakest requirement of set_value.
  cc2.set_value(10)
{cc2.Cell(val: _, cnt: _)}
```

While the client has a `Cell` predicate, the argument tagged by `cnt` and returned by *count* reflects precisely how many times the value has been set. If the client tries to manipulate the count by calling *increment*, then it can never regain the needed capability to call *value* and *set_value*, and must forever treat the object as a simple counter in the code. The combination of abstract predicate relationships and method specifications enforces this protocol in verified code.

```

class CCEL2 inherit CELL COUNTER
define
  x.CellCCEL2(val: v, cnt: c) as x.CellCELL(val: v) * x.CnCOUNTER(cnt: c)
  x.CnCCEL2(cnt: c) as x.CnCOUNTER(cnt: c)
export
   $\forall x. x : \text{CCEL2} \Rightarrow [\forall v, c. x.\text{Cell}(\text{val}: v, \text{cnt}: c) \Rightarrow x.\text{Cn}(\text{cnt}: c)]$  where {}
feature
  introduce CCEL2(v: INT)
  dynamic {value  $\hookrightarrow$  _ * count  $\hookrightarrow$  _} - {Cell(val: v, cnt: 0)}
  do Precursor{CELL}(v); Precursor{COUNTER}() end

  inherit value(): INT
  dynamic {Cell(val: v, cnt: c)} - {Cell(val: v, cnt: c) * Result = v}

  override set_value(v: INT)
  dynamic {Cell(val: _, cnt: c)} - {Cell(val: v, cnt: c + 1)}
  do CCEL2::increment(); CELL::set_value(v) end

  inherit count(): INT
  dynamic {Cell(val: v, cnt: c)} - {Cell(val: v, cnt: c) * Result = c}
    also {Cn(cnt: c)} - {Cn(cnt: c) * Result = c}
end

```

Figure 3.3: The CCEL2 class.

```

class PERSON
  define x. $P_{PERSON}$ (age: a) as x.age  $\hookrightarrow$  a
  export  $\forall x, a. x.P_{PERSON}(age: a) \Leftrightarrow x.age \hookrightarrow a$  where {}
  feature
    introduce PERSON(a: INT)
    dynamic {age  $\hookrightarrow$  -}_{-}{ $P(age: a)$ }
    do age := a end

    introduce age(): INT
    dynamic { $P(age: a)$ }_{-}{ $P(age: a) * \text{Result} = a$ }
    do Result := age end

    introduce set_age(a: INT)
    dynamic { $P(age: -)$ }_{-}{ $P(age: a)$ }
    do age := a end

    introduce celebrate_birthday()
    static { $P(age: a)$ }_{-}{ $P(age: a + 1)$ }
    do tmp: INT. tmp := age(); tmp := tmp+1; set_age(tmp) end

  age: INT
end

```

Figure 3.4: The PERSON class.

3.1.3 Diamond inheritance

Verification of multiple inheritance requires proper handling of data from several parent classes. Diamond inheritance complicates matters because common ancestor fields are shared. This is unproblematic for our proof system, although the abstraction of the shared data is typically lost. Diamond inheritance can moreover require relationships between several abstractions, which this example achieves with axiom clauses.

An axiom clause consists of a name and a predicate. The name identifies the clause and allows subclasses to refine the predicate. We propagate axiom clauses down the hierarchy to save specification overhead: if a class does not list an axiom clause with the same name as one in a parent, then it is assumed to list an identical clause. To avoid ambiguity in the presence of multiple inheritance, we require that if clauses with the same name are present in multiple parents, then they must all be identical. An axiom clause copied in

```

class STUDENT inherit PERSON define
x. $P_{STUDENT}$ (age:a) as x. $P_{PERSON}$ (age:a)
x. $S_{STUDENT}$ (age:a, exm:e) as x. $P_{STUDENT}$ (age:a) * x.exams  $\hookrightarrow$  e
x. $RestStoP_{STUDENT}$ (exm:e) as x.exams  $\hookrightarrow$  e
export  $\forall x,a,e. [x.P_{PERSON}(age:a) * x.RestStoP_{STUDENT}(exm:e)] \Leftrightarrow$ 
      x. $S_{STUDENT}$ (age:a, exm:e) where {}
axiom S_P:  $\forall a, e. S(age:a, exm:e) \Leftrightarrow [P(age:a) * RestStoP(exm:e)]$ 
feature
  introduce STUDENT(a: INT, e: INT)
  dynamic {age  $\hookrightarrow$  _ * exams  $\hookrightarrow$  _} - {S(age:a, exm:e)}
  do Precursor{PERSON}(a); exams := e end

  introduce exams(): INT
  dynamic {S(age:a, exm:e)} - {S(age:a, exm:e) * Result = e}
  do Result := exams end

  introduce take_exam()
  dynamic {S(age:a, exm:e)} - {S(age:a, exm:e + 1)}
  do tmp: INT. tmp := exams; exams := tmp + 1 end

  exams: INT
end

```

Figure 3.5: The STUDENT class.

```

class SMUSICIAN inherit STUDENT MUSICIAN
define
  x.PSMUSICIAN(age: a) as x.PPERSON(age: a)
  x.SSMUSICIAN(age: a, exm: e) as x.SSTUDENT(age: a, exm: e)
  x.MSMUSICIAN(age: a, pfm: p) as x.MMUSICIAN(age: a, pfm: p)
  x.SMSMUSICIAN(age: a, exm: e, pfm: p) as x.PPERSON(age: a) *
    x.RestStoPSTUDENT(exm: e) * x.RestMtoPMUSICIAN(pfm: p)
  x.RestStoPSMUSICIAN(exm: e) as x.RestStoPSTUDENT(exm: e)
  x.RestMtoPSMUSICIAN(pfm: p) as x.RestMtoPMUSICIAN(pfm: p)
  x.RestSMtoSSMUSICIAN(pfm: p) as x.RestMtoPMUSICIAN(pfm: p)
  x.RestSMtoMSMUSICIAN(exm: e) as x.RestStoPSTUDENT(exm: e)
axiom
  SM_S:  $\forall a, e, p. \text{SM}(\text{age}: a, \text{exm}: e, \text{pfm}: p) \Leftrightarrow$ 
     $[\text{S}(\text{age}: a, \text{exm}: e) * \text{RestSMtoS}(\text{pfm}: p)]$ 
  SM_M:  $\forall a, e, p. \text{SM}(\text{age}: a, \text{exm}: e, \text{pfm}: p) \Leftrightarrow$ 
     $[\text{M}(\text{age}: a, \text{pfm}: e) * \text{RestSMtoM}(\text{exm}: e)]$ 
feature
  introduce SMUSICIAN(a: INT, e: INT, p: INT)
  dynamic {age  $\hookrightarrow$  _ * exams  $\hookrightarrow$  _ * performances  $\hookrightarrow$  _}
    {SM(age: a, exm: e, pfm: p)}
  do Precursor{STUDENT}(a,e); Precursor{MUSICIAN}(a,p) end

  introduce do_exam_performance()
  static {SM(age: a, exm: e, pfm: p)}-{SM(age: a, exm: e + 1, pfm: p + 1)}
  do take_exam(); perform() end
end

```

Figure 3.6: The SMUSICIAN class.

this way is not refined in the subclass and automatically consistent with its parent versions. In the general case where a subclass refines an axiom clause, Parent consistency must be proven as indicated in the later formalisation.

The focus of this example is class SMUSICIAN, shown in Figure 3.6. It inherits from STUDENT and MUSICIAN, both which inherit from PERSON. The STUDENT and PERSON classes are shown in Figures 3.4 and 3.5; MUSICIAN is similar to STUDENT and not shown. A diamond is formed with PERSON at the top, and an instance of SMUSICIAN has one ‘age’ field, one *set_age* method, etc. under shared multiple inheritance semantics. The classes use axiom clauses to specify relationships between abstractions P , S , M and SM .

Since SMUSICIAN is non-abstract, we must prove that axiom SM_S holds for its direct instances. This proof obligation for axiom clauses is called Implication in the formalisation that will follow. It holds indeed, since under the standard apf assumptions of SMUSICIAN, exported information of all classes, and the assumption **Current** : SMUSICIAN, we have:

```

SM(age: a, exm: e, pfm: p)
⇔ // Standard apf assumptions, Current : SMUSICIAN
SMSMUSICIAN(age: a, exm: e, pfm: p)
⇔ // Standard apf assumptions.
PPERSON(age: a) * RestStoPSTUDENT(exm: e) *
  RestMtoPMUSICIAN(pfm: p)
⇔ // Exported information from STUDENT.
SSTUDENT(age: a, exm: e) * RestMtoPMUSICIAN(pfm: p)
⇔ // Standard apf assumptions.
SSMUSICIAN(age: a, exm: e) * RestSMtoSSMUSICIAN(pfm: p)
⇔ // Standard apf assumptions, Current : SMUSICIAN
S(age: a, exm: e) * RestSMtoS(pfm: p)

```

The export clause in STUDENT is not closely connected to multiple inheritance. In fact, any class C inheriting from STUDENT which defines P_C to be P_{PERSON} and S_C to be $S_{STUDENT}$ will need the export clause to prove Implication of S_P , which we omit here for SMUSICIAN. Axiom verification frequently requires exported information of this kind. What is vital about the export clause in the shared multiple inheritance setting is that it isolates the shared ancestor state of SMUSICIAN, namely P_{PERSON} . This allows SMUSICIAN to relate ancestor abstractions in a fairly abstract way. Only the constructor’s Body verification proof needs the export clause in PERSON³:

³Unless a class manipulates fields of its ancestors directly, this export clause would not be needed in languages where constructors of common ancestor classes cannot be called more than once.

$$\begin{aligned}
& \{\text{age} \hookrightarrow _ * \text{exams} \hookrightarrow _ * \text{performances} \hookrightarrow _ \} \\
& \quad \mathbf{Precursor}\{\text{STUDENT}\}(a, e) \\
& \{\text{S}_{\text{STUDENT}}(\text{age}: a, \text{exm}: e) * \text{performances} \hookrightarrow _ \} \\
& \{\text{P}_{\text{PERSON}}(\text{age}: a) * \text{RestStoP}_{\text{STUDENT}}(\text{exm}: e) * \text{performances} \hookrightarrow _ \} \\
& \{\text{age} \hookrightarrow a * \text{RestStoP}_{\text{STUDENT}}(\text{exm}: e) * \text{performances} \hookrightarrow _ \} \\
& \quad \mathbf{Precursor}\{\text{MUSICIAN}\}(a, p) \\
& \{\text{M}_{\text{MUSICIAN}}(\text{age}: a, \text{pfm}: p) * \text{RestStoP}_{\text{STUDENT}}(\text{exm}: e) \} \\
& \{\text{P}_{\text{PERSON}}(\text{age}: a) * \text{RestMtoP}_{\text{MUSICIAN}}(\text{pfm}: p) * \\
& \quad \text{RestStoP}_{\text{STUDENT}}(\text{exm}: e) \} \\
& \{\text{SM}_{\text{SMUSICIAN}}(\text{age}: a, \text{exm}: e, \text{pfm}: p) \}
\end{aligned}$$

Note that class `SMUSICIAN` would not have needed exported information if it ignored the parent constructors and simply overrode everything. The same is true for proof systems with less abstraction where method bodies are reverified in subclasses.

Since **Current** in `SMUSICIAN` will always reference an object whose dynamic type is a subtype of `SMUSICIAN`, the Body verification proof of *do_exam_performance* can use axiom information to infer **SM**-specs for *take_exam* and *perform*:

$$\begin{aligned}
& \{\text{SM}(\text{age}: a, \text{exm}: e, \text{pfm}: p) \} \\
& \{\text{S}(\text{age}: a, \text{exm}: e) * \text{RestSMtoS}(\text{pfm}: p) \} \\
& \quad \text{take_exam}() \\
& \{\text{S}(\text{age}: a, \text{exm}: e + 1) * \text{RestSMtoS}(\text{pfm}: p) \} \\
& \{\text{SM}(\text{age}: a, \text{exm}: e + 1, \text{pfm}: p) \} \\
& \{\text{M}(\text{age}: a, \text{pfm}: p) * \text{RestSMtoM}(\text{exm}: e + 1) \} \\
& \quad \text{perform}() \\
& \{\text{M}(\text{age}: a, \text{pfm}: p + 1) * \text{RestSMtoM}(\text{exm}: e + 1) \} \\
& \{\text{SM}(\text{age}: a, \text{exm}: e + 1, \text{pfm}: p + 1) \}
\end{aligned}$$

The specification overhead incurred by axiom clauses is offset by specification inference gains: **SM**, **S** and **M** specifications can be inferred for *age*, *set_age* and *celebrate_birthday*, while **SM** specifications can be inferred for *exams*, *take_exam*, *performances* and *get_performance* – a total of 13 specifications for methods of `SMUSICIAN`. These inferred specifications are guaranteed to be implemented by all subclasses, and no dynamic type information is needed to use them⁴. Yet the system is still flexible – a subclass can always choose

⁴The technique used by Chin et al. [14] of inheriting static method specifications and deriving dynamic specifications from them implements “internal specification inference”, i.e. a class infers and publishes dynamic specifications for its methods which external clients can use. In contrast to this, the technique of equipping classes with export/axiom clauses implements “external specification inference”, i.e. clients infer specifications for methods based on published export/axiom information. A benefit of the

to satisfy such constraints vacuously by defining selected apf entries as false. Class DCell in [58] provides an example of this.

3.2 MultiStar

This section sketches notable aspects of the MultiStar implementation. MultiStar has a two-tier architecture: a front-end that translates Eiffel programs and specifications into a simpler form for verification, and a language-independent back-end based on jStar [23] which implements our proof system.

3.2.1 Front-end

The front-end provides a graphical user interface within the EVE integrated development environment, and is part of the standard EVE download [28]. It translates Eiffel code and specifications into the back-end's input format, and provides access to verification results. Verification is triggered by picking and dropping an annotated class on the MultiStar tool. Class annotations consist of apf entry definitions, export/axiom clauses and method specifications.

To simplify the proofs and formalisation in this chapter, constructor preconditions explicitly mention fields and break information hiding. The front-end translation of MultiStar injects them automatically. For example, a user would write the specification of CCELL's constructor as

dynamic {true} - {Cc(val: v, cnt: 0)}

instead of

dynamic {value $\hookrightarrow$ _ * count $\hookrightarrow$ _} - {Cc(val: v, cnt: 0)}

In detail, the front-end facilitates this by:

1. Using jStar's **new** statement, whose specification is given by the triple {true} $x := \mathbf{new} C\{x : C\}$. This allows us to omit the fields in the dynamic precondition⁵.
2. Adding all fields (including ancestor ones) to the static precondition when checking Body verification, and consuming all fields of a parent

external approach is that clients can infer valid specifications for library code without re-verifying it. For example, knowing only the inferred specifications for methods of class SMUSICIAN is not enough for proving the triple {x.SM(age: a, exm: e, pfm: p)} use_student(x) {x.SM(age: a + 1, exm: e + 4, pfm: p)} without looking at the implementation of the library routine use_student(st: STUDENT) whose dynamic specification is {st.S(age: a, exm: e)} - {st.S(age: a + 1, exm: e + 4)}.

⁵The dynamic specification of the constructor is used for object initialisation. $x := \mathbf{new} C(\bar{e})$ abbreviates $x := \mathbf{new} C$; $x.C(\bar{e})$ if x is not free in $\bar{e}$.

class and its ancestors right before the parent constructor is called. This is communicated to the back-end by emitting special instructions, and allows us to omit the fields in the static precondition.

The Dynamic dispatch proof obligation, which checks that the static and dynamic specifications are consistent with each other, is unaffected because fields are omitted in both static and dynamic preconditions. In languages where no fields are shared by ancestors, or constructors of common ancestors are called only once, the manipulation does not have to add ancestor fields to the static precondition and consume fields when a parent constructor is called. This is the approach jStar uses for Java verification. A front-end for C++ can use a similar technique because every ancestor constructor is called exactly once when virtual base classes are used [27].

3.2.2 Back-end

The MultiStar back-end extends jStar with support for export and axiom clauses, abstract classes and multiple inheritance. The latter two demand generalised proof obligation checking for methods.

Export and axiom clauses

The background theory used by the jStar theorem prover is encoded as a list of sequent rules. A sequent is of the form $P \mid Q \vdash R$, meaning $(P * Q) \Rightarrow (P * R)$. Each sequent rule has the form

$A \mid B \vdash C$
if
 $D \mid E \vdash F$

If the prover is trying to prove a sequent that matches the rule's conclusion $A \mid B \vdash C$, it suffices to prove the premise sequent $D \mid E \vdash F$ where the appropriate substitutions due to matching have been made. A new proof goal is thus obtained, and the proof is complete when the goal is of the form $G \mid H \vdash \cdot$. For details the reader is referred to [23].

Exported information is written as sets of implications. Before verifying an export clause, the background theory is temporarily extended with the definitions of all predicates in its **where** part. For each definition of the form $w(x) = P$, the following two rules are generated:

$\mid w(x) \vdash$ if $\mid P \vdash$	$\mid \vdash w(x)$ if $\mid \vdash P$
--	--

After all exported implications in the clause have been checked, the definitions are removed from the background theory. After all export clauses have been verified, each exported implication $P \Rightarrow (Q_1 * \dots * Q_n)$ is added to the background theory as a set of n rules, where rule $i \in 1..n$ has the form

$$\begin{array}{l} | P \vdash Q_i \\ \text{if} \\ Q_i \mid Q_1 * \dots * Q_{i-1} * Q_{i+1} * \dots * Q_n \vdash \end{array}$$

This rule form retains information about Q_i in its premise, and removal of Q_i from the goal sequent's right-hand side brings the proof closer to completion.

The background theory augmented with export information is then used to verify axiom clauses. The predicates in axiom clauses are written as implications. After all axiom clauses have been verified, an axiom implication $P \Rightarrow (Q_1 * \dots * Q_n)$ written in class C is encoded as n rules, with rule $i \in 1..n$ of the form

$$\begin{array}{l} | P \vdash Q_i \\ \text{if} \\ Q_i \mid Q_1 * \dots * Q_{i-1} * Q_{i+1} * \dots * Q_n \vdash x <: C \end{array}$$

where x is the pattern variable substituted for **Current**.

The background theory augmented with export and axiom information is then used for method verification.

Method proof obligations

The back-end accommodates abstract classes and abstract methods in addition to shared multiple inheritance. An abstract method has no body and hence no static specification. The back-end takes this into account when expanding specification shorthands. After shorthand expansion, verification of method m in class C proceeds as follows (the formalisation contains details about the proof obligations):

- If m has a static specification and C can be instantiated (i.e. is non-abstract), then check Dynamic dispatch.
- If m has a body in C , then check Body verification.
- Always check Behavioural subtyping. This succeeds trivially if m is introduced in C : the set of dynamic specifications for m in C 's parents is empty, and therefore all its elements are preserved by the new specification.
- If m has a static specification but no body in C , then check Inheritance.




Figure 3.7: The core Gobo cursor hierarchy.

The treatment subsumes interface inheritance – interfaces are treated as abstract classes with only abstract methods and no fields.

3.3 Case study

The Gobo data structure library [29] is an open-source Eiffel library covering data structures and algorithms. It contains classic data structures such as lists, stacks and sets, and provides several implementations of each structure. The library is stable and a popular choice among Eiffel developers.

Data structures such as lists and sets can be traversed with iterators. The iterator (or *cursor*) hierarchy is characterised by relatively simple algorithms and extensive use of multiple inheritance, which makes it an ideal candidate for evaluating the novel aspects of our proof system and its implementation. The core classes are shown in Figure 3.7: a `LINEAR_CURSOR` can traverse a data structure forwards, a `BILINEAR_CURSOR` can traverse both forwards and backwards, an `INDEXED_CURSOR` offers random data structure access with an integer position or index, and a `DYNAMIC_CURSOR` can modify the data structure being traversed.

We successfully verified the core cursor hierarchy of Figure 3.7 with MultiStar. The overall effort for specification and verification was five person-days. Most of the time was spent on finding and revising specifications, since we did not modify the code. Table 3.1 shows the experimental results. The total time taken by MultiStar is reported, which includes translating Eiffel code, expanding specification shorthands and checking all proof obligations.

Since iterators rely on properties of the data structures (containers) they traverse, we annotated the container classes with the required specifications. Particularly interesting are the axiom clauses that iterators demand.

Class	LOC ₁	LOC ₂	Time(s)
BILINEAR_CURSOR	99	124	1.306
BILINEAR_SET_CURSOR	44	50	0.841
CURSOR	130	158	1.039
DYNAMIC_CURSOR	50	66	1.070
INDEXED_CURSOR	46	57	0.698
LINEAR_CURSOR	98	123	1.327
LIST_CURSOR	238	271	1.643
SET_CURSOR	38	44	0.738
8 classes	743	893	8.662

Table 3.1: Experimental results of the Gobo iterator case study. LOC₁ and LOC₂ denote the lines of code before and after specification respectively. MultiStar was executed on a 2.53 GHz Intel Core 2 Duo with 4 GB RAM.

Consider for example the simplified extract of DYNAMIC_CURSOR in Figure 3.8. The [G] denotes that DYNAMIC_CURSOR has a generic parameter G. Method *swap* takes another cursor referencing the same container, and additionally requires that there are data elements (items) at both cursor positions (the cursors are not ‘off’). The Body verification proof of *swap* uses several properties of containers that can be expressed as axioms, including the following one:

$$\begin{aligned}
& \forall r1, r2, iter1, iter2, i, c1, c2, c3 \cdot \\
& [\text{ItemAt}(\text{res}:r1, \text{ref}:iter1, \text{iters}:i, \text{content}:c1) * \\
& \quad \text{ItemAt}(\text{res}:r2, \text{ref}:iter2, \text{iters}:i, \text{content}:c1) * \\
& \quad \text{Replaced}(\text{ref}:iter1, \text{value}:r2, \text{newcontent}:c2, \text{oldcontent}:c1, \text{iters}:i) * \\
& \quad \text{Replaced}(\text{ref}:iter2, \text{value}:r1, \text{newcontent}:c3, \text{oldcontent}:c2, \text{iters}:i)] \\
& \Rightarrow \\
& \text{Swapped}(\text{ref1}:iter1, \text{ref2}:iter2, \text{iters}:i, \text{oldcontent}:c1, \text{newcontent}:c3)
\end{aligned}$$

This invariant property relates the [ItemAt](#), [Replaced](#) and [Swapped](#) abstractions. Informally, it states that if a data structure has items *r1* and *r2* at iterators *iter1* and *iter2*, and we replace the item at *iter1* with *r2* and the item at *iter2* with *r1*, then the resulting data structure has the same contents as the original one except that the items at *iter1* and *iter2* have been swapped. The example involves neither multiple inheritance nor splitting of superclass and subclass state, and illustrates the generality and expressive power of axiom clauses.

The complete specifications and code of the case study are included in the download [28].

```

abstract class DYNAMIC_CURSOR [G] inherit CURSOR [G]
feature
  introduce abstract replace(v: G) dynamic
  { Cursor(ds: d) * d.DS(content: c1, iters: i) *
    d.IsOff(res: False, ref: Current, iters: i, content: c1) }_
  { Cursor(ds: d) * d.DS(content: c2, iters: i) *
    d.Replaced(ref: Current, value: v, newcontent: c2, oldcontent: c1, iters: i) }

  inherit item(): G static
  { Cursor(ds: d) * d.DS(content: c, iters: i) *
    d.IsOff(res: False, ref: Current, iters: i, content: c) }_
  { Cursor(ds: d) * d.DS(content: c, iters: i) *
    d.ItemAt(res: Result, ref: Current, iters: i, content: c) }

  introduce swap(other: DYNAMIC_CURSOR [G]) static
  { Cursor(ds: d) * d.DS(content: c1, iters: i) * other.Cursor(ds: d) *
    d.IsOff(res: False, ref: Current, iters: i, content: c1) *
    d.IsOff(res: False, ref: other, iters: i, content: c1) }_
  { Cursor(ds: d) * d.DS(content: c2, iters: i) * other.Cursor(ds: d) *
    d.Swapped(ref1: Current, ref2: other, iters: i, oldcontent: c1, newcontent: c2) }
do
  v: G. w: G.
  v := item(); w := other.item();
  replace(w); replace(v)
end
end

```

Figure 3.8: A simplified extract of DYNAMIC_CURSOR

3.4 Formalisation

This section contains a formal presentation of the programming language and its proof system. It extends the material of chapter 2 with abstract classes, multiple inheritance, and export and axiom specifications. To keep things simple, the presentation concentrates on the new extensions.

3.4.1 *Language syntax*

The grammar of the extended kernel language appears in Figure 3.9. New and changed productions are highlighted: abstract classes, multiple parents, export and axiom clauses, and abstract methods are now supported. The letter *p* ranges over apf names, *w* over auxiliary predicate names, and *a* over axiom names.

We assume the absence of clashes when names are introduced. This guarantees that axioms, methods and fields with the same name in parent classes stem from common ancestors.

The shared semantics of multiple inheritance is used, which is popular in Eiffel [26] and known as inheritance with *virtual base classes* in C++ [27]. Common ancestor fields are shared, and method overriding overrides all ancestor versions. To avoid ambiguity, a class can inherit a method only if its body (if there is one) is the same along all inheritance paths. Direct method calls can encode language mechanisms which allow a particular ancestor implementation to be chosen, so no generality is lost.

3.4.2 *Operational semantics*

The shared semantics of multiple inheritance ensures that 1) only dynamic type information is needed at runtime (in contrast to what ‘select’ clauses of Eiffel’s replicated inheritance demand), and 2) the usual semantics of casts can be adopted (in contrast to replicated inheritance in C++, where casting can change pointer values [27]).

The operational semantics is therefore similar to Parkinson’s semantics of Java [57]. Configurations contain a stack, a heap and a sequence of statements under execution. The stack maps variables to values which include object ids. The heap maps object ids to records containing a dynamic type *G* and field-value mappings.

L ::= Ab class G inherit $\bar{H}$ define $\bar{D}$ export $\bar{E}$ axiom $\bar{A}$ feature $\bar{M}$ $\bar{F}$ end	
Ab ::= abstract ϵ	
D ::= x.p_G(t:y) as P	<i>Define clause</i>
E ::= P where {$\bar{W}$}	<i>Export clause</i>
W ::= w($\bar{x}$) = P	<i>Where clause</i>
A ::= a: P	<i>Axiom clause</i>
M ::= introduce m(Args) Rt Sd Ss B	<i>Method declaration</i>
override m(Args) Rt Sd Ss B	
inherit m(Args) Rt Sd Ss	
introduce abstract m(Args) Rt Sd	
inherit abstract m(Args) Rt Sd	
F ::= f: Type	<i>Field declaration</i>
Sd ::= dynamic Spec	<i>Dynamic specification</i>
Ss ::= static Spec	<i>Static specification</i>
Spec ::= {P}-{Q} {P}-{Q} also Spec	<i>Specification</i>
B ::= do s end	<i>Method body</i>
s ::= skip	<i>Statement</i>
x := e	<i>Assignment</i>
s ; s	<i>Sequential composition</i>
if e then s else s end	<i>Conditional</i>
while e do s end	<i>Loop</i>
x: Type. s	<i>Local variable block</i>
x := new G	<i>Object allocation</i>
x := y.f	<i>Field lookup</i>
x.f := e	<i>Field update</i>
x := y.m($\bar{e}$) y.m($\bar{e}$)	<i>Dynamically dispatched call</i>
x := y.G::m($\bar{e}$) y.G::m($\bar{e}$)	<i>Direct method call</i>
e ::= x e + e e = e Void 0 1 2 ...	<i>Expression</i>
Type ::= INT BOOL G	
Args ::= $\overline{x: \text{Type}}$	<i>Formal arguments</i>
Rt ::= ϵ : Type	<i>Return type</i>

Figure 3.9: The grammar of the extended kernel language.

3.4.3 Logic syntax and semantics

The assertion language is extended with auxiliary predicates. These predicates are defined in the **where** part of export clauses.

$$\begin{aligned}
 P, Q, S, T, \Delta ::= & \forall x.P \mid P \Rightarrow Q \mid \text{false} \mid e = e' \mid x : G \mid x <: G \mid x.f \hookrightarrow e \mid P * Q \\
 & \mid x.p(\overline{t:e}) \quad \text{Apf predicate} \\
 & \mid x.p_G(\overline{t:e}) \quad \text{Apf entry} \\
 & \mid w(\overline{x}) \quad \text{Auxiliary predicate}
 \end{aligned}$$

3.4.4 Specification refinement

The formalisation of specification refinement of Section 2.3 remains valid here. However, the assumptions Δ now contain the standard apf assumptions of a class as well as export and axiom information of all other classes.

3.4.5 The specification environment

The specification environment Γ is extended to map axiom names to their associated predicates for all classes in a program:

$$\begin{aligned}
 \Gamma ::= & G.a \mapsto P \quad \text{Axiom specification} \\
 & \mid G.m \mapsto (\overline{x}, \{P\} - \{Q\}) \quad \text{Method dynamic specification} \\
 & \mid G::m \mapsto (\overline{x}, \{S\} - \{T\}) \quad \text{Method static specification} \\
 & \mid \overline{\Gamma}
 \end{aligned}$$

For well-formed programs, Γ is still guaranteed to be a partial function, and we write $\Gamma(G.a) = P$ for $G.a \mapsto P \in \Gamma$.

3.4.6 Export information verification

A class can make information about itself available to other classes in an export clause. Export clauses are frequently used to specify relationships between apfs or their entries, and to expose apf entry definitions. Information can be hidden in predicates defined after the keyword **where**: the definitions are not exported, so other classes must treat these predicates abstractly.

Export information must be verified since other classes use it for reasoning. Under the predicate definitions following **where**, the assumptions about a class must imply exported information. This is captured by the following proof rule:

$$\frac{[\Delta \wedge (\forall \overline{x}_1 \cdot w_1(\overline{x}_1) \Leftrightarrow Q_1) \wedge \dots \wedge (\forall \overline{x}_n \cdot w_n(\overline{x}_n) \Leftrightarrow Q_n)] \Rightarrow P}{\Delta \vdash_e P \text{ where } \{w_1(\overline{x}_1) = Q_1; \dots; w_n(\overline{x}_n) = Q_n\}}$$

Since the assumptions about a class do not include assumptions about the exact dynamic type of **Current**, export information can be used to verify axioms and methods of other classes in a program.

3.4.7 Axiom verification

Information about a class and all its subclasses can be made available in an axiom clause. This knowledge can be used later to verify methods. To simplify the treatment, we require that a class explicitly lists all axiom clauses applicable to it (in MultiStar and the examples, specification shorthands achieve this).

In the rule for axiom verification, the assumptions Δ include information about class G and export information from all other classes. A subclass must preserve all axioms of its parents and may refine the predicate associated with an axiom name (the Parent consistency [P.c.] obligation). A non-abstract class must also show that the predicate holds for its direct instances (the Implication [Imp.] obligation).

$$\begin{array}{c}
 \forall i \in I. \Gamma(H_i.a) = Q_i \quad \wedge \quad \forall j \in (1..n \setminus I). H_j.a \notin \text{dom}(\Gamma) \\
 (\Delta \wedge P) \Rightarrow \bigwedge_{i \in I} Q_i \quad \text{[P.c.]} \\
 \text{Ab} \neq \epsilon \quad \vee \quad (\Delta \wedge \mathbf{Current} : G) \Rightarrow P \quad \text{[Imp.]} \\
 \hline
 \Delta; \Gamma \vdash_a a: P \text{ in Ab } G \text{ parents } H_1 \dots H_n
 \end{array}$$

3.4.8 Statement verification

The assumptions Δ used to verify statements contain information about the enclosing class as well as export and axiom information from all other classes. The rules remain the same as those in Section 2.5.

3.4.9 Method verification

The rules for method verification of Section 2.6 are extended here to deal with multiple inheritance and abstract classes and methods. As for statement verification, the assumptions Δ used to verify method definitions contain information about the method's enclosing class as well as export and axiom information of all other classes.

A newly introduced method's static and dynamic specifications must be consistent if the class is not abstract:

$$\begin{array}{c}
B = \mathbf{do\ s\ end} \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
Ab \neq \epsilon \vee \Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\Delta; \Gamma \vdash_s \{S_G\}s\{T_G\} \quad [\text{B.v.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{introduce}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ B\ \text{in}\ Ab\ G\ \text{parents}\ \overline{H}
\end{array}$$

An abstract method can be introduced without any proof obligations, since there is only a dynamic specification and no method body.

$$\Delta; \Gamma \vdash_m \mathbf{introduce\ abstract}\ m(\text{Args})\ \text{Rt}\ Sd\ \text{in}\ Ab\ G\ \text{parents}\ \overline{H}$$

The next rule is used whenever an abstract method is implemented or a method body is redefined. The $H_1 \dots H_n$ are the immediate superclasses of G .

$$\begin{array}{c}
\forall i \in I. \Gamma(H_i.m) = (\overline{x}, \{P_{H_i}\}\text{-}\{Q_{H_i}\}) \\
\forall j \in (1..n \setminus I). H_j.m \notin \text{dom}(\Gamma) \\
B = \mathbf{do\ s\ end} \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
\Delta \vdash \{P_G\}\text{-}\{Q_G\} \implies (\mathbf{also}_{i \in I} \{P_{H_i}\}\text{-}\{Q_{H_i}\}) \quad [\text{B.s.}] \\
Ab \neq \epsilon \vee \Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\Delta; \Gamma \vdash_s \{S_G\}s\{T_G\} \quad [\text{B.v.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{override}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ B\ \text{in}\ Ab\ G\ \text{parents}\ H_1 \dots H_n
\end{array}$$

When a non-abstract method is inherited, its static specification must follow from those in parents. The Inheritance [Inh.] obligation enforces it.

$$\begin{array}{c}
\forall i \in I. \Gamma(H_i.m) = (\overline{x}, \{P_{H_i}\}\text{-}\{Q_{H_i}\}) \\
\forall k \in (1..n \setminus I). H_k.m \notin \text{dom}(\Gamma) \\
\forall j \in J. \Gamma(H_j::m) = (\overline{x}, \{S_{H_j}\}\text{-}\{T_{H_j}\}) \\
\forall l \in (1..n \setminus J). H_l::m \notin \text{dom}(\Gamma) \\
Sd = \mathbf{dynamic}\ \{P_G\}\text{-}\{Q_G\} \\
Ss = \mathbf{static}\ \{S_G\}\text{-}\{T_G\} \\
\Delta \vdash \{P_G\}\text{-}\{Q_G\} \implies (\mathbf{also}_{i \in I} \{P_{H_i}\}\text{-}\{Q_{H_i}\}) \quad [\text{B.s.}] \\
\Delta \vdash (\mathbf{also}_{j \in J} \{S_{H_j}\}\text{-}\{T_{H_j}\}) \implies \{S_G\}\text{-}\{T_G\} \quad [\text{Inh.}] \\
Ab \neq \epsilon \vee \Delta \vdash \{S_G\}\text{-}\{T_G\} \xRightarrow{\mathbf{Current}:G} \{P_G\}\text{-}\{Q_G\} \quad [\text{D.d.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{inherit}\ m(\text{Args})\ \text{Rt}\ Sd\ Ss\ \text{in}\ Ab\ G\ \text{parents}\ H_1 \dots H_n
\end{array}$$

The next rule applies whenever an abstract method is inherited or a non-abstract method is inherited and made abstract. Such a method has no static specification, so only the consistency of its dynamic specification w.r.t. those in parent classes is required with the Behavioural subtyping proof obligation.

$$\begin{array}{l}
\forall i \in I. \Gamma(H_i.m) = (\bar{x}, \{P_{H_i}\} - \{Q_{H_i}\}) \\
\forall j \in (1..n \setminus I). H_j.m \notin \text{dom}(\Gamma) \\
Sd = \mathbf{dynamic} \{P_G\} - \{Q_G\} \\
\Delta \vdash \{P_G\} - \{Q_G\} \implies (\mathbf{also}_{i \in I} \{P_{H_i}\} - \{Q_{H_i}\}) \quad [\text{B.s.}] \\
\hline
\Delta; \Gamma \vdash_m \mathbf{inherit abstract} \ m(\text{Args}) \text{ Rt } Sd \text{ in Ab } G \text{ parents } H_1 \dots H_n
\end{array}$$

Several intuitive relationships hold between the rules. For example, introducing a non-abstract method is the same as overriding the version of parent classes $I = \emptyset$. Likewise, introducing an abstract method is the same as inheriting it from the parent set $I = \emptyset$. Section 3.2.2 gives yet another perspective the rationale behind the proof obligations of the various rules.

3.4.10 Class and program verification

For class verification, different assumptions are used to verify the various class sections. The formula Δ_{APF} contains class-specific information and is used to verify export clauses. The assumptions Δ_E contain export information from all classes, and are used together with Δ_{APF} to verify axioms. The formula Δ_A contains axiom information of all classes, and is used with Δ_{APF} and Δ_E in method definition verification.

$$\begin{array}{l}
\forall E_i \in \bar{E}. \Delta_{APF} \vdash_e E_i \\
\forall A_i \in \bar{A}. (\Delta_{APF} \wedge \Delta_E); \Gamma \vdash_a A_i \text{ in Ab } G \text{ parents } \bar{H} \\
\forall M_i \in \bar{M}. (\Delta_{APF} \wedge \Delta_E \wedge \Delta_A); \Gamma \vdash_m M_i \text{ in Ab } G \text{ parents } \bar{H} \\
\hline
\Delta_{APF}, \Delta_E, \Delta_A; \Gamma \vdash_c \\
\text{Ab class } G \mathbf{inherit} \ \bar{H} \mathbf{define} \ \bar{D} \mathbf{export} \ \bar{E} \mathbf{axiom} \ \bar{A} \mathbf{feature} \ \bar{M} \ \bar{F} \mathbf{end}
\end{array}$$

Finally, here is the rule for program verification:

$$\begin{array}{l}
\forall i \in 1..n. L_i = \dots \mathbf{class} \ G_i \dots \mathbf{export} \ \bar{E}_i \mathbf{axiom} \ \bar{A}_i \mathbf{feature} \dots \mathbf{end} \\
\Delta_E = \bigwedge_{i \in 1..n} \bigwedge_{E_{i_k} \in \bar{E}_i} \text{exportinfo}(E_{i_k}) \\
\Delta_A = \bigwedge_{i \in 1..n} \bigwedge_{A_{i_k} \in \bar{A}_i} \text{axiominfo}(G_i, A_{i_k}) \\
\Gamma = \text{specs}(L_1 \dots L_n) \\
\forall i \in 1..n. \text{apf}(L_i), \Delta_E, \Delta_A; \Gamma \vdash_c L_i \\
\Delta_E \wedge \Delta_A; \Gamma \vdash_s \{\text{true}\} \{\text{true}\} \\
\hline
\vdash_p L_1 \dots L_n \ s
\end{array}$$

$$\begin{aligned} \text{exportinfo}(P \textbf{ where } \dots) &\stackrel{\text{def}}{=} P \\ \text{axiominfo}(G, a: P) &\stackrel{\text{def}}{=} \forall x <: G \cdot P[x/\mathbf{Current}], \text{ where } x \text{ is fresh.} \end{aligned}$$

Predicate definitions following the **where** keyword are hidden by *exportinfo*, and the definition of *axiominfo* reflects the fact that subclasses preserve axioms.

The function *apf* translates the abstract predicate family definitions of a class into its standard apf assumptions in the same way as before.

MultiStar and the examples assume that every class implicitly exports tag reduction information. In other words, for every entry $(x.p_G(Y) \textbf{ as } P)$ in the **define** section of a class G , $(\forall x <: G \cdot TR(p, x, Y) \textbf{ where } \{\})$ is implicitly exported.

Theorem. The program verification rule is sound. (The proof, sketched in Appendix A, depends on the layered assumption structure of export and axiom clauses that avoids circularity in reasoning⁶.)

3.5 Conclusions and related work

The presented proof system supports two complementary mechanisms that can express relationships between abstractions in the logic. Such relationships are pervasive in OO programs, and facilitate flexible client reasoning, access control, specification inference, and constraints on the implementation of abstractions. Moreover, the system offers a sound way to verify various forms and uses of shared multiple inheritance. By virtue of extending Parkinson and Bierman’s system, the examples in [58] illustrate that it can also deal with behaviour extension, restriction and modification, as well as representation replacement in subclasses. It is modular and every method body is verified only once. MultiStar implements these features in an automatic tool that, as the Gobo case study shows, holds good promise for verifying real-world software.

We are not aware of any other proof system or tool that can verify our examples and case study. Nevertheless, there are many relationships with other work:

⁶The layered assumption structure does not rule out axioms that depend on other axioms. If we want an axiom Q in class C_2 that depends on axiom P in class C_1 , we can write the axiom $(\forall x <: C_1 \cdot P[x/\mathbf{Current}]) \Rightarrow Q$ in class C_2 instead, where x is a fresh variable. After axiom verification, method and statement verification can use $(\forall y <: C_2 \cdot Q[y/\mathbf{Current}])$ where y is fresh. The proof system enables the specification and verification of invariants for aggregate structures involving multiple objects of different types.

Class invariants	Axiom clauses
Related to object consistency	No notion of object consistency
Hold at particular program points	Hold everywhere
Expressed i.t.o. fields and pure method calls	Expressed as logical predicates
Constrain operations	Constrain logical abstractions of data
Verification involves methods	Verification cannot involve methods

Table 3.2: The main differences between class invariants and axiom clauses.

Axiom clauses We do not know of any existing specification mechanisms that are closely related to axiom clauses, in the sense that they describe and enforce relationships between logical abstractions of data in a similar way for inheritance hierarchies.

Class invariants form the basis of several OO specification and verification approaches, including Spec# [3] and JML [40]. Two main flavours of class invariants exist: private invariants, as exemplified by the object invariants of Spec#, and public invariants [41], which include JML’s derived invariants and the invariants of Jacobs and Piessens for describing relationships between inspector methods [33]. Class invariants, like axiom clauses, constrain subclasses. However, there are several important differences between them. Class invariants either define exactly when an object is consistent (private invariants), or describe abstract properties of consistent objects (public invariants). Axiom clauses have no notion of object consistency. Class invariants are expected to hold at particular points in a program and may be broken at others, according to the employed invariant protocol [25]. Axiom clauses are true invariants in the sense that they hold everywhere. The relationships they describe cannot be violated by assignment statements, and hence there are no problems with e.g. method callbacks [44, 39]. Class invariants are expressed in terms of fields (including model fields) and pure method calls. Public invariants constrain operations that must establish or preserve them. Private invariants constrain the pure methods they use and indirectly also other operations that depend on the private invariant. Axioms are expressed as logical predicates, and they constrain logical abstractions of data. The verification of private invariants involves the inspection of method bodies (for the proof obligations see e.g. [45]), and the verification of public invariants involves private invariants and the specifications and/or bodies of pure methods. Axiom clauses are verified prior to methods and do not depend on methods in any way. The main differences between class invariants and axiom clauses are summarised in Table 3.2.

Export clauses There is some correspondence between the class axioms of Kassios [37, 36] and apfs/export clauses. Class axioms can be used to axiomatise the specification and program attributes of a class. In a class implementation, class axioms typically describe abstract state (represented by specification attributes) as a function of the concrete implementation state (represented by program attributes). This corresponds loosely to the way apf entry definitions relate the concrete state to apf arguments. Furthermore, class axioms can describe consequences of a specification attribute such as a class invariant, which typically include framing properties and relationships between other specification attributes. This corresponds somewhat to export clauses that describe properties of apf arguments or relationships between them. The framework of class axioms does not include inheritance, so despite the similar names, it is appropriate to compare them with export and not axiom clauses. Class axioms can also define method specification/implementations, which apfs and export clauses are incapable of.

In jStar [23], the modifier ‘export’ can be added to an apf entry definition to expose it to all other classes. Even though our export clauses are more general and flexible than this mechanism, MultiStar supports it as a useful shorthand.

The rules for lossless casting by Chin et al. [14] describe relationships between predicates that provide full and partial views of objects. A view predicate describes the contents of the fields of an object directly: a full view of object *o* provides full knowledge of all *o*’s fields, while a partial view with respect to class *C* describes only values of fields introduced by *C* and its ancestors. View predicates and relationships between them are generated automatically. The relationships do not have to be verified and do not constrain subclasses.

Krishnaswami et al. use so-called ‘static specifications’ in [38] to specify relationships between abstract predicates. Although not presented in an OO context, these relationships must be satisfied by implementations and are thus related to our export clauses.

The lemma functions of VeriFast [34] record proofs of relationships between predicates. The relationships are then used in reasoning; the proof of the Composite pattern in [35] provides a good example. Lemma functions, like export clauses, do not constrain subclasses.

Multiple inheritance Surprisingly few systems exist for reasoning about multiple inheritance. The system in [46] also uses separation logic, but without abstraction mechanisms such as apfs. Most of the paper is devoted to elementary separation logic proof rules that also apply in a single-inheritance

context. Diamond inheritance is never treated, and the bodies of inherited methods are re-verified in subclasses.

The focus of [24] is on behavioural subtyping. It proposes to verify behavioural subtyping of methods lazily, i.e. only to the extent demanded by client code. Supplier code is then continually re-verified as a client's use of it grows.

The restricted form of interface inheritance is easily handled by our proof system: an interface is simply an abstract class with only abstract methods and no fields. Many verification tools for OO programs, including Spec# [3] and the JML toolset [9], provide support for specifying and verifying interface inheritance. Both Spec# and JML use pure expressions of the programming language for specification, and follow a class invariant-based approach to verification.

CHAPTER 4

VERIFYING EXECUTABLE CONTRACTS

Many conventional OO program specification approaches, including Eiffel [26], Spec# [3] and JML [40], use Boolean expressions of the programming language to specify routine preconditions and postconditions, class invariants, and other assertions which hold at particular points such as loop invariants. Explaining the operational meaning of such executable assertions to programmers is easy, as they already know the meaning of programming language expressions. Even if such specifications are not formally verified, their runtime checking and value for testing provide significant benefits. Their popularity in software development is therefore not surprising.

Verifying that executable specifications will always hold at runtime is not easy. Expressions used for assertions typically include calls to methods which depend on the heap and which may also cause side-effects such as heap mutation. Consider for example a class SLIST which implements a sorted list of integers. The postcondition of method `insert(i)` can be specified as `has(i)`, meaning that the list has or contains an element `i`. Yet `has(i)` might allocate a new iterator, traverse the linked structure, and temporarily affect bookkeeping of active iterators. Translating expressions such as `has(i)` into logical formulae for static verification is hard. Moreover, programming language expressions might have preconditions and are generally not guaranteed to terminate. This further complicates matters for verification.

In contrast to executable specifications, separation logic predicates are not computations. They specify the effects of computations in terms of program states, i.e. the contents of the stack and heap, and not in terms of outcomes of further computations. Proof systems based on separation logic (such as the ones we have seen) are promising for OO verification and can

successfully verify common programming patterns such as the Visitor [23] and Composite [35]. Unfortunately, the semantics of separation logic is not yet widely known by OO programmers.

How separation logic specifications can be used to tame executable ones is the topic of this chapter. The guarantee is simple: if a program satisfies its separation logic specification and the executable assertions are verified with the separation logic ones, then all executable assertions are guaranteed to hold at runtime. We do not verify the program with respect to its executable assertions, we verify the program *and* its executable assertions with respect to a separation logic specification. For example, if a program containing class SLIST satisfies its separation logic specification, and the separation logic postcondition of `insert(i)` is Q , then we are guaranteed that the executable postcondition, `has(i)`, will always hold at runtime if we can prove the *connecting implication* $Q \Rightarrow \text{has}(i)$. A connecting implication connects the world of separation logic predicates with the one of Boolean programming language expressions. It is proved in this case by deriving the Hoare-style triple $\{Q\}v := \text{has}(i)\{Q \wedge v = \text{True}\}$, where v is a fresh variable. A novel notion called *relative purity* is embedded in connecting implications and embraces side-effects in expressions such as `has(i)`. The logical framework depends on the non-faulting semantics of separation logic triples and its $*$ -connective, as we shall see.

We believe that executable and separation logic specifications can be complementary in OO software development. The proposed formal framework accommodates both ordinary programmers and proof experts. Programmers can express their intentions in the form of executable assertions and use runtime checking to identify faults. While software designs are still evolving, it is probably also easier to change executable assertions than more elaborate specifications. We envisage that in a second phase, once the software has stabilised and many faults have been removed, proof experts would annotate the critical parts with separation logic for verification. At this point the executable assertions do not have to be discarded: our framework integrates the specification approaches and can verify whether the expectations recorded in executable assertions are fulfilled. Problems in the verification of executable assertions can indicate discrepancies between the two types of specifications, e.g. a misunderstanding by the separation logic specifier as to when a routine may be called. Executable and separation logic specifications can therefore complement each other even in verification.

This chapter makes several contributions in the area of OO specification and verification:

1. It describes a simple technique based on connecting implications for

verifying executable preconditions and postconditions. It gives executable assertions a semantics based on their semantics as expressions. The formalisation can even be applied to contracts for non-OO non-garbage collected languages such as C.

2. It presents simple techniques to verify class invariants. If an invariant is specified with a separation logic predicate, then properties of the invariant, such as the fact that it holds in all visible states [52], often follow as a consequence. The framework can help to devise flexible and sound class invariant protocols.
3. It illustrates the framework’s applicability to model-based specifications [12], where model classes and model queries are used to strengthen contracts [63].
4. It shows how the novel notion of relative purity tolerates side-effects in executable assertions to a high degree.

The techniques are well-suited to separation logic proof tools. We implemented them in the MultiStar tool [65] and verified the examples automatically.

Chapter outline A discussion of background material follows in Section 4.1. Precondition and postcondition verification are the topics of Sections 4.2 and 4.3 respectively. Section 4.4 contains an exposition on class invariants. The verification techniques are then applied to model-based specifications in Section 4.5. Relative purity and predicate extraction are considered in Section 4.6. A discussion of the MultiStar implementation follows in Section 4.7. Finally, Section 4.8 concludes and mentions related work.

4.1 Background

This chapter presents the verification framework in an abstract setting, which can be instantiated with concrete languages and proof systems. The abstract setting is not committed to a particular separation logic, proof system or language type. It is even applicable to non-OO non-garbage-collected languages such as C. The presentation uses the OO language and proof system of Chapter 2 to illustrate the abstract ideas with concrete examples.

4.1.1 The abstract setting: triples and footprints

An abstract triple is of the form $\{P\}s\{Q\}$. The partial correctness meaning of a triple is as follows: statement s does not fault when executed in a state satisfying P , and if it terminates then the resulting state satisfies Q . Faulting occurs when s accesses unallocated memory. In OO terms, ‘unallocated memory’ means heap storage which is not necessarily present in the initial state¹ and not allocated by s before being accessed. O’Hearn [54] uses the term *footprint* of s to describe the minimal state from which s can be executed safely. So s when executed in a state satisfying P will not fault if P describes at least the footprint of s . The notion of footprint has been expanded since [59], yet this is of little concern here.

4.1.2 The concrete setting

The OO language and proof system of Chapter 2 are convenient for illustrating the abstract concepts. To improve readability, we drop empty argument lists in method declarations and calls, and adhere to the uniform access principle [26] where queries can be implemented by fields or result-returning methods. As usual in Hoare-style logics, predicates may use inductive data types and functions involving them. This chapter employs only sequences of integers, where α and β are sequence-valued variables. ϵ denotes the empty sequence, and $[e]$ denotes the singleton sequence whose only element is e . The length of α is written $|\alpha|$, the i ’th element of α is α_i , and $\alpha \mathbin{++} \beta$ denotes the sequence obtained by appending α and β .

The Boolean expressions of the executable contract language can be more complex than the expressions of the kernel language of Chapter 2. In particular, feature calls are allowed in contract expressions. To facilitate formal reasoning about executable assertions, the function $\langle E \rangle_v$ yields a kernel language statement which captures the meaning of the contract expression E in the variable v . Said otherwise, $\langle E \rangle_v$ translates the sugared statement $v := E$ into an appropriate kernel language statement. Formally, assume that v has the same static type as E and that $t, t_1, \dots, t_n$ are fresh variables:

$$\begin{aligned}
 \langle x \rangle_v &\stackrel{\text{def}}{=} v := x, \text{ if } x \text{ is a variable or a constant.} \\
 \langle e_0.f(e_1, \dots, e_n) \rangle_v &\stackrel{\text{def}}{=} \langle e_0 \rangle_{t_0}; \dots; \langle e_n \rangle_{t_n}; v := t_0.f(t_1, \dots, t_n) \\
 \langle \text{not } e \rangle_v &\stackrel{\text{def}}{=} \langle e \rangle_t; v := \text{not } t \\
 \langle e_1 \text{ op } e_2 \rangle_v &\stackrel{\text{def}}{=} \langle e_1 \rangle_{t_1}; \langle e_2 \rangle_{t_2}; v := t_1 \text{ op } t_2, \text{ if } \text{op} \in \{\mathbf{and}, +, -, =, >=\}.
 \end{aligned}$$

¹Separation logic systems do not rely on well-formedness of the heap.

The translation function fixes the semantics of contract language expressions. For example, if we want the contract language to use short-circuit evaluation of e_1 **and** e_2 instead, then its translation would involve a conditional statement. The definition above is sufficient for our purposes, and we use it in the examples that follow.

4.2 Precondition verification

Given a separation logic predicate P and an executable assertion B which should hold at the same program point, the question is whether B somehow follows from P . To this end we define the connecting implication $P \Rightarrow B$, which informally means that P is sufficient to evaluate B to `True`:

$$P \Rightarrow B \stackrel{\text{def}}{=} \{P\}v := B\{P \wedge v = \text{True}\}, \text{ where } v \text{ is a fresh variable.}$$

The intuition behind the definition is that in every state satisfying P , we can evaluate B into v without faulting, and if this computation terminates, then P will hold in the resulting state and v will have value `True`. The routine body relies on P , so it must be re-established by the evaluation of B . This enforces a notion of purity on B , i.e. it prevents B from performing certain kinds of side-effects while allowing others. Section 4.6.1 contains more about this.

In the concrete setting of the OO proof system, there is a sufficient condition² for $P \Rightarrow B$:

$$\boxed{\text{If } \Delta; \Gamma \vdash_s \{P\} \langle B \rangle_v \{P * v = \text{True}\}, \text{ then } P \Rightarrow B \text{ under } \Delta \text{ and } \Gamma.}$$

In intuitionistic separation logic, which is used for garbage-collected OO languages such as our concrete one, $(P * e = e') \Leftrightarrow (P \wedge e = e')$. The $\Delta; \Gamma$ used to verify a B appearing in class C is the same $\Delta; \Gamma$ under which statements and methods of C are verified. In the examples of this chapter we omit explicit reference to Δ and Γ , since none of them uses the additional logical assumptions in Δ , and method specifications in Γ can be read directly from code listings.

Examples. Consider the method preconditions of class *SLIST* in Figure 4.1.

1. The executable precondition of the constructor *SLIST* and features *insert*, *has*, *count* and *is_empty* is `True`. For any separation logic predicate P it holds that $P \Rightarrow \text{True}$:

²Whether or not this condition is also necessary depends on the completeness of the proof system.

```

class SLIST
  define x.LSLIST(l:  $\alpha$ ) as ...
  axiom
    L_sorted:  $\forall \alpha. \mathbf{L}(l: \alpha) \Rightarrow [\forall i, j \in (1..|\alpha|). i < j \Rightarrow \alpha_i \leq \alpha_j]$ 
  feature
    SLIST
    dynamic {True} - {L(l:  $\epsilon$ )}
    executable {True} - {is_empty}

    insert(i: INT)
    dynamic {L(l:  $\alpha$ )} - { $\exists \alpha_F, \alpha_S. \mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha$ }
    executable {True} - {has(i) and count = old(count) + 1}

    remove_first
    dynamic {L(l: [e] ++  $\alpha$ )} - {L(l:  $\alpha$ )}
    executable {not is_empty} - {count = old(count) - 1}

    first: INT
    dynamic {L(l: [e] ++  $\alpha$ )} - {L(l: [e] ++  $\alpha$ ) * Result = e}
    executable {not is_empty} - {not is_empty}

    has(i: INT): BOOL
    dynamic {L(l:  $\alpha$ )} - {L(l:  $\alpha$ ) * Result = ( $\exists j \in 1..|\alpha|. \alpha_j = i$ ) }
    executable {True} - {True}

    count: INT
    dynamic {L(l:  $\alpha$ )} - {L(l:  $\alpha$ ) * Result =  $|\alpha|$ }
    executable {True} - {True}

    is_empty: BOOL
    dynamic {L(l:  $\alpha$ )} - {L(l:  $\alpha$ ) * Result = ( $\alpha = \epsilon$ ) }
    executable {True} - {True}
  invariant
    count_non_negative: {L(l:  $\alpha$ )} count  $\geq 0$ 
    empty_definition: {L(l:  $\alpha$ )} is_empty = (count = 0)
end

```

Figure 4.1: The interface of a sorted list class.

$\Delta; \Gamma \vdash_s \{\text{True} = \text{True}\} v := \text{True} \{v = \text{True}\}$	Assignment axiom
$\Delta; \Gamma \vdash_s \{\text{True}\} v := \text{True} \{v = \text{True}\}$	Consequence
$\Delta; \Gamma \vdash_s \{\text{True} * P\} v := \text{True} \{v = \text{True} * P\}$	Frame rule
$\Delta; \Gamma \vdash_s \{P\} v := \text{True} \{P * v = \text{True}\}$	Consequence

The side-condition of the Frame rule is satisfied: $\text{modifies}(v := \text{True}) = \{v\}$ and $\{v\} \cap \text{FV}(P) = \emptyset$ (remember that v is fresh). The second application of Consequence used the commutativity of $*$ and the fact that $P \Rightarrow (P * \text{True})$ in intuitionistic separation logic.

Instead of such detailed proofs, the rest of the chapter uses proof outlines where statements are interspersed between assertions. Explicitly mentioning the freshness of certain variables is omitted if it is clear from the context. The above proof looks as follows in outline form:

$\{P\}$
 $\quad v := \text{True}$
 $\{P * v = \text{True}\}$

2. The preconditions of *remove_first* and *first* are identical. Here is a proof of $\mathbf{L}(l:[e] \uparrow \alpha) \Rightarrow \mathbf{not} \text{ is_empty}$:

$\{\mathbf{L}(l:[e] \uparrow \alpha)\}$
 $\quad t := \text{is_empty}$
 $\{\mathbf{L}(l:[e] \uparrow \alpha) * t = ([e] \uparrow \alpha = \epsilon)\}$
 $\{\mathbf{L}(l:[e] \uparrow \alpha) * t = \text{False}\}$
 $\quad v := \mathbf{not} \ t$
 $\{\mathbf{L}(l:[e] \uparrow \alpha) * t = \text{False} * v = \neg t\}$
 $\{\mathbf{L}(l:[e] \uparrow \alpha) * v = \text{True}\}$

3. The executable precondition does not always follow from the separation logic one. Suppose that the executable precondition of *has* is not True but instead $\mathbf{not} \text{ is_empty}$. A proof attempt of $\mathbf{L}(l:\alpha) \Rightarrow (\mathbf{not} \text{ is_empty})$ proceeds as follows:

$\{\mathbf{L}(l:\alpha)\}$
 $\quad t := \text{is_empty}$
 $\{\mathbf{L}(l:\alpha) * t = (\alpha = \epsilon)\}$
 $\quad v := \mathbf{not} \ t$
 $\{\mathbf{L}(l:\alpha) * t = (\alpha = \epsilon) * v = \neg t\}$
 $\{\mathbf{L}(l:\alpha) * v = (\alpha \neq \epsilon)\}$

Since $\alpha \neq \epsilon$ is not a consequence, the proof cannot conclude with $\mathbf{L}(l:\alpha) * v = \text{True}$. The separation logic specification is too weak to demonstrate the executable one.³

Such a clash of specifications has several potential causes. The executable assertion is maybe overly restrictive, or it captures important semantic properties of the domain that the separation logic one ignores. If the issue is not resolved, then there is of course no guarantee that the executable assertion will hold at runtime. $\square$

The rest of the chapter contains more examples of the form $P \Rightarrow B$.

4.3 Postcondition verification

Executable postconditions for routines may contain so-called *old-expressions*. The old-expression $\mathbf{old}(E)$ denotes the value of expression E in the state right before the routine is executed, i.e. the state which satisfies the precondition.

Given two separation logic predicates P and Q and an executable assertion B which may contain old-expressions, the connecting implication $P, Q \Rightarrow B$ informally means that any pre-state satisfying P and any post-state satisfying Q are sufficient to evaluate B to True :

Let $\mathbf{old}(E_1), \dots, \mathbf{old}(E_n)$ be the list of old-expressions in B and let $v, v_1, \dots, v_n$ be fresh variables.

$$\begin{aligned} P, Q \Rightarrow B &\stackrel{\text{def}}{=} \exists R. \{P\}(v_1 := E_1, \dots, v_n := E_n) \{P * R\} \text{ and} \\ &\quad \{Q * R\}v := B[v_1; \dots; v_n] \{Q \wedge v = \text{True}\} \\ B[v_1; \dots; v_n] &\stackrel{\text{def}}{=} B[v_1/\mathbf{old}(E_1), \dots, v_n/\mathbf{old}(E_n)] \end{aligned}$$

Intuitively, R contains the result of evaluating $(v_1 := E_1, \dots, v_n := E_n)$ in the pre-state where P holds. Old-expression evaluation must re-establish P because the routine body relies on it. The existence of old-expression results depends on runtime checking, so R is not contained in Q and the body is not allowed to access state described by R . The body is verified only with respect to P and Q , which guarantees that it will establish Q upon termination and never touch R . So $Q * R$ holds right after the body's execution. Evaluating $B[v_1; \dots; v_n]$ in this state should yield True and re-establish Q , since clients

³In fact we have shown that $(\mathbf{not} \text{ is_empty})$ evaluates to False if $\alpha = \epsilon$, so the only way it can also evaluate to True is when it loops forever. A verified implementation of SLIST (and classes transitively used by it) for which $(\mathbf{not} \text{ is_empty})$ terminates when executed in some initial state satisfying $\mathbf{L}(l:\alpha)$ comprises a counterexample for $\mathbf{L}(l:\alpha) \Rightarrow (\mathbf{not} \text{ is_empty})$.




Figure 4.2: (a) Routine r and its specification. (b) Resulting proof obligations. (c) Proof obligations as triples in a proof outline.

depend on the fact that Q holds. Figure 4.2 summarises the proof obligations of a routine.

The predicate R will almost certainly have $v_1, \dots, v_n$ as free variables, but should not have free variables which are modified by the routine body. This allows the transfer of R over the routine body with the Frame rule.

An assertion language's semantics will stipulate how to prove a triple of the form $\{P\}(v_1 := E_1, \dots, v_n := E_n)\{Q\}$. For example, if old-expressions are evaluated in an unspecified order, then $\{P\}s\{Q\}$ must be proved for every permutation s of $(v_1 := E_1, \dots, v_n := E_n)$.

In our concrete OO setting, the assertion language specification states that old-expressions will be evaluated in an arbitrary order. The following condition is sufficient (but not necessary⁴) for concluding $P, Q \Rightarrow B$:

If $\forall i \in (1..n). \Delta; \Gamma \vdash_s \{P\}\langle E_i \rangle_{v_i} \{P * R_i\}$ with **Result** $\notin FV(R_i)$ and $\Delta; \Gamma \vdash_s \{Q * R_1 * \dots * R_n\} v := B[v_1; \dots; v_n] \{Q * v = \text{True}\}$ then $P, Q \Rightarrow B$ under Δ and Γ .

Once again, an assertion B appearing in class C is verified with the same $\Delta; \Gamma$ as the statements and methods of C .

⁴Suppose we take $(\exists e. \text{o.f} \hookrightarrow e * \text{even}(e))$ for P and $\text{o.f}++$ for both E_1 and E_2 . Since there exists no R_i such that $\Delta; \Gamma \vdash_s \{P\}\langle E_i \rangle_{v_i} \{P * R_i\}$ for $i \in 1..2$, the rule cannot prove $P, \text{True} \Rightarrow \text{odd}(\text{old}(E_1) + \text{old}(E_2))$.

Relationships between $P, Q \Rightarrow B$ and $Q \Rightarrow B$ exist, which are convenient in proofs because many executable postconditions contain no old-expressions.

1. If $P, Q \Rightarrow B$ then $Q \Rightarrow B$ whenever B contains no old-expression.
 2. If $Q \Rightarrow B$ then $P, Q \Rightarrow B$ for any P .
- So $P, Q \Rightarrow B$ iff $Q \Rightarrow B$ whenever B contains no old-expression.

Examples. Consider class *SLIST* in Figure 4.1.

1. For *remove_first*, it is the case that $\mathbf{L}(l:[e] ++ \alpha), \mathbf{L}(l:\alpha) \Rightarrow (\text{count} = \mathbf{old}(\text{count}) - 1)$. The executable postcondition contains one old-expression and $E_1 = \text{count}$. Here is the first part of the proof:

$$\begin{aligned} & \{\mathbf{L}(l:[e] ++ \alpha)\} \\ & \quad t1 := \text{count} \\ & \{\mathbf{L}(l:[e] ++ \alpha) * t1 = |[e] ++ \alpha|\} \end{aligned}$$

Notice that $\text{FV}(t1 = |[e] ++ \alpha|) = \{t1, e, \alpha\}$. Using $(t1 = |[e] ++ \alpha|)$ for R_1 completes the proof:

$$\begin{aligned} & \{\mathbf{L}(l:\alpha) * t1 = |[e] ++ \alpha|\} \\ & \quad t2 := \text{count} \\ & \{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t1 = |[e] ++ \alpha|\} \\ & \{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t1 = |\alpha| + 1\} \\ & \quad t3 := t1 - 1 \\ & \{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t1 = |\alpha| + 1 * t3 = |\alpha|\} \\ & \{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t3 = |\alpha|\} \\ & \quad v := t2 = t3 \\ & \{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t3 = |\alpha| * v = (t2 = t3)\} \\ & \{\mathbf{L}(l:\alpha) * v = \text{True}\} \end{aligned}$$

2. The executable postcondition of the constructor *SLIST* contains no old-expression, so it suffices to prove $\mathbf{L}(l:\epsilon) \Rightarrow \text{is_empty}$:

$$\begin{aligned} & \{\mathbf{L}(l:\epsilon)\} \\ & \quad v := \text{is_empty} \\ & \{\mathbf{L}(l:\epsilon) * v = (\epsilon = \epsilon)\} \\ & \{\mathbf{L}(l:\epsilon) * v = \text{True}\} \end{aligned}$$

3. The executable postcondition of *first* contains no old-expression, so we only need $(\mathbf{L}(l:[e] ++ \alpha) * \mathbf{Result} = e) \Rightarrow \mathbf{not} \text{ is_empty}$. Example 2 in the previous section established the inner triple of the proof:

$$\begin{aligned}
& \{\mathbf{L}(l: [e] ++ \alpha) * \mathbf{Result} = e\} \\
& \quad \{\mathbf{L}(l: [e] ++ \alpha)\} \\
& \quad (\mathbf{not\ is_empty})_v \\
& \quad \{\mathbf{L}(l: [e] ++ \alpha) * v = \mathbf{True}\} \\
& \{\mathbf{L}(l: [e] ++ \alpha) * v = \mathbf{True} * \mathbf{Result} = e\} \\
& \{\mathbf{L}(l: [e] ++ \alpha) * \mathbf{Result} = e * v = \mathbf{True}\}
\end{aligned}$$

The Frame rule is key to the proof.

4. Consider the postcondition of *insert*. It contains one old-expression, namely **old**(count):

$$\begin{aligned}
& \{\mathbf{L}(l: \alpha)\} \\
& \quad t1 := \text{count} \\
& \{\mathbf{L}(l: \alpha) * t1 = |\alpha|\}
\end{aligned}$$

The second part of the proof consists of small pieces which are put together.

$$\begin{aligned}
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S)\} \\
& \quad t2 := \text{has}(i) \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t2 = (\exists j \in 1..|\alpha_F ++ [i] ++ \alpha_S|. (\alpha_F ++ [i] ++ \alpha_S)_j = i)\} \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t2 = \mathbf{True}\}
\end{aligned}$$

Applying Frame and Consequence to this triple yields Piece 1:

$$\begin{aligned}
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha|\} \\
& \quad t2 := \text{has}(i) \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha| * t2 = \mathbf{True}\}
\end{aligned}$$

Next, we prove

$$\begin{aligned}
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha|\} \\
& \quad t3 := \text{count} \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t3 = |\alpha_F ++ [i] ++ \alpha_S| * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha|\} \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t3 = |\alpha| + 1 * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha|\} \\
& \quad t4 := t1 + 1 \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t3 = |\alpha| + 1 * \alpha_F ++ \alpha_S = \alpha * t1 = |\alpha| * t4 = t1 + 1\} \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t3 = |\alpha| + 1 * \alpha_F ++ \alpha_S = \alpha * t4 = |\alpha| + 1\} \\
& \quad t5 := t3 = t4 \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * t3 = |\alpha| + 1 * \alpha_F ++ \alpha_S = \alpha * t4 = |\alpha| + 1 * t5 = (t3=t4)\} \\
& \{\mathbf{L}(l: \alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha * t5 = \mathbf{True}\}
\end{aligned}$$

Applying Frame establishes Piece 2:

$$\begin{aligned}
& \{L(l:\alpha_F \mathbin{++} [i] \mathbin{++} \alpha_S) * \alpha_F \mathbin{++} \alpha_S = \alpha * t1 = |\alpha| * t2 = \text{True}\} \\
& \quad t3 := \text{count} \\
& \quad t4 := t1 + 1 \\
& \quad t5 := t3 = t4 \\
& \{L(l:\alpha_F \mathbin{++} [i] \mathbin{++} \alpha_S) * \alpha_F \mathbin{++} \alpha_S = \alpha * t5 = \text{True} * t2 = \text{True}\}
\end{aligned}$$

Here is Piece 3:

$$\begin{aligned}
& \{L(l:\alpha_F \mathbin{++} [i] \mathbin{++} \alpha_S) * \alpha_F \mathbin{++} \alpha_S = \alpha * t5 = \text{True} * t2 = \text{True}\} \\
& \quad v := t2 \text{ and } t5 \\
& \{L(l:\alpha_F \mathbin{++} [i] \mathbin{++} \alpha_S) * \alpha_F \mathbin{++} \alpha_S = \alpha * t5 = \text{True} * t2 = \text{True} * v = (t2 \text{ and } t5)\} \\
& \{L(l:\alpha_F \mathbin{++} [i] \mathbin{++} \alpha_S) * \alpha_F \mathbin{++} \alpha_S = \alpha * v = \text{True}\}
\end{aligned}$$

The proof is completed by putting the three pieces together, eliminating α_F and α_S with the Auxiliary variable elimination rule, and applying Consequence to move $(t1 = |\alpha|)$ and $(v = \text{True})$ out from under the quantifiers to the top level in the precondition and postcondition respectively. $\square$

4.4 Class invariant verification

Consider an executable assertion B which should hold at program point pp . If the separation logic predicate P holds at pp and $P \Rightarrow B$, then B is verified for pp . This technique can be used to verify executable assert statements and loop invariants, for example.

Class invariants⁵ (henceforth simply called invariants) are verified in a similar way. Since an invariant protocol [64] always specifies the points in a program where an invariant, say B , should hold, B is verified if it is verified for all such points.

Another technique is to annotate an invariant B with a separation logic predicate P which characterises the states in which the invariant should hold. If $P \Rightarrow B$, then B is verified. This offers a flexible scheme if annotations are given on the level of individual invariant clauses.

Examples. Consider the following invariant clauses, reproduced from the bottom of class SLIST in Figure 4.1:⁶

⁵Also called *object invariants*.

⁶These are public invariants, as opposed to private representation invariants (which is how the term ‘invariant’ is used in the Spec# literature, e.g. [45]). Private invariants can be verified in a similar way.

invariant

count_non_negative: $\{\mathbf{L}(l:\alpha)\} \text{ count} \geq 0$
 empty_definition: $\{\mathbf{L}(l:\alpha)\} \text{ is_empty} = (\text{count} = 0)$

1. For the clause named *count_non_negative*, $\mathbf{L}(l:\alpha) \Rightarrow (\text{count} \geq 0)$ holds:

$$\begin{aligned} &\{\mathbf{L}(l:\alpha)\} \\ &\quad t := \text{count} \\ &\{\mathbf{L}(l:\alpha) * t = |\alpha|\} \\ &\quad v := t \geq 0 \\ &\{\mathbf{L}(l:\alpha) * t = |\alpha| * v = (t \geq 0)\} \\ &\{\mathbf{L}(l:\alpha) * v = \text{True}\} \end{aligned}$$

2. The invariant clause *empty_definition* is verified similarly:

$$\begin{aligned} &\{\mathbf{L}(l:\alpha)\} \\ &\quad t1 := \text{is_empty} \\ &\{\mathbf{L}(l:\alpha) * t1 = (\alpha = \epsilon)\} \\ &\quad t2 := \text{count} \\ &\{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t1 = (\alpha = \epsilon)\} \\ &\quad t3 := t2 = 0 \\ &\{\mathbf{L}(l:\alpha) * t2 = |\alpha| * t1 = (\alpha = \epsilon) * t3 = (t2 = 0)\} \\ &\{\mathbf{L}(l:\alpha) * t1 = (\alpha = \epsilon) * t3 = (|\alpha| = 0)\} \\ &\quad v := t1 = t3 \\ &\{\mathbf{L}(l:\alpha) * t1 = (\alpha = \epsilon) * t3 = (|\alpha| = 0) * v = (t1 = t3)\} \\ &\{\mathbf{L}(l:\alpha) * v = ((\alpha = \epsilon) = (|\alpha| = 0))\} \\ &\{\mathbf{L}(l:\alpha) * v = \text{True}\} \end{aligned}$$

□

The fact that the structural rules of separation logic are sound provides a simple way to show that a verified invariant clause holds at a particular program point. The following rules can also be used when verifying other executable assertions:^{7 8}

⁷Similar rules can be given for $P, Q \Rightarrow B$.

⁸These rules are independent of the structure of B . Rules also exist which do depend on its structure. For example, if $P_1 \Rightarrow B_1$ and $P_2 \Rightarrow B_2$, then $(P_1 * P_2) \Rightarrow (B_1 \text{ **par_and** } B_2)$. This rule allows concurrency in executable assertions.

$$\boxed{
\begin{array}{c}
\frac{P \Rightarrow B}{(P * Q) \Rightarrow B} \text{ Frame}' \\
\\
\frac{P \Leftrightarrow Q \quad P \Rightarrow B}{Q \Rightarrow B} \text{ Conseq}' \\
\\
\frac{P \Rightarrow B}{(\exists x. P) \Rightarrow B} \text{ AuxVarElim}' \\
\\
\frac{P \Rightarrow B \quad Q \Rightarrow B}{(P \vee Q) \Rightarrow B} \text{ Disj}' \\
\\
\text{and others.}
\end{array}
}$$

The Frame' rule does not need any side-condition, since the translation of B will modify only fresh variables, i.e. variables not free in Q . Note the way Conseq' is written: if $P \Rightarrow B$ and $Q \Rightarrow P$, then $Q \Rightarrow B$ does *not* necessarily hold (example 3 in Section 4.6.1 provides a counterexample). In the $\text{AuxVarElim}'$ rule, the variable x may not be free in B .

Example. If a program containing class *SLIST* from Figure 4.1 is verified, then it follows as a consequence of the separation logic specification that both invariant clauses hold in the visible states [52] of *SLIST*. In other words, both clauses are established by the constructor and hold on entry and exit of all other public features. Since each clause B is verified, we know $L(l:\alpha) \Rightarrow B$. In the following proofs the ' $\Rightarrow B$ ' part is omitted.

1. Upon exit from *SLIST*:

$$\begin{array}{c}
\frac{L(l:\alpha)}{\quad} \text{Frame}' \\
\frac{L(l:\alpha) * \alpha = \epsilon}{\quad} \text{Conseq}' \\
\frac{L(l:\epsilon) * \alpha = \epsilon}{\quad} \text{AuxVarElim}' \\
\frac{\exists \alpha. L(l:\epsilon) * \alpha = \epsilon}{\quad} \text{Conseq}' \\
\frac{L(l:\epsilon) * \text{True}}{\quad} \text{Conseq}' \\
\hline
L(l:\epsilon)
\end{array}$$

2. Upon entry to *remove_first*:

$\frac{}{\mathbf{L}(l:\alpha)}$	Frame'
$\frac{}{\mathbf{L}(l:\alpha) * \alpha = \beta}$	Conseq'
$\frac{}{\mathbf{L}(l:\beta) * \alpha = \beta}$	AuxVarElim'
$\frac{}{\exists \alpha. \mathbf{L}(l:\beta) * \alpha = \beta}$	Conseq'
$\frac{}{\mathbf{L}(l:\beta)}$	Frame'
$\frac{}{\mathbf{L}(l:\beta) * \beta = [e] ++ \alpha}$	Conseq'
$\frac{}{\mathbf{L}(l:[e] ++ \alpha) * \beta = [e] ++ \alpha}$	AuxVarElim'
$\frac{}{\exists \beta. \mathbf{L}(l:[e] ++ \alpha) * \beta = [e] ++ \alpha}$	Conseq'
$\frac{}{\mathbf{L}(l:[e] ++ \alpha)}$	

3. For the exit of *insert*, we start with $\mathbf{L}(l:\beta)$ which was derived in the previous proof.

$\frac{}{\mathbf{L}(l:\beta)}$	Frame'
$\frac{}{\mathbf{L}(l:\beta) * \beta = (\alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha}$	AuxVarElim'
$\frac{}{\exists \beta. \mathbf{L}(l:\beta) * \beta = (\alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha}$	Conseq'
$\frac{}{\mathbf{L}(l:\alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha}$	AuxVarElim'
$\frac{}{\exists \alpha_S. \mathbf{L}(l:\alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha}$	AuxVarElim'
$\frac{}{\exists \alpha_F, \alpha_S. \mathbf{L}(l:\alpha_F ++ [i] ++ \alpha_S) * \alpha_F ++ \alpha_S = \alpha}$	

□

In effect we create fine-grained invariant protocols by annotating invariant clauses: the predicate abstractly specifies the program points where they should hold, and invariants can be verified, i.e. guaranteed, to hold there. One can refine this idea to create more sophisticated invariant protocols.

Examples.

1. Here is an invariant protocol for single-inheritance programs where subclasses can refine individual invariant clauses. Γ maps a class name and invariant clause name to the invariant specification predicate and Boolean expression. Two rules distinguish the case where class C introduces an invariant clause named *inv* from the case where C inherits and refines *inv* from its parent C' :

$$\frac{\forall i \in (1..n). (P * \mathbf{Current} : C) \Rightarrow B_i}{\Gamma \vdash_i \mathbf{introduce} \text{ inv: } \{P\} \langle B_1 \rangle \text{ and } \dots \text{ and } \langle B_n \rangle \text{ in } C}$$

$$\begin{array}{l}
C' \prec_1 C \\
\Gamma(C.\text{inv}) = \{P\} B \\
B = \langle B_1 \rangle \text{ and } \dots \text{ and } \langle B_n \rangle \\
P \Leftrightarrow (P' * R) \\
\forall i \in (1..n). (P' * \mathbf{Current} : C') \Rightarrow B_i \\
\forall j \in (1..m). (P' * \mathbf{Current} : C') \Rightarrow B'_j \\
\hline
\Gamma \vdash_i \mathbf{inherit} \text{ inv: } \{P'\} B \text{ and } \langle B'_1 \rangle \text{ and } \dots \text{ and } \langle B'_m \rangle \text{ in } C'
\end{array}$$

The above invariant protocol guarantees that whenever ‘inv: $\{P\} B$ ’ appears in class C , then $\forall x <: C. P[x/\mathbf{Current}] \Rightarrow B[x/\mathbf{Current}]$ where x is fresh. Enclosing an expression in angle-brackets provides it with a side-effect scope. The next example explains why it works.

2. The previous invariant protocol is an instance of a more general scheme. Here, invariant clauses are equipped with both a dynamic and a static specification predicate. When a class C introduces an invariant clause:

introduce inv: $\{P\} \{S\} \langle B_1 \rangle \text{ and } \dots \text{ and } \langle B_n \rangle$

then it is considered a shorthand for introducing the methods $C.\text{inv}_i$, where $i \in 1..n$:

```

introduce  $C.\text{inv}_i$ : BOOL
dynamic  $\{P\} \{P * \mathbf{Result} = \text{True}\}$ 
static  $\{S\} \{S * \mathbf{Result} = \text{True}\}$ 
do  $\langle B_i \rangle_{\mathbf{Result}}$  end

```

When the direct subclass C' refines the invariant clause by listing:

inherit inv: $\{P'\} \{S'\} B \text{ and } \langle B'_1 \rangle \text{ and } \dots \text{ and } \langle B'_m \rangle$

where $B = \langle B_1 \rangle \text{ and } \dots \text{ and } \langle B_n \rangle$, then it is considered a shorthand for overriding all the methods $C.\text{inv}_i$, where $i \in 1..n$:

```

override  $C.\text{inv}_i$ : BOOL
dynamic  $\{P'\} \{P' * \mathbf{Result} = \text{True}\}$ 
static  $\{S'\} \{S' * \mathbf{Result} = \text{True}\}$ 
do  $\langle B'_i \rangle_{\mathbf{Result}}$  end

```

and introducing the methods $C'.\text{inv}_j$, where $j \in 1..m$:

```

introduce  $C'.\text{inv}_j$ : BOOL
dynamic  $\{P'\} \{P' * \mathbf{Result} = \text{True}\}$ 
static  $\{S'\} \{S' * \mathbf{Result} = \text{True}\}$ 
do  $\langle B'_j \rangle_{\mathbf{Result}}$  end

```

Looking at the method verification rules of 2.6, we see that class C must satisfy:

$$\Delta \vdash \{S\}_{-}\{S * \mathbf{Result} = \text{True}\} \xRightarrow{\mathbf{Current}:C} \{P\}_{-}\{P * \mathbf{Result} = \text{True}\}$$

$$\forall i \in (1..n) \cdot S \Rightarrow B_i$$

and class C' must satisfy:

$$\Delta \vdash \{P'\}_{-}\{P' * \mathbf{Result} = \text{True}\} \Rightarrow \{P\}_{-}\{P * \mathbf{Result} = \text{True}\}$$

$$\Delta \vdash \{S'\}_{-}\{S' * \mathbf{Result} = \text{True}\} \xRightarrow{\mathbf{Current}:C'} \{P'\}_{-}\{P' * \mathbf{Result} = \text{True}\}$$

$$\forall i \in (1..n) \cdot S' \Rightarrow B_i$$

$$\forall j \in (1..m) \cdot S' \Rightarrow B'_j$$

By taking $S = P * \mathbf{Current} : C$, and $S' = P' * \mathbf{Current} : C'$, it is simple to prove that the two rules of the previous example are theorems. It is also easy to verify the protocol's guarantee.

3. Instead of overriding all the methods C_inv_i in class C' , as in the previous example, a protocol might choose to inherit them:

```
inherit C_inv_i: BOOL
dynamic {P'}_{-}{P' * Result = True}
static {S'}_{-}{S' * Result = True}
```

Doing so changes the proof obligations of class C' – instead of:

$$\forall i \in (1..n) \cdot S' \Rightarrow B_i$$

we must establish that:

$$\Delta \vdash \{S\}_{-}\{S * \mathbf{Result} = \text{True}\} \Rightarrow \{S'\}_{-}\{S' * \mathbf{Result} = \text{True}\} \quad \square$$

4.5 Model-based specifications

Since executable specifications are frequently not very expressive, model classes and model-based contracts are sometimes used to strengthen them [12, 63].

Example. Consider the interface of model class `SEQUENCE` in Figure 4.3. It provides an abstraction of immutable sequences for specification purposes. Class `SLIST` can be specified in terms of `SEQUENCE`, as the interface extract in Figure 4.4 shows. Note that `SLIST` now has a model query, namely *model*, which returns the immutable sequence abstraction of an `SLIST` instance at the point when it is called. A comparison of *remove_first*'s specification in Figures 4.1 and 4.4 shows that the model-based specification involves the

```

class SEQUENCE
define x.SEQSEQUENCE(s:  $\alpha$ ) as ...
feature
  SEQUENCE
  dynamic {True} - {SEQ(s:  $\epsilon$ )}

  cons(i: INT): SEQUENCE
  dynamic {SEQ(s:  $\alpha$ )} - {SEQ(s:  $\alpha$ ) * Result.SEQ(s: [i] ++  $\alpha$ )}

  head: INT
  dynamic {SEQ(s: [e] ++  $\alpha$ )} - {SEQ(s: [e] ++  $\alpha$ ) * Result = e}

  tail: SEQUENCE
  dynamic {SEQ(s: [e] ++  $\alpha$ )} - {SEQ(s: [e] ++  $\alpha$ ) * Result.SEQ(s:  $\alpha$ )}

  is_nil: BOOL
  dynamic {SEQ(s:  $\alpha$ )} - {SEQ(s:  $\alpha$ ) * Result = ( $\alpha = \epsilon$ )}

  eq(o: SEQUENCE): BOOL
  dynamic {SEQ(s:  $\alpha$ ) * o.SEQ(s:  $\beta$ )} -
    {SEQ(s:  $\alpha$ ) * o.SEQ(s:  $\beta$ ) * Result = ( $\alpha = \beta$ )}
end

```

Figure 4.3: The interface of model class SEQUENCE.

```

class SLIST
  ...
feature
  ...
  model: SEQUENCE
  dynamic { $\mathbf{L}(l:\alpha)$ }-{ $\mathbf{L}(l:\alpha) * \mathbf{Result}.\mathbf{SEQ}(s:\alpha)$ }
  executable {True}-{True}

  remove_first
  dynamic { $\mathbf{L}(l:[e] ++ \alpha)$ }-{ $\mathbf{L}(l:\alpha)$ }
  executable {not model.is_nil}-{model.eq(old(model).tail)}
  ...
invariant
  ...
  empty_inv: { $\mathbf{L}(l:\alpha)$ } is_empty = model.is_nil
end

```

Figure 4.4: An extract from class SLIST which uses model-based contracts.

element values stored in an SLIST instance and not just their number. $\square$

The following three problems are typically not easy to solve with conventional techniques [39, 19]:

1. Devising semantics for model classes and proving their implementations correct.
2. Giving a semantics to model queries, such as *model* in Figure 4.4, and proving their implementations correct.
3. Verifying model-based specifications, such as the model-based contract of *remove_first* and the model-based invariant clause *empty_inv* in Figure 4.4.

The first two problems can be solved with separation logic. Figures 4.3 and 4.4 show separation logic specifications for a model class and model query. A conventional separation logic proof system, such as the one of Chapter 2, can be used to verify their implementations. The third problem can then be addressed with the framework of this chapter.

Examples. Consider the model-based specifications of class SLIST in Figure 4.4. Suppose that omitted features have the same separation logic specifications as in Figure 4.1.

1. For the invariant clause *empty_inv*:

$$\begin{aligned}
& \{\mathbf{L}(l:\alpha)\} \\
& \quad t1 := \text{is_empty} \\
& \{\mathbf{L}(l:\alpha) * t1 = (\alpha = \epsilon)\} \\
& \quad t2 := \text{model} \\
& \{\mathbf{L}(l:\alpha) * t2.\text{SEQ}(s:\alpha) * t1 = (\alpha = \epsilon)\} \\
& \quad t3 := t2.\text{is_nil} \\
& \{\mathbf{L}(l:\alpha) * t2.\text{SEQ}(s:\alpha) * t3 = (\alpha = \epsilon) * t1 = (\alpha = \epsilon)\} \\
& \{\mathbf{L}(l:\alpha) * t3 = (\alpha = \epsilon) * t1 = (\alpha = \epsilon)\} \\
& \quad v := t1 = t3 \\
& \{\mathbf{L}(l:\alpha) * t3 = (\alpha = \epsilon) * t1 = (\alpha = \epsilon) * v = (t1 = t3)\} \\
& \{\mathbf{L}(l:\alpha) * v = \text{True}\}
\end{aligned}$$

2. We next verify the postcondition of *remove_first* which contains one old-expression:

$$\begin{aligned}
& \{\mathbf{L}(l:[e] ++ \alpha)\} \\
& \quad v1 := \text{model} \\
& \{\mathbf{L}(l:[e] ++ \alpha) * v1.\text{SEQ}(s:[e] ++ \alpha)\}
\end{aligned}$$

The proof is completed by using $v1.\text{SEQ}(s:[e] ++ \alpha)$ for R_1 :

$$\begin{aligned}
& \{\mathbf{L}(l:\alpha) * v1.\text{SEQ}(s:[e] ++ \alpha)\} \\
& \quad t1 := \text{model} \\
& \{\mathbf{L}(l:\alpha) * t1.\text{SEQ}(s:\alpha) * v1.\text{SEQ}(s:[e] ++ \alpha)\} \\
& \quad t2 := v1.\text{tail} \\
& \{\mathbf{L}(l:\alpha) * t1.\text{SEQ}(s:\alpha) * v1.\text{SEQ}(s:[e] ++ \alpha) * t2.\text{SEQ}(s:\alpha)\} \\
& \{\mathbf{L}(l:\alpha) * t1.\text{SEQ}(s:\alpha) * t2.\text{SEQ}(s:\alpha)\} \\
& \quad v := t1.\text{eq}(t2) \\
& \{\mathbf{L}(l:\alpha) * t1.\text{SEQ}(s:\alpha) * t2.\text{SEQ}(s:\alpha) * v = (\alpha = \alpha)\} \\
& \{\mathbf{L}(l:\alpha) * v = \text{True}\}
\end{aligned}$$

□

4.6 Relative purity and predicate extraction

4.6.1 Relative purity

Side-effects in executable specifications conventionally complicate verification, since the logical predicates extracted from them cannot be imperative. Most techniques therefore impose some form of purity, i.e. side-effect freeness, on specification expressions. Purity comes in many flavours, such as strong purity, weak purity and observational purity [18, 53].

The verification techniques of this chapter tolerate side-effects in executable specifications to a high degree. In fact, they use a new notion of purity, namely *relative* purity. An executable assertion is pure or impure with respect to a given separation logic specification:

B is pure relative to P $\stackrel{\text{def}}{=} \{P\}v := B\{P\}$, where v is a fresh variable.

This means informally that B can be evaluated in a state satisfying P, and that the evaluation will preserve P. Similarly, for an executable assertion B containing the list of **old**-expressions **old**(E₁), ..., **old**(E_n):

B is pure relative to P, Q $\stackrel{\text{def}}{=} \exists R. \{P\}(v_1 := E_1, \dots, v_n := E_n)\{P * R\}$ and $\{Q * R\}v := B[v_1; \dots; v_n]\{Q\}$

where v, v₁, ..., v_n are fresh variables and B[v₁; ...; v_n] is defined as before.

The following lemmas follow from the definitions by the rule of Consequence:

If $P \Rightarrow B$, then B is pure relative to P.

If $P, Q \Rightarrow B$, then B is pure relative to P, Q.

Examples. Consider the specification of class SLIST in Figure 4.1.

1. Suppose we replace the executable postcondition of *insert* with a hypothetical assertion B for which $\llbracket B \rrbracket_v$ is given by:

```
t := remove_first
insert(t)
v := has(i)
```

B is pure relative to the separation logic postcondition of *insert* because one can prove⁹ $(\exists \alpha_F, \alpha_S. \mathbf{L}(l: \alpha_F \uparrow\uparrow [i] \uparrow\uparrow \alpha_S) * \alpha_F \uparrow\uparrow \alpha_S = \alpha) \Rightarrow B$.

⁹The proof uses the axiom information expressed in *L.sorted*. The rule of Disjunction can combine subproofs arising from a case split on $\alpha_F = \epsilon$.

2. Suppose B is a hypothetical assertion for which $\langle B \rangle_v$ is:

```
insert(4)
v := not is_empty
```

B is not pure relative to $L(l:\alpha)$. We cannot complete the following proof attempt, since the needed implication does not necessarily hold:

```
{L(l:α)}
  insert(4)
{∃αF, αS. L(l:αF ++ [4] ++ αS) * αF ++ αS = α}
  v := not is_empty
{∃αF, αS. L(l:αF ++ [4] ++ αS) * αF ++ αS = α}
// An implication is needed to establish:
{L(l:α)}
```

A counterexample showing that B is not pure relative to P in the context of class C comprises:

- (a) a verified implementation of C and all classes transitively used by it;
 - (b) an execution trace where B is evaluated in an initial state satisfying P that either faults or terminates in a state not satisfying P.
3. We use B from the previous example and show $(\exists\alpha. L(l:\alpha)) \Rightarrow B$:

```
{∃α. L(l:α)}
  insert(4)
{∃α, αF, αS. L(l:αF ++ [4] ++ αS) * αF ++ αS = α}
  v := not is_empty
{∃α, αF, αS. L(l:αF ++ [4] ++ αS) * v = True * αF ++ αS = α}
{∃α. L(l:α) * v = True}
```

So B is pure with respect to $\exists\alpha. L(l:\alpha)$.

Since $L(l:\alpha) \Rightarrow (\exists\alpha. L(l:\alpha))$, this example and the previous one show that if $P \Rightarrow Q$ and $Q \Rightarrow B$, then $P \Rightarrow B$ does not hold in general.

4. Let B be any executable assertion without old-expressions. B is always pure relative to False. If B never faults, then it is also pure relative to True. $\square$

Any notion of purity which classifies side-effects as intrinsically harmful or not will rule out executable assertions that programmers might want to write. Relative purity judges whether a side-effect is harmful or not only with

respect to the properties one cares about. For example, the side-effects in B of examples 2 and 3 are harmless with respect to $\exists\alpha. \mathbf{L}(l:\alpha)$, but harmful with respect to $\mathbf{L}(l:\alpha)$. This is so because the side-effect of inserting 4 in a sorted list will maintain a list structure (what the first purity specification states), but not a list structure with the same elements (what the second specification demands). The concept of relativity pervades verification already, since code is only correct or incorrect relative to a specification. We adopt it also for purity.

Relative purity guarantees soundness of reasoning and does not impose unnecessary constraints on executable assertions. An executable assertion is free to perform any side-effect as long as nothing happened from the perspective of the separation logic specification. Runtime assertion checking does nothing relative to the specification of a verified program. An executable assertion can even print ‘Hello World!’ if the specification permits it. Purity is in the eye of the assenter.

4.6.2 Predicate extraction

Programmers are conventionally encouraged to specify sets of states with computations. Hoare-style verification cannot do much with computations – it needs predicates. So conventional verification approaches, including the Spec# system [3] and JML toolset [9], have to extract predicates from computations. This can be achieved in various ways: as long as there is an agreed-upon mapping from computations to predicates, the computations can be seen as syntactic sugar for predicates.

When given an executable assertion B without old-expressions¹⁰, one possibility is to extract a predicate P which follows from B . This is described by the connecting implication $B \Rightarrow P$, which informally means that B evaluating to True is sufficient to conclude P . Here is a formal definition:

$$B \Rightarrow P \stackrel{\text{def}}{=} \forall Q. \text{ if } Q \Rightarrow B \text{ then } Q \Rightarrow P.$$

If $B \Rightarrow P$, then $P \Rightarrow B$ does not necessarily hold. However, the best predicate that can be extracted in this way, i.e. the strongest P such that $B \Rightarrow P$, is precisely the weakest P such that $P \Rightarrow B$. So if we define

$$\text{Best}(B) \stackrel{\text{def}}{=} \bigvee \{P \mid P \Rightarrow B\}$$

then the following lemmas hold:

¹⁰The treatment can be generalised to executable assertions with old-expressions.

1. $Best(B) \Rightarrow B$ and $B \Rightarrow Best(B)$, i.e. $Best(B) \Leftrightarrow B$
2. $B \Rightarrow P$ iff $Best(B) \Rightarrow P$
3. If $P \Rightarrow B$ then $P \Rightarrow Best(B)$

Since $Best(B) \Leftrightarrow B$, the predicate $Best(B)$ can be seen as the exact logical counterpart of B .

Examples. Let us ignore inheritance and suppose an implementation of class SLIST in Figure 4.1 maintains its stored elements in a list of linked nodes.

1. If *is_empty* is implemented as a field in class SLIST, then $Best(is_empty) = (is_empty \hookrightarrow \text{True})$.
2. If *is_empty* tests whether a field named ‘head’, which points to the first node, is **Void**, then $Best(\text{not } is_empty) = (\exists x. \text{head} \hookrightarrow x * x \neq \text{Void})$.
3. If *has* is implemented by a simple linear traversal of the nodes, then the predicate $Best(has(i))$ denotes that *i* is encountered inside a node before the next node pointer becomes **Void**. Note that $Best(has(i))$ is too weak to rule out a cyclic or frying-pan [8] list. Such a list does not even have to contain *i* in order to satisfy $Best(has(i))$, because the definition does not demand the termination of *has*(*i*). $\square$

As these examples show, the logical counterparts of executable assertions are often very weak. Predicate extraction techniques will have to map computations into stronger predicates in order to verify method bodies. Such predicates are similar in flavour to the ones used in this chapter, and do not follow from the executable ones. Section 4.8 discusses this in more depth.

Our approach uses predicates rather than computations to characterise states. It views the separation logic specification as the primary one for verification. There is no extraction of predicates, and verification treats executable assertions as computations which should evaluate to **True**. We do not verify the program *using* its executable assertions, we verify the program *and* its executable assertions.

4.7 Implementation

We extended MultiStar to provide fully automatic verification of executable Eiffel assertions with respect to separation logic specifications. The implementation leverages mechanisms which are present in MultiStar and most separation logic proof tools: symbolic execution, implication checking and

frame inference. Symbolic execution and implication checking are the basic ingredients for proving connecting implications. For example, when proving $P \Rightarrow B$, MultiStar executes $v := B$ in the symbolic state P to obtain a symbolic state P' . It then checks whether $P' \Rightarrow (P * v = \text{True})$. For executable postconditions, the results of old-expressions are inferred automatically with frame inference and not specified manually. Suppose MultiStar must prove the connecting implication $P, Q \Rightarrow C$, where C contains a single old-expression **old**(E). Executing $v_1 := E$ in the symbolic state P yields a resulting symbolic state P' . Frame inference next determines the symbolic result R of E in $P' \vdash P * R$. Then $w := C[v_1]$ is executed in the symbolic state $Q * R$ and yields a symbolic state Q' . Finally, MultiStar checks whether $Q' \Rightarrow (Q * w = \text{True})$.

MultiStar can easily handle the examples presented in this chapter, which are also available in the EVE download [28].

4.8 Conclusions and related work

The presented framework offers a sound and simple way to verify various executable specifications with separation logic. The notion of relative purity is central to the framework and embraces side-effects in executable assertions, thereby allowing programmers more freedom of expression compared to conventional verification approaches. The framework is well-suited to separation logic proof tools and is implemented in MultiStar.

Regarding related work, separation logic [54, 60] offers local reasoning for heap-manipulating programs and is central to the framework. Its adaptation to object-orientation [57, 56, 58, 65] is especially relevant for the presentation.

Executable specifications are embodied in programming languages such as Eiffel [26] and Spec# [3]. Dedicated specification languages such as JML [40] also use them. JML includes model classes and allows model-based contracts [12]. Another library with model classes is MML [63]. Executable assertions offer several benefits in software development, including runtime checking [26, 13] and automated testing [47, 15].

The problem of obtaining the strongest P such that $B \Rightarrow P$ amounts to finding the weakest footprint of B preserved by B which ensures B evaluates to True. This problem appears to be similar in flavour but more general than abduction [10, 30].

Conventional approaches to verification, including the Spec# system [3] and JML toolset [9], extract predicates from assertions in a different way [45, 18]:

1. They do not use footprints. Well-formedness of the heap is exploited and any chain of references can be followed. Since a method can po-

tentially modify any reachable object, the absence of footprints makes reasoning more global. Specifications must describe which objects are modified, because there might be external references to reachable objects. The *aliasing problem* arises because objects maintaining such references, or aliases, can depend on properties of the aliased objects. Approaches to the problem typically restrict or prevent aliases and/or operations on references. They include confinement, sharing and access control, ownership, immutability, uniqueness, information flow and escape analyses [16].

2. They impose purity constraints on the methods used in assertion expressions. Several notions of purity exist which prevent methods from changing the state [18, 53]. A strongly pure method has no side-effect at all. A weakly pure method does not change existing objects, but might allocate and modify new ones. An observationally pure method may modify existing objects, provided that the change is sufficiently encapsulated such that no other class can observe a change. Proving method purity becomes harder as the notion becomes more permissive. Weak purity can be proved with a combination of pointer and escape analysis [62], while a method's observational purity can be shown by proving that it simulates a weakly pure one from the perspective of other classes [53].
3. They encode pure methods and their contracts in first-order logic as uninterpreted functions and axioms respectively [18, 45, 17]. Checking well-formedness of pure method specifications is vital in this step, because inconsistent axiomatisations can result from unsatisfiable specifications or recursive specifications which are ill-founded [61]. A pure method call in an assertion is encoded as an application of the corresponding uninterpreted function. A loop, written in the stylised form of quantification over a finite domain (as in JML and Spec#), is directly encoded as a quantified formula. The reader is referred to [45] for assertion encoding details. Model class and model-based contract encoding is discussed in [19].

CHAPTER 5

CORRECTNESS BY CONSTRUCTION

The previous two chapters proposed extensions to the separation logic proof system: features for reasoning about related abstractions, multiple inheritance and executable contracts. Instead of adding new features, this chapter investigates correctness by construction – an alternative way to develop verified programs. Dijkstra, who advocated this approach, described it well in his Turing Award lecture [22]:

The only effective way to raise the confidence level of a program significantly is to give a convincing proof of its correctness. But one should not first make the program and then prove its correctness, because then the requirement of providing the proof would only increase the poor programmer's burden. On the contrary: the programmer should let correctness proof and program grow hand in hand.


Stepwise refinement [50, 2, 31] is an important technique in correctness by construction. (It is often all that is needed, but developing correct OO code will require more machinery.) The bulk of this chapter is consequently devoted to freefinement – a new algorithm that constructs a sound refinement calculus from a verification system under certain conditions.

In fact many theories in computer science are presented, or approximated, by compositional *verification systems*. A verification system in this chapter is any formal system for establishing whether an inductively defined term, typically a program, satisfies a specification. For example, Hoare logics and type systems can be viewed as verification systems. In the case of Hoare logics, the system proves that a statement satisfies certain specifications given as preconditions and postconditions. In the case of type systems, the system proves that a term has a certain type in a type context.

Refinement systems play a similar role to verification systems, the main difference being that they relate terms to other terms, instead of terms and specifications. Another difference is that they typically include so-called specification terms. Intuitively, a term refines another if it is ‘better’, i.e. if it satisfies more specifications. *Refinement calculi* are formal systems for establishing refinements. For example, the calculus of Morgan [50] derives refinements between statements based on total correctness specifications. Starting from an appropriate specification statement, one can derive a correct algorithm for computing the factorial of a number by applying Morgan’s refinement rules.

The work on freefinement started from the observation that a Hoare logic and a refinement calculus for a command language do not have to be independent entities: once the Hoare logic is extended with specification statements, the two systems can be accommodated in a single theory. Moreover, there is a strong relation between the two systems. This chapter explains that this relation is not a coincidence: it is possible to analyse the structure of the inference rules defining a verification system, and automatically generate a related refinement calculus. Freefinement is an algorithm that implements this transformation. Surprisingly, freefinement is not limited to Hoare logics, but can be applied to any verification system whose inference rules satisfy certain conditions. Several refinement rules proposed in the literature in different contexts arise in this way.

The freefinement algorithm works as follows. Given a term language and an accompanying verification system V_1 that satisfies certain conditions, freefinement extends the term language with specification terms and builds a verification system V_2 for extended terms. The conditions on V_1 ensure that it is possible to extend the terms without breaking the inference rules; V_2 is consequently a sound and conservative extension of V_1 . Moreover, freefinement proposes a sound refinement system R that is in harmony with V_2 . Harmony means that the two formal systems can interoperate smoothly. It entails, for example, that a term satisfies a specification according to V_2 if and only if it is possible to refine the specification into the term with R . In fact, proof translation between V_2 and R becomes possible because harmony is demonstrated constructively. Freefinement internally constructs the refinement calculus by ‘linearising’ V_2 in a series of steps. The conditions on V_1 ensure that successful linearisation is possible. According to the presentation below, at most six steps are needed for this ‘refinement of refinement systems’. The situation is summarised as follows:



Freefinement requires no human intervention. The conditions it imposes are fulfilled by many program logics and type systems: examples include Hoare logic, separation logic, the simply-typed lambda calculus and System F. Freefinement defines the semantics of refinement at an abstract level: it associates each term of the extended language with a set of terms from the original language, and refinement simply reduces this set.

With freefinement, tools that are based on verification systems can readily include refinement as a complementary or alternative development style. In cases where the verification system is part of a larger system for establishing program correctness, it may well be possible to use refinement in a more elaborate calculus for constructing correct programs. The second part of this chapter shows how to use the output of freefinement in a calculus for developing correct OO code. This calculus is based on the separation logic proof system of Chapter 2, and maintains close relationships with it.

Chapter outline Section 5.1 describes the freefinement algorithm, which is applied in Section 5.2 to a simple type system for the lambda calculus and also to Hoare logic. Section 5.3 shows how the output of freefinement can be used in a calculus for constructing correct OO programs. Section 5.4 concludes with related work.

5.1 Freefinement

5.1.1 The Inputs

Freefinement requires four things as input:

1. A set of constructors $\mathbb{K}$. The constructors give rise to a term language T , where an arbitrary term t of T is defined by the grammar:

$$t ::= C(t_1, \dots, t_n)$$

where $C \in \mathbb{K}$.

2. A set of specifications $\mathbb{S}$.

3. A binary relation $\models_{V_1} \text{Sat } _$ between terms and specifications. Intuitively, $\models_{V_1} t \text{ Sat } S$ denotes that term $t \in T$ satisfies specification $S \in \mathbb{S}$.
4. A formal system $V_1(\mathbb{K}, \mathbb{S}, \models_{V_1} \text{Sat } _)$, which consists of a set of inference rules for proving sentences of the form $t \text{ Sat } S$. Each rule of V_1 must have the form A_1 or B_1 :

$$\begin{array}{c}
 A_1 \frac{t_1 \text{ Sat } S_1 \quad \dots \quad t_n \text{ Sat } S_n}{C(t_1, \dots, t_n) \text{ Sat } S} \\
 \text{provided } \text{Pred}(C, S_1, \dots, S_n, S). \\
 \\
 B_1 \frac{t \text{ Sat } S_1 \quad \dots \quad t \text{ Sat } S_m}{t \text{ Sat } S} \\
 \text{provided } \text{Pred}(S_1, \dots, S_m, S).
 \end{array}$$

The t 's, S 's and C in the rule forms indicate where the rules of V_1 must use metavariables. Thus a rule of form A_1 has only the freedom to choose a concrete n and a definition for its proviso predicate Pred ; the proviso predicate implements the side condition of the rule based on the arguments $C, S_1, \dots, S_n$ and S . A rule of form B_1 is also a pair: a concrete m and a definition of a predicate with arguments $S_1, \dots, S_m$ and S . Freefinement requires that the rules must be sound with respect to the following semantics:

Definition 5.1 (Semantics of the Inference Rules).

- 5.1.1. For rules of the form A_1 :
 $\text{Pred}(C, S_1, \dots, S_n, S) \Rightarrow [\forall t_1, \dots, t_n \in T \cdot \models_{V_1} t_1 \text{ Sat } S_1 \wedge \dots \wedge \models_{V_1} t_n \text{ Sat } S_n \Rightarrow \models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S]$
- 5.1.2. For rules of the form B_1 :
 $\text{Pred}(S_1, \dots, S_m, S) \Rightarrow [\forall t \in T \cdot \models_{V_1} t \text{ Sat } S_1 \wedge \dots \wedge \models_{V_1} t \text{ Sat } S_m \Rightarrow \models_{V_1} t \text{ Sat } S]$

The rule forms stipulate that the rules of V_1 must be highly compositional – a requirement that freefinement will exploit. For example, rules cannot inspect or constrain the t 's that appear in premises. This will allow freefinement to reuse the rules after specification terms are added to the term language.

Consider the following three rules over $\mathbb{K} = \{0, \text{succ}, \text{pred}\}$ and $\mathbb{S} = \{\mathbb{N}\}$, where n is a metavariable:

$$\begin{array}{lll}
 1 \frac{n : \mathbb{N}}{\text{succ}(n) : \mathbb{N}} & 2 \frac{\text{succ}(n) : \mathbb{N}}{\text{pred}(\text{succ}(n)) : \mathbb{N}} & 3 \frac{n : \mathbb{N}}{\text{pred}(n) : \mathbb{N}} \\
 & & \text{provided } \text{positive}(n).
 \end{array}$$

Rule 1 can be written in form A_1 with $n = 1$ by defining the proviso $\text{Pred}(C, S_1, S)$ as $C = \text{succ} \wedge S_1 = S = \mathbb{N}$. Rule 2 is unacceptable, because its premise inspects the term and requires it to match $\text{succ}(n)$. Rule 3 is also unacceptable, because it constrains the term in its proviso.

It will become clear later that the ‘structural’ rules of Hoare logic, such as the rule of consequence, are examples of rules of form B_1 . Other rules of Hoare logic, such as the assignment axiom and rule for sequential composition, have the form A_1 .

Let $\vdash_{V_1} t \text{ Sat } S$ denote that $t \text{ Sat } S$ is derivable with V_1 . The soundness of the rules with respect to the semantics of Definition 5.1 implies the soundness of V_1 :

Theorem 5.1 (Soundness of V_1). $\vdash_{V_1} t \text{ Sat } S \Rightarrow \models_{V_1} t \text{ Sat } S$

Proof. By induction on the derivation of $t \text{ Sat } S$:

- A rule of the form A_1 was last applied. Assume $\text{Pred}(C, S_1, \dots, S_n, S)$ and the induction hypothesis $\models_{V_1} t_1 \text{ Sat } S_1 \wedge \dots \wedge \models_{V_1} t_n \text{ Sat } S_n$. Then $\models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S$ by Definition 5.1.1.
- A rule of the form B_1 was last applied. Assume $\text{Pred}(S_1, \dots, S_m, S)$ and also the induction hypothesis $\models_{V_1} t \text{ Sat } S_1 \wedge \dots \wedge \models_{V_1} t \text{ Sat } S_m$. From Definition 5.1.2 follows $\models_{V_1} t \text{ Sat } S$. $\square$

Freefinement does not assume the completeness of V_1 , i.e. it never assumes $\models_{V_1} t \text{ Sat } S \Rightarrow \vdash_{V_1} t \text{ Sat } S$.

5.1.2 The Extended Language and Formal System

This section extends the language T with specification terms that are useful for refinement. It gives a semantics to the resulting language U , and extends V_1 in a sound and conservative way to prove sentences of the form $u \text{ Sat } S$ where $u \in U$.

The Extended Language U

Suppose $\mathbb{K}$ and $\mathbb{S}$ are disjoint (if they are not, then they can always be decorated to become disjoint) and do not contain a symbol $\sqcup$. The extended set of constructors

$$\mathbb{K}' = \mathbb{K} \cup \mathbb{S} \cup \{\sqcup \text{ with arity } n \mid n \in \mathbb{N}\}$$

gives rise to an extended language U , which can also be written as:

$$u ::= C(u_1, \dots, u_n) \mid S \mid \sqcup(u_1, \dots, u_n)$$

A term of the form S is called a *spec* term, and a term of the form $\sqcup(u_1, \dots, u_n)$ is called the *join* of $u_1, \dots, u_n$. Intuitively, S is a generic term that satisfies S , and $\sqcup(u_1, \dots, u_n)$ is a generic term that satisfies any S that any of the $u_1, \dots, u_n$ satisfy. Although the details will become clear later, the reasons for adding these terms are simple: the refinement system should be able to refine spec terms into other terms for top-down development, and join terms will be important for simplifying rules of the form B_1 where $m > 1$. If there are no rules of the form B_1 where $m > 1$, then join terms and their consequent treatment can be omitted.

A couple of constructs are used for giving a semantics to U . Let X denote a subset of T , and let Y denote a subset of $\mathcal{S}$. $\text{Specs}(X)$ is the set of all specifications that all the terms in X satisfy, and $\text{Terms}(Y)$ is the set of terms of T that satisfy all the specifications in Y :

Definition 5.2 (Specs and Terms).

- $\text{Specs}(X) \stackrel{\text{def}}{=} \{S \mid \forall t \in X. \models_{v_1} t \text{ Sat } S\}$
- $\text{Terms}(Y) \stackrel{\text{def}}{=} \{t \mid \forall S \in Y. \models_{v_1} t \text{ Sat } S\}$

An antitone Galois connection¹ exists between Specs and Terms:

Lemma 5.2. $X \subseteq \text{Terms}(Y) \Leftrightarrow Y \subseteq \text{Specs}(X)$

Proof. $X \subseteq \text{Terms}(Y)$
 $\Leftrightarrow \{\text{definition of Terms and } \subseteq\}$
 $\forall t \in X. \forall S \in Y. \models_{v_1} t \text{ Sat } S$
 $\Leftrightarrow \{\text{predicate calculus}\}$
 $\forall S \in Y. \forall t \in X. \models_{v_1} t \text{ Sat } S$
 $\Leftrightarrow \{\text{definition of Specs and } \subseteq\}$
 $Y \subseteq \text{Specs}(X)$ □

Antitone Galois connections have several well-known properties. For instance, $(\text{Terms} \circ \text{Specs})$ and $(\text{Specs} \circ \text{Terms})$ are extensive, increasing and idempotent and therefore closure operators. Freefinement relies on the following properties (their proofs appear in Appendix B):

Corollary 5.3.

5.3.1. $X \subseteq \text{Terms}(\text{Specs}(X))$

¹Also known as an *order-reversing* or *contravariant* Galois connection.

$$5.3.2. \text{Terms}(\text{Specs}(\text{Terms}(Y))) = \text{Terms}(Y)$$

$$5.3.3. \text{Specs}(X) \subseteq \text{Specs}(X') \Leftrightarrow \text{Terms}(\text{Specs}(X)) \supseteq \text{Terms}(\text{Specs}(X'))$$

$$5.3.4. \text{Terms}(Y \cup Y') = \text{Terms}(Y) \cap \text{Terms}(Y')$$

The following auxiliary definition provides a shorthand for the set of all terms of the form $C(t_1, \dots, t_n)$ where $t_1 \in X_1, \dots, t_n \in X_n$:

$$C(X_1, \dots, X_n) \stackrel{\text{def}}{=} \{C(t_1, \dots, t_n) \mid \bigwedge_{i \in 1..n} t_i \in X_i\}$$

For example, it yields a singleton set for nullary constructors:

$$\{C() \mid \bigwedge_{i \in 1..0} t_i \in X_i\} = \{C() \mid \text{True}\} = \{C()\}$$

The semantics of U is given by the function $\llbracket _ \rrbracket$ of type $U \rightarrow \mathcal{P}(T)$, i.e. every term in U denotes a set of terms from T :

Definition 5.3 (Semantics of U).

$$\begin{aligned} \llbracket C(u_1, \dots, u_n) \rrbracket &\stackrel{\text{def}}{=} \text{Terms}(\text{Specs}(C(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket))) \\ \llbracket S \rrbracket &\stackrel{\text{def}}{=} \text{Terms}(\{S\}) \\ \llbracket \bigsqcup(u_1, \dots, u_n) \rrbracket &\stackrel{\text{def}}{=} \bigcap_{i \in 1..n} \llbracket u_i \rrbracket \end{aligned}$$

If the relation $\models_{V_1} \text{Sat}$ is well-behaved in a sense that will be made precise later, then $\llbracket u \rrbracket$ has a simple intuitive explanation: it denotes the set of all primitive terms, i.e. terms from T , that refine u . For a term $C(u_1, \dots, u_n)$, first consider $C(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket)$ – the set of terms of the form $C(t_1, \dots, t_n)$ where $t_1 \in \llbracket u_1 \rrbracket$ (i.e. t_1 refines u_1) and $\dots$ and $t_n \in \llbracket u_n \rrbracket$. All the specifications that all these terms implement are then collected, and any primitive term that satisfies all such specifications refines $C(u_1, \dots, u_n)$. The primitive terms that refine S are exactly those that satisfy S . Finally, $\bigsqcup(u_1, \dots, u_n)$ is refined by any primitive term that refines all $u_1, \dots, u_n$.

For all u , the set $\llbracket u \rrbracket$ is a fixpoint of $\text{Terms} \circ \text{Specs}$ and hence a closed element:

Lemma 5.4. $\text{Terms}(\text{Specs}(\llbracket u \rrbracket)) = \llbracket u \rrbracket$

Proof. By induction on the structure of u :

- If u has the form $C(u_1, \dots, u_n)$ or S , then $\llbracket u \rrbracket = \text{Terms}(Y)$ for some Y and the result follows by Corollary 5.3.2.

- If u has the form $\sqcup(u_1, \dots, u_n)$, assume $\llbracket u_i \rrbracket = \text{Terms}(\text{Specs}(\llbracket u_i \rrbracket))$ for all $i \in 1..n$. So $\llbracket \sqcup(u_1, \dots, u_n) \rrbracket = \bigcap_{i \in 1..n} \text{Terms}(\text{Specs}(\llbracket u_i \rrbracket)) = \text{Terms}(\bigcup_{i \in 1..n} \text{Specs}(\llbracket u_i \rrbracket))$ by Corollary 5.3.4, and Corollary 5.3.2 concludes the proof. $\square$

The rest of the chapter introduces further properties of the semantics as needed.

Extending V_1 : Preliminaries

The next section will extend V_1 to obtain a formal system V_2 for proving sentences of the form $u \text{ Sat } S$. The aim is to construct a sound and conservative extension of V_1 . Informally, a sound extension of V_1 must have equal or more power:

Definition 5.4 (Sound Extension). $V_2(\mathbb{K}', \mathbb{S}', \models_{V_2} \text{ Sat } -)$ is a sound extension of $V_1(\mathbb{K}, \mathbb{S}, \models_{V_1} \text{ Sat } -)$ if and only if

1. V_2 uses richer terms and specifications:
 $\mathbb{K} \subseteq \mathbb{K}'$ and $\mathbb{S} \subseteq \mathbb{S}'$
2. V_2 can prove everything that V_1 can prove:
 $\forall t \in T, S \in \mathbb{S} \cdot \vdash_{V_1} t \text{ Sat } S \Rightarrow \vdash_{V_2} t \text{ Sat } S$
3. V_2 uses a richer semantics:
 $\forall t \in T, S \in \mathbb{S} \cdot \models_{V_2} t \text{ Sat } S \Rightarrow \models_{V_1} t \text{ Sat } S$
4. V_2 is sound:
 $\forall u \in U, S' \in \mathbb{S}' \cdot \vdash_{V_2} u \text{ Sat } S' \Rightarrow \models_{V_2} u \text{ Sat } S'$

As a consequence, $\forall t \in T, S \in \mathbb{S} \cdot \vdash_{V_2} t \text{ Sat } S \Rightarrow \models_{V_1} t \text{ Sat } S$, which intuitively means that V_2 restricted to $\mathbb{K}$ and $\mathbb{S}$ is sound with respect to the semantics of V_1 .

In a sound and conservative extension, the converse of requirement 2 also holds:

Definition 5.5 (Sound and Conservative Extension). A formal system $V_2(\mathbb{K}', \mathbb{S}', \models_{V_2} \text{ Sat } -)$ is a sound and conservative extension of $V_1(\mathbb{K}, \mathbb{S}, \models_{V_1} \text{ Sat } -)$ if and only if

1. V_2 is a sound extension of V_1 .
2. V_1 and V_2 restricted to $\mathbb{K}$ and $\mathbb{S}$ have equal derivability:
 $\forall t \in T, S \in \mathbb{S} \cdot \vdash_{V_1} t \text{ Sat } S \Leftrightarrow \vdash_{V_2} t \text{ Sat } S$

Although a sound and conservative extension cannot prove more sentences of the form $t \text{ Sat } S$, it is still useful for extending the term language and installing a richer semantics. It can also extend the specifications, but the V_2 of the next section will simply use S .

The Extended Formal System V_2

The construction of V_2 starts with the empty set of rules and proceeds in two steps:

1. For each rule of V_1 , replace t 's by u 's and add the resulting rule. This change of metavariables yields the rule forms A_2 and B_2 in V_2 :

$$A_2 \frac{u_1 \text{ Sat } S_1 \quad \dots \quad u_n \text{ Sat } S_n}{C(u_1, \dots, u_n) \text{ Sat } S}$$

provided $\text{Pred}(C, S_1, \dots, S_n, S)$.

$$B_2 \frac{u \text{ Sat } S_1 \quad \dots \quad u \text{ Sat } S_m}{u \text{ Sat } S}$$

provided $\text{Pred}(S_1, \dots, S_m, S)$.

2. Add the following rules for spec terms and joins:

$$\text{SPEC} \frac{}{S \text{ Sat } S}$$

$$\text{JOIN} \frac{u \text{ Sat } S}{\bigsqcup(\dots, u, \dots) \text{ Sat } S}$$

By induction on the derivation, V_1 and V_2 are equivalent with respect to derivability on T , i.e. $\vdash_{V_1} t \text{ Sat } S \Leftrightarrow \vdash_{V_2} t \text{ Sat } S$. So for V_2 to be a sound and conservative extension of V_1 , it will suffice to equip V_2 with a richer semantics and to prove it sound.

The *Sat* relation between U and S is defined as follows:

Definition 5.6 (Extended Satisfaction).

$$\models_{V_2} u \text{ Sat } S \stackrel{\text{def}}{=} \forall t \in \llbracket u \rrbracket \cdot \models_{V_1} t \text{ Sat } S$$

Furthermore, the U -semantics of t contains t as an element:

Lemma 5.5 (Term Embedding). $\forall t \in T \cdot t \in \llbracket t \rrbracket$

Proof. By induction on the structure of t . Suppose $t = C(t_1, \dots, t_n)$ and assume $t_1 \in \llbracket t_1 \rrbracket, \dots, t_n \in \llbracket t_n \rrbracket$. So $t \in C(\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket)$, which is a subset of $\text{Terms}(\text{Specs}(C(\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket)))$ by Corollary 5.3.1. $\square$

Therefore $\models_{V_2} t \text{ Sat } S \Rightarrow \models_{V_1} t \text{ Sat } S$ holds, and the soundness proof of V_2 establishes that V_2 is a sound and conservative extension of V_1 :

Theorem 5.6 (Soundness of V_2). $\vdash_{V_2} u \text{ Sat } S \Rightarrow \models_{V_2} u \text{ Sat } S$

Proof. By induction on the structure of the derivation:

- For each rule of the form A_2 , assume $\text{Pred}(C, S_1, \dots, S_n, S)$ and assume

$$\begin{aligned} \forall t_1 \in \llbracket u_1 \rrbracket \cdot \models_{V_1} t_1 \text{ Sat } S_1 \\ \vdots \\ \forall t_n \in \llbracket u_n \rrbracket \cdot \models_{V_1} t_n \text{ Sat } S_n \end{aligned}$$

So $\forall t_1 \in \llbracket u_1 \rrbracket, \dots, t_n \in \llbracket u_n \rrbracket \cdot \models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S$ because the corresponding rule of the form A_1 in V_1 is sound with respect to Definition 5.1.1. So $S \in \text{Specs}(C(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket))$ and hence $\forall t \in \text{Terms}(\text{Specs}(C(\llbracket u_1 \rrbracket, \dots, \llbracket u_n \rrbracket))) \cdot \models_{V_1} t \text{ Sat } S$.

- For each rule of the form B_2 , assume $\text{Pred}(S_1, \dots, S_m, S)$ and assume $\forall t \in \llbracket u \rrbracket \cdot \models_{V_1} t \text{ Sat } S_1 \wedge \dots \wedge \models_{V_1} t \text{ Sat } S_m$.

Now $\forall t \in \llbracket u \rrbracket \cdot \models_{V_1} t \text{ Sat } S$ because the corresponding rule of the form B_1 in V_1 is sound with respect to Definition 5.1.2.

- SPEC: $\forall t \in \text{Terms}(\{S\}) \cdot \models_{V_1} t \text{ Sat } S$ by definition.
- JOIN: Assume $\forall t \in \llbracket u \rrbracket \cdot \models_{V_1} t \text{ Sat } S$. If $t \in \llbracket \sqcup(\dots, u, \dots) \rrbracket$ then $t \in \llbracket u \rrbracket$ and hence $\models_{V_1} t \text{ Sat } S$. $\square$

Extended satisfaction has an alternative characterisation that freefine-ment will also use:

Lemma 5.7. $\models_{V_2} u \text{ Sat } S \Leftrightarrow S \in \text{Specs}(\llbracket u \rrbracket)$

Proof.

$$\begin{aligned} & \models_{V_2} u \text{ Sat } S \\ \Leftrightarrow & \text{ \{definition\} } \\ & \forall t \in \llbracket u \rrbracket \cdot \models_{V_1} t \text{ Sat } S \\ \Leftrightarrow & \text{ \{definition of Specs\} } \\ & S \in \text{Specs}(\llbracket u \rrbracket) \end{aligned}$$

$\square$

5.1.3 System V_2 and Refinement

The next section will construct several refinement systems, or calculi, that are based on V_2 . These refinement systems are formal systems for proving sentences of the form $u \sqsubseteq u'$. The definition of the refinement relation makes the semantics of refinement precise:

Definition 5.7 (Refinement). $\models u \sqsubseteq u' \stackrel{\text{def}}{=} \llbracket u \rrbracket \supseteq \llbracket u' \rrbracket$

This definition leads to simple proofs, and is equivalent to several other formulations. The following theorem states one such alternative, and its proof mentions others:

Lemma 5.8 (Equivalent Characterisation of Refinement).

$$\models u \sqsubseteq u' \Leftrightarrow \forall S \cdot \models_{V_2} u \text{ Sat } S \Rightarrow \models_{V_2} u' \text{ Sat } S$$

Proof.

$$\begin{aligned} & \llbracket u \rrbracket \supseteq \llbracket u' \rrbracket \\ \Leftrightarrow & \{ \text{Lemma 5.4} \} \\ & \text{Terms}(\text{Specs}(\llbracket u \rrbracket)) \supseteq \text{Terms}(\text{Specs}(\llbracket u' \rrbracket)) \\ \Leftrightarrow & \{ \text{Corollary 5.3.3} \} \\ & \text{Specs}(\llbracket u \rrbracket) \subseteq \text{Specs}(\llbracket u' \rrbracket) \\ \Leftrightarrow & \{ \text{definition of } \sqsubseteq \} \\ & \forall S \cdot S \in \text{Specs}(\llbracket u \rrbracket) \Rightarrow S \in \text{Specs}(\llbracket u' \rrbracket) \\ \Leftrightarrow & \{ \text{Lemma 5.7} \} \\ & \forall S \cdot \models_{V_2} u \text{ Sat } S \Rightarrow \models_{V_2} u' \text{ Sat } S \quad \square \end{aligned}$$

If $\models_{V_1} \text{ Sat } _$ is well-behaved, then there is also another explanation for defining $\models u \sqsubseteq u'$ as $\llbracket u \rrbracket \supseteq \llbracket u' \rrbracket$: u' refines u iff every primitive term that refines u' also refines u . Put differently, u' refines u iff u' constrains the set of eventual primitive terms that refinement can produce to the same or higher degree compared to u . So u can be seen as a placeholder for any of the primitive terms in $\llbracket u \rrbracket$, and the role of refinement is to reduce the uncertainty.

Many examples of refinements will follow later, so here is a small one: a join term implements the least upper bound (join) of its immediate subterms with respect to $\sqsubseteq$, hence the name. In particular:

1. $\forall i \in 1..n \cdot \models u_i \sqsubseteq \bigsqcup(u_1, \dots, u_n)$
2. If $(\forall i \in 1..n \cdot \models u_i \sqsubseteq u)$, then $\models \bigsqcup(u_1, \dots, u_n) \sqsubseteq u$.

The notation $u \equiv u'$ is a shorthand for $\llbracket u \rrbracket = \llbracket u' \rrbracket$, which is equivalent to $\models u \sqsubseteq u' \wedge \models u' \sqsubseteq u$.

A refinement system R will be sound if and only if $\vdash_R u \sqsubseteq u'$ implies $\models u \sqsubseteq u'$. In the next section, freefinement will construct several sound refinement systems where each system R is related to V_2 by the properties Harmony 1 and 2 below.

Harmony 1. If $\vdash_{V_2} u \text{ Sat } S$ and $\vdash_R u \sqsubseteq u'$, then $\vdash_{V_2} u' \text{ Sat } S$.

Intuitively, Harmony 1 says that V_2 contains sufficient machinery to prove the same properties about u' that it could prove about u . In other words, R is not too powerful for V_2 .

Harmony 2. If $\vdash_{V_2} u \text{ Sat } S$, then $\vdash_R S \sqsubseteq u$.

Intuitively, Harmony 2 means that the refinement system R contains sufficient machinery to refine a specification into any term that satisfies it according to V_2 . In other words, V_2 is embedded in R and hence R is not too weak.

Harmony 1 is stronger than the converse of Harmony 2:

Theorem 5.9. If V_2 and a refinement system R are related by Harmony 1, then $\vdash_R S \sqsubseteq u \Rightarrow \vdash_{V_2} u \text{ Sat } S$.

Proof. Assume $\vdash_R S \sqsubseteq u$. Since $\vdash_{V_2} S \text{ Sat } S$ by SPEC, it follows from Harmony 1 that $\vdash_{V_2} u \text{ Sat } S$. $\square$

A refinement system R is called *harmonic* iff it satisfies Harmony 1 and 2. Harmonic refinement systems interoperate nicely with V_2 . In fact, the proofs of Harmony 1 and 2 in the next section are constructive in the sense that they enable proof translation. Given a V_2 -proof of $u \text{ Sat } S$ and an R -proof of $u \sqsubseteq u'$, they describe a V_2 -proof of $u' \text{ Sat } S$. Based on a V_2 -proof of $u \text{ Sat } S$, they show how to build an R -proof for $S \sqsubseteq u$. Since Harmony 1 is established constructively, given an R -proof of $S \sqsubseteq u$, the proof of Theorem 5.9 shows how to build a V_2 -proof for $u \text{ Sat } S$.

The final refinement system that free refinement produces will also have a specific desired form. This form guarantees that refinement proofs are ‘linear’ developments where terms can be refined in-place. Formally, a refinement system has the desired form if the rules with premises describe either the transitivity or the monotonicity of refinement. All the other rules must be axioms, i.e. without any premise.

5.1.4 The Refinement of Refinement Systems

V_2 can be linearised in a series of steps to obtain a sound and harmonic refinement system of the desired form. At most six steps are necessary according to this presentation – the exact number depends on V_1 . The steps make it easy to prove and maintain soundness and harmony, which would otherwise be more complex to establish for the final refinement calculus.

Many of the steps take a previously constructed refinement system and add or remove rules to obtain a new system. If a sound and harmonic refinement system is extended with a rule that is sound and respects Harmony 1, then the resulting system will be sound and harmonic. There is no need to prove Harmony 2 again, because the new refinement system can still derive all sentences that the old one could derive. If a rule is removed from a sound and harmonic refinement system, then the resulting system remains sound and will also be harmonic if it satisfies Harmony 2. A simple way of showing that Harmony 2 still holds is to show that any application of the old rule can be achieved by a combination of rules that remain in the system.

Getting Started: R_1

The first refinement system R_1 is obtained from V_2 by a simple syntactic transformation: each sentence $u \text{ Sat } S$ becomes $S \sqsubseteq u$. R_1 has rules of the form A_3 and B_3 , a SPEC rule and also a JOIN rule if join terms were needed:

$$\begin{array}{c}
 A_3 \frac{S_1 \sqsubseteq u_1 \quad \dots \quad S_n \sqsubseteq u_n}{S \sqsubseteq C(u_1, \dots, u_n)} \\
 \text{provided } \text{Pred}(C, S_1, \dots, S_n, S). \\
 \\
 B_3 \frac{S_1 \sqsubseteq u \quad \dots \quad S_m \sqsubseteq u}{S \sqsubseteq u} \\
 \text{provided } \text{Pred}(S_1, \dots, S_m, S). \\
 \\
 \text{SPEC } \frac{}{S \sqsubseteq S} \\
 \\
 \text{JOIN } \frac{S \sqsubseteq u}{S \sqsubseteq \sqcup(\dots, u, \dots)}
 \end{array}$$

V_2 and R_1 are isomorphic: a proof of $u \text{ Sat } S$ in V_2 corresponds to a proof of $S \sqsubseteq u$ in R_1 and vice versa, so $\vdash_{V_2} u \text{ Sat } S \Leftrightarrow \vdash_{R_1} S \sqsubseteq u$. The soundness proof of R_1 relies on the following equivalence:

Lemma 5.10. $\models_{V_2} u \text{ Sat } S \Leftrightarrow \models S \sqsubseteq u$

Proof.

$$\begin{aligned}
 & \models_{V_2} u \text{ Sat } S \\
 \Leftrightarrow & \quad \{\text{Lemma 5.7}\} \\
 & \{S\} \subseteq \text{Specs}(\llbracket u \rrbracket) \\
 \Leftrightarrow & \quad \{\text{Lemma 5.2}\} \\
 & \llbracket u \rrbracket \subseteq \text{Terms}(\{S\})
 \end{aligned}$$

□

Theorem 5.11 (Soundness of R_1). $\vdash_{R_1} u \sqsubseteq u' \Rightarrow \models u \sqsubseteq u'$

Proof. If $\vdash_{R_1} u \sqsubseteq u'$, then u has the form S and $\vdash_{V_2} u' \text{ Sat } S$. The soundness of V_2 implies $\models_{V_2} u' \text{ Sat } S$, and Lemma 5.10 in turn implies $\models S \sqsubseteq u'$. $\square$

Theorem 5.12. R_1 is harmonic.

Proof. Harmony 2 holds by construction. For Harmony 1, assume $\vdash_{R_1} u \sqsubseteq u'$. Then u has the form S'' and $\vdash_{V_2} u' \text{ Sat } S''$ by construction. That $\vdash_{V_2} u \text{ Sat } S'$ (i.e. $\vdash_{V_2} S'' \text{ Sat } S'$) implies $\vdash_{V_2} u' \text{ Sat } S'$ for all S' follows by induction on the derivation of $S'' \text{ Sat } S'$:

- SPEC: S' and S'' are the same. Since $\vdash_{V_2} u' \text{ Sat } S''$, it holds that $\vdash_{V_2} u' \text{ Sat } S'$.
- For each rule of the form B_2 : S and S' are the same. Assume $\text{Pred}(S_1, \dots, S_m, S)$, $\vdash_{V_2} u \text{ Sat } S_1, \dots, \vdash_{V_2} u \text{ Sat } S_m$, and by the induction hypothesis also $\vdash_{V_2} u' \text{ Sat } S_1, \dots, \vdash_{V_2} u' \text{ Sat } S_m$. So the rule being considered is applicable and $\vdash_{V_2} u' \text{ Sat } S$. Hence $\vdash_{V_2} u' \text{ Sat } S'$. $\square$

Note: if V_2 has only rules of the form A_2 where $n = 0$ and/or rules of the form B_2 where $m = 0$, then R_1 is a refinement system of the desired form and refinement stops.

Adding Transitivity: R_2

The refinement system R_2 extends R_1 with the rule TRANS which states that refinement is transitive:

$$\text{TRANS} \frac{u_1 \sqsubseteq u_2 \quad u_2 \sqsubseteq u_3}{u_1 \sqsubseteq u_3}$$

TRANS is sound because $\supseteq$ is transitive, and it maintains Harmony 1 since implication is transitive. So R_2 is sound and harmonic.

Simplification: R_3

The presence of SPEC and TRANS in R_2 allows the simplification of rules of the form B_3 with $m = 1$:

$$B_3 \frac{S_1 \sqsubseteq u}{S \sqsubseteq u} \\ \text{provided } \text{Pred}(S_1, S).$$

For an arbitrary rule of this form, consider the derivation

$$\begin{array}{c} \text{SPEC} \frac{}{\overline{S_1 \sqsubseteq S_1}} \\ B_3 \frac{}{\overline{S \sqsubseteq S_1}} \\ \text{provided } \text{Pred}(S_1, S). \end{array}$$

By virtue of having been derived, the new rule

$$\begin{array}{c} B_3 \frac{}{\overline{S \sqsubseteq S_1}} \\ \text{provided } \text{Pred}(S_1, S). \end{array}$$

is sound and maintains Harmony 1, and can therefore be added to R_2 to obtain a sound and harmonic refinement system. In fact, it can replace the old version without breaking Harmony 2, since removing the old version will not decrease the derivable set of sentences: every application of the old B_3 can be changed into:

$$\text{TRANS} \frac{B_3 \frac{}{\overline{S \sqsubseteq S_1}} \quad S_1 \sqsubseteq u}{\overline{S \sqsubseteq u}}$$

since $\text{Pred}(S_1, S)$ is guaranteed.

The refinement system R_3 is the same as R_2 , except that the rules of the form B_3 with $m = 1$ are replaced by their simplified versions. R_3 is sound and harmonic.

Note: if V_2 has only rules of the form A_2 where $n = 0$ and rules of the form B_2 where $m \leq 1$, then R_3 is a refinement system of the desired form and freefinement stops.

Adding Monotonicity: R_4

All the constructors of U are monotone with respect to $\sqsubseteq$, i.e. the following rules are sound:

$$\begin{array}{c} C-i \frac{u_i \sqsubseteq u'_i}{\overline{C(u_1, \dots, u_i, \dots, u_n) \sqsubseteq C(u_1, \dots, u'_i, \dots, u_n)}} \\ \\ \text{JOIN-}i \frac{u_i \sqsubseteq u'_i}{\overline{\sqcup(u_1, \dots, u_i, \dots, u_n) \sqsubseteq \sqcup(u_1, \dots, u'_i, \dots, u_n)}} \end{array}$$

Moreover, these rules maintain harmony:

Lemma 5.13. $C-i$ maintains Harmony 1.

Proof. Assume $\forall S' \cdot \vdash_{V_2} u_i \text{ Sat } S' \Rightarrow \vdash_{V_2} u'_i \text{ Sat } S'$. That $\forall S \cdot \vdash_{V_2} C(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S \Rightarrow \vdash_{V_2} C(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$ follows by induction on the derivation of $C(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S$:

- A₂: Suppose $\vdash_{V_2} u_j \text{ Sat } S_j$ for $j \in 1..n$, and also suppose $\text{Pred}(C, S_1, \dots, S_n, S)$ holds. Since $\vdash_{V_2} u'_i \text{ Sat } S_i$, the same rule A₂ can be applied to derive $C(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$.
- B₂: Suppose $\vdash_{V_2} C(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S_j$ for $j \in 1..m$, and suppose $\text{Pred}(S_1, \dots, S_m, S)$. The induction hypothesis is the assumption $\vdash_{V_2} C(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S_j$ for $j \in 1..m$. Since $\text{Pred}(S_1, \dots, S_m, S)$, the same rule B₂ is applicable and hence $\vdash_{V_2} C(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$. $\square$

Lemma 5.14. JOIN-*i* maintains Harmony 1.

Proof. Assume $\forall S' \cdot \vdash_{V_2} u_i \text{ Sat } S' \Rightarrow \vdash_{V_2} u'_i \text{ Sat } S'$. That $\forall S \cdot \vdash_{V_2} \sqcup(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S \Rightarrow \vdash_{V_2} \sqcup(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$ follows by induction on the derivation of $\sqcup(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S$:

- JOIN: Suppose $u_j \text{ Sat } S$ was the premise for some $j \in 1..n$. If $j \neq i$, then apply JOIN to the premise $u_j \text{ Sat } S$ to derive the required $\sqcup(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$. If $j = i$, then by assumption $\vdash_{V_2} u'_i \text{ Sat } S$ holds, and the result follows by JOIN.
- B₂: Suppose $\vdash_{V_2} \sqcup(u_1, \dots, u_i, \dots, u_n) \text{ Sat } S_j$ for $j \in 1..m$, and suppose $\text{Pred}(S_1, \dots, S_m, S)$. The induction hypothesis is the assumption $\vdash_{V_2} \sqcup(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S_j$ for $j \in 1..m$. Since $\text{Pred}(S_1, \dots, S_m, S)$, the same rule B₂ is applicable and hence $\vdash_{V_2} \sqcup(u_1, \dots, u'_i, \dots, u_n) \text{ Sat } S$. $\square$

Let the notation $v[u]$ denote a term in U whose parse tree is factored into two parts: a core tree v with a ‘hole’ where the sub-tree for u fits. The rule MONO packages C-*i* and JOIN-*i* in a single convenient form:

$$\text{MONO} \frac{u \sqsubseteq u'}{v[u] \sqsubseteq v[u']}$$

Informally, the rule MONO allows in-place refinement: if u_0 can be factored as $v[u]$, and u' refines u , then $v[u']$ refines u_0 .

MONO is sound and maintains harmony because C-*i* and JOIN-*i* are sound and maintain harmony. The refinement system R_4 extends R_3 with MONO. It is sound and harmonic.

Simplification: R_5

The rule MONO makes it possible to simplify:

- The JOIN rule:

$$\text{JOIN} \frac{S \sqsubseteq u}{S \sqsubseteq \bigsqcup(\dots, u, \dots)}$$

- Rules of the form A_3 with $n \geq 1$:

$$A_3 \frac{S_1 \sqsubseteq u_1 \quad \dots \quad S_n \sqsubseteq u_n}{S \sqsubseteq C(u_1, \dots, u_n)} \\ \text{provided } \text{Pred}(C, S_1, \dots, S_n, S).$$

Consider the derivation:

$$\text{JOIN} \frac{\text{SPEC} \overline{S \sqsubseteq S}}{S \sqsubseteq \bigsqcup(\dots, S, \dots)}$$

By virtue of having been derived, the simplified rule

$$\text{JOIN} \overline{S \sqsubseteq \bigsqcup(\dots, S, \dots)}$$

is sound and respects Harmony 1. It can replace the old version of JOIN without decreasing derivability, because any application of the old version can be achieved by:

$$\text{TRANS} \frac{\text{JOIN} \overline{S \sqsubseteq \bigsqcup(\dots, S, \dots)} \quad \text{MONO} \frac{S \sqsubseteq u}{\bigsqcup(\dots, S, \dots) \sqsubseteq \bigsqcup(\dots, u, \dots)}}{S \sqsubseteq \bigsqcup(\dots, u, \dots)}$$

Likewise, for each rule of the form A_3 , the derived rule

$$A_3 \overline{S \sqsubseteq C(S_1, \dots, S_n)} \\ \text{provided } \text{Pred}(C, S_1, \dots, S_n, S).$$

is sound and respects harmony. It makes the old version redundant, since any application of the old rule can be replaced by:

$$\begin{array}{c} A_3 \overline{S \sqsubseteq C(S_1, \dots, S_n)} \\ \boxed{E_1} \\ \dots \\ \boxed{E_n} \end{array}$$

where E_i is given by:

$$\boxed{\text{TRANS} \frac{S \sqsubseteq C(u_1, \dots, u_{i-1}, S_i, \dots, S_n) \quad P_i}{S \sqsubseteq C(u_1, \dots, u_i, S_{i+1}, \dots, S_n)}}$$

and P_i is the proof tree:

$$\text{MONO} \frac{S_i \sqsubseteq u_i}{C(u_1, \dots, u_{i-1}, S_i, \dots, S_n) \sqsubseteq C(u_1, \dots, u_i, S_{i+1}, \dots, S_n)}$$

Apart from these simplifications, the refinement system R_5 is the same as R_4 . It is sound and harmonic.

Note: if V_2 does not include rules of the form B_2 where $m > 1$, then R_5 has the desired form and freefinement stops.

Wrapping Up: R_6

It remains to simplify rules of the form B_3 with $m > 1$:

$$B_3 \frac{S_1 \sqsubseteq u \quad \dots \quad S_m \sqsubseteq u}{S \sqsubseteq u} \\ \text{provided } \text{Pred}(S_1, \dots, S_m, S).$$

If $\text{Pred}(S_1, \dots, S_m, S)$, then R_5 can derive:

$$B_3 \frac{\text{JOIN} \frac{}{S_1 \sqsubseteq \sqcup(S_1, \dots, S_m)} \quad \dots \quad \text{JOIN} \frac{}{S_m \sqsubseteq \sqcup(S_1, \dots, S_m)}}{S \sqsubseteq \sqcup(S_1, \dots, S_m)}$$

The derived rule

$$B_3 \frac{}{S \sqsubseteq \sqcup(S_1, \dots, S_m)} \\ \text{provided } \text{Pred}(S_1, \dots, S_m, S).$$

is therefore sound and respects Harmony 1. Together with the rule:

$$\text{UNJOIN} \frac{}{\sqcup(u, \dots, u) \sqsubseteq u}$$

which is trivially sound and respects Harmony 1, it can replace the old B_3 because any application of the old rule can be rewritten as:

$$\begin{array}{c}
\text{B}_3 \frac{\overline{S \sqsubseteq \bigsqcup(S_1, \dots, S_m)}}{\boxed{F_1}} \\
\vdots \\
\text{TRANS} \frac{\boxed{F_m} \quad G}{S \sqsubseteq u}
\end{array}$$

where G is UNJOIN, F_i is given by:

$$\boxed{\text{TRANS} \frac{S \sqsubseteq \bigsqcup(\overbrace{u, \dots, u}^{i-1}, S_i, \dots, S_m) \quad Q_i}{S \sqsubseteq \bigsqcup(u, \dots, u, S_{i+1}, \dots, S_m)}}$$

and Q_i is the proof tree:

$$\text{MONO} \frac{S_i \sqsubseteq u}{\bigsqcup(u, \dots, u, S_i, \dots, S_m) \sqsubseteq \bigsqcup(\underbrace{u, \dots, u}_i, S_{i+1}, \dots, S_m)}$$

R_6 is the same as R_5 , except that it includes UNJOIN and replaces rules of the form B_3 where $m > 1$ with their simplified versions. R_6 is sound, harmonic and of the desired form.

5.1.5 Discussion

R_6 can be made more powerful in several ways. For example, the following generalisation of JOIN is sound and preserves Harmony 1:

$$\text{JOIN}' \frac{}{u \sqsubseteq \bigsqcup(\dots, u, \dots)}$$

The same holds for the reflexivity of refinement, which generalises SPEC, and other rules such as UNNEST:

$$\text{UNNEST} \frac{}{\bigsqcup(u_1, \dots, u_n) \sqsubseteq \bigsqcup(u_1, \dots, u_{i-1}, u'_1, \dots, u'_m, u_{i+1}, \dots, u_n)}$$

provided $1 \leq i \leq n$ and $u_i = \bigsqcup(u'_1, \dots, u'_m)$.

In specific applications of freefinement, it might also be useful to add derived rules to R_6 . Examples of this will follow later.

Freefinement assumes as little as possible about $\models_{V_1} \text{Sat } _$ and is consequently very generic. As one might expect, additional assumptions can help to construct more powerful refinement systems. For example, suppose ‘plus’ is a constructor that is commutative in the sense that

$$\forall t_1, t_2 \in T, S \in \mathbb{S} \cdot \models_{V_1} \text{plus}(t_1, t_2) \text{ Sat } S \Leftrightarrow \models_{V_1} \text{plus}(t_2, t_1) \text{ Sat } S$$

Then $\text{Specs}(\text{plus}(\llbracket u_1 \rrbracket, \llbracket u_2 \rrbracket)) = \text{Specs}(\text{plus}(\llbracket u_2 \rrbracket, \llbracket u_1 \rrbracket))$ because

$$\begin{aligned} & S \in \text{Specs}(\text{plus}(\llbracket u_1 \rrbracket, \llbracket u_2 \rrbracket)) \\ \Leftrightarrow & \forall t_1 \in \llbracket u_1 \rrbracket, t_2 \in \llbracket u_2 \rrbracket \cdot \models_{V_1} \text{plus}(t_1, t_2) \text{ Sat } S \\ \Leftrightarrow & \forall t_1 \in \llbracket u_1 \rrbracket, t_2 \in \llbracket u_2 \rrbracket \cdot \models_{V_1} \text{plus}(t_2, t_1) \text{ Sat } S \\ \Leftrightarrow & S \in \text{Specs}(\text{plus}(\llbracket u_2 \rrbracket, \llbracket u_1 \rrbracket)) \end{aligned}$$

So $\llbracket \text{plus}(u_1, u_2) \rrbracket = \llbracket \text{plus}(u_2, u_1) \rrbracket$ and therefore the refinement rule $\text{plus}(u_1, u_2) \equiv \text{plus}(u_2, u_1)$ is sound. Depending on the rules of V_1 , it might also preserve harmony.

As mentioned before, the semantic function $\llbracket _ \rrbracket$ and the refinement order $\sqsubseteq$ have nice interpretations when $\models_{V_1} \text{Sat } _$ is well-behaved. Here is the definition:

Definition 5.8 (Well-behavedness). $\models_{V_1} \text{Sat } _$ is well-behaved iff
 $\forall C \in \mathbb{K}, t_1, \dots, t_n \in T, S \in \mathbb{S} \cdot \models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S \Rightarrow$
 $\forall t \in C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))) \cdot \models_{V_1} t \text{ Sat } S$

There is also an alternative characterisation of well-behavedness:

Lemma 5.15. $\models_{V_1} \text{Sat } _$ is well-behaved iff
 $\forall C \in \mathbb{K}, t_1, \dots, t_n \in T \cdot \text{Terms}(\text{Specs}(\{C(t_1, \dots, t_n)\})) =$
 $\text{Terms}(\text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))))$

Proof. $t_i \in \text{Terms}(\text{Specs}(\{t_i\}))$ for $i \in 1..n$ by Corollary 5.3.1, so
 $\{C(t_1, \dots, t_n)\} \subseteq C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\})))$. Hence
 by Corollary B.1.3 in Appendix B, $\text{Specs}(\{C(t_1, \dots, t_n)\}) \supseteq$
 $\text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))))$.

Therefore:

$$\begin{aligned} & \models_{V_1} \text{Sat } _ \text{ is well-behaved} \\ \Leftrightarrow & \\ & \forall C \in \mathbb{K}, t_1, \dots, t_n \in T, S \in \mathbb{S} \cdot \models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S \Rightarrow \\ & \forall t \in C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))) \cdot \models_{V_1} t \text{ Sat } S \\ \Leftrightarrow & \\ & \forall C \in \mathbb{K}, t_1, \dots, t_n \in T, S \in \mathbb{S} \cdot S \in \text{Specs}(\{C(t_1, \dots, t_n)\}) \Rightarrow \\ & S \in \text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\})))) \\ \Leftrightarrow & \\ & \forall C \in \mathbb{K}, t_1, \dots, t_n \in T \cdot \text{Specs}(\{C(t_1, \dots, t_n)\}) \subseteq \\ & \text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\})))) \\ \Leftrightarrow & \quad \{\text{by the reasoning above}\} \\ & \forall C \in \mathbb{K}, t_1, \dots, t_n \in T \cdot \text{Specs}(\{C(t_1, \dots, t_n)\}) = \\ & \text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\})))) \end{aligned}$$

The result then follows by Corollary 5.3.3. □

Freefinement does not require well-behavedness of $\models_{V_1-} \text{Sat } _$, but the next theorem shows that the intuitions behind the definitions are simple when $\models_{V_1-} \text{Sat } _$ is well-behaved. For example, Theorem 5.16.3 says that $\llbracket u \rrbracket$ is the set of all primitive terms that refine u .

Theorem 5.16. If $\models_{V_1-} \text{Sat } _$ is well-behaved, then

$$5.16.1. \forall t \in T \cdot \llbracket t \rrbracket = \text{Terms}(\text{Specs}(\{t\}))$$

$$5.16.2. \forall t \in T, S \in \mathbb{S} \cdot \models_{V_1} t \text{ Sat } S \Leftrightarrow \models_{V_2} t \text{ Sat } S$$

$$5.16.3. \forall t \in T, u \in U \cdot t \in \llbracket u \rrbracket \Leftrightarrow \models u \sqsubseteq t$$

Proof.

5.16.1. By induction on the structure of t . Suppose $t = C(t_1, \dots, t_n)$ and assume $\llbracket t_i \rrbracket = \text{Terms}(\text{Specs}(\{t_i\}))$ for $i \in 1..n$. Then:
 $\llbracket t \rrbracket = \text{Terms}(\text{Specs}(C(\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket)))$
 $= \{\text{induction hypothesis}\}$
 $\text{Terms}(\text{Specs}(C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))))$
 $= \{\text{Lemma 5.15}\}$
 $\text{Terms}(\text{Specs}(\{C(t_1, \dots, t_n)\})) = \text{Terms}(\text{Specs}(\{t\})).$

5.16.2. $\models_{V_1} t \text{ Sat } S$
 $\Leftrightarrow \{\text{definition of Specs}\}$
 $S \in \text{Specs}(\{t\})$
 $\Leftrightarrow \{\text{Corollary B.1.7 in Appendix B}\}$
 $S \in \text{Specs}(\text{Terms}(\text{Specs}(\{t\})))$
 $\Leftrightarrow \{\text{Theorem 5.16.1}\}$
 $S \in \text{Specs}(\llbracket t \rrbracket)$
 $\Leftrightarrow \{\text{Lemma 5.7}\}$
 $\models_{V_2} t \text{ Sat } S$

5.16.3. The $\Leftarrow$ proof is trivial since $t \in \llbracket t \rrbracket$. For $\Rightarrow$, assume $\{t\} \subseteq \llbracket u \rrbracket$. Then $\text{Terms}(\text{Specs}(\{t\})) \subseteq \text{Terms}(\text{Specs}(\llbracket u \rrbracket))$ by Corollary B.1.5 in Appendix B, and $\llbracket t \rrbracket \subseteq \llbracket u \rrbracket$ by Theorem 5.16.1 and Lemma 5.4. $\square$

Whether $\models_{V_1-} \text{Sat } _$ is well-behaved depends partly on the expressivity of specifications. For example, suppose

$$\mathbb{K} = \{x := e \mid e \text{ is an arithmetic expression}\} \cup \{-\circ-\}$$

i.e. there is a nullary constructor $x := e$ for all arithmetic expressions e , and a binary constructor for sequential composition. Suppose $\mathbb{S} = \{\text{Even}_x\}$, and $\models_{V_1} t \text{ Sat Even}_x$ holds iff, if t is executed in any state where x is even, then x is even in every resulting state. So $\models_{V_1} x := x + 1 \ ; \ x := x + 1 \text{ Sat Even}_x$, but it is not the case that $\models_{V_1} x := x + 1 \text{ Sat Even}_x$. In fact, $x := x + 1$ does not satisfy any specification. This implies that $\text{Terms}(\text{Specs}(\{x := x + 1\})) = T$, so $x := 1 \ ; \ x := 1 \in \text{Terms}(\text{Specs}(\{x := x + 1\})) \ ; \ \text{Terms}(\text{Specs}(\{x := x + 1\}))$. But $\models_{V_1} x := 1 \ ; \ x := 1 \text{ Sat Even}_x$ does not hold, hence $\models_{V_1} \text{ Sat } _$ is not well-behaved.

Even though $\models_{V_1} \text{ Sat } _$ is not well-behaved, it is still possible to have inference rules that are amenable to freefinement, for example:

$$\begin{array}{l} 1 \frac{}{x := e \text{ Sat Even}_x} \\ \text{provided } e \in \{\dots, -2, 0, 2, \dots\}. \end{array} \quad 2 \frac{t \text{ Sat Even}_x \quad t' \text{ Sat Even}_x}{t \ ; \ t' \text{ Sat Even}_x}$$

If $\mathbb{S}$ is instead a set of specifications of the form $[P, Q]$, where P is a precondition and Q a postcondition, and

$$\models_{V_1} t \ ; \ t' \text{ Sat } [P, Q] \Leftrightarrow \exists R. \models_{V_1} t \text{ Sat } [P, R] \wedge \models_{V_1} t' \text{ Sat } [R, Q]$$

then it is easy to show that this $\models_{V_1} \text{ Sat } _$ is well-behaved.

The completeness of V_1 is a sufficient condition for the well-behavedness of $\models_{V_1} \text{ Sat } _$:

Theorem 5.17. If V_1 is complete, then $\models_{V_1} \text{ Sat } _$ is well-behaved.

Proof. If V_1 is complete, then $\models_{V_1} C(t_1, \dots, t_n) \text{ Sat } S \Leftrightarrow \vdash_{V_1} C(t_1, \dots, t_n) \text{ Sat } S$. The well-behavedness of $\models_{V_1} \text{ Sat } _$ follows by induction on the derivation of $C(t_1, \dots, t_n) \text{ Sat } S$:

- For each rule of the form A_1 , assume $\text{Pred}(C, S_1, \dots, S_n, S)$ and $S_i \in \text{Specs}(\{t_i\})$ for all $i \in 1..n$. So $\forall t'_i \in \text{Terms}(\text{Specs}(\{t_i\})). \models_{V_1} t'_i \text{ Sat } S_i$ for all $i \in 1..n$. The rule is sound with respect to Definition 5.1.1, hence $\forall t \in C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))). \models_{V_1} t \text{ Sat } S$.
- For each rule of the form B_1 , assume $\text{Pred}(S_1, \dots, S_m, S)$ and $\forall t \in C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))). \models_{V_1} t \text{ Sat } S_i$ for all $i \in 1..m$. The rule is sound w.r.t. Definition 5.1.2, so $\forall t \in C(\text{Terms}(\text{Specs}(\{t_1\})), \dots, \text{Terms}(\text{Specs}(\{t_n\}))). \models_{V_1} t \text{ Sat } S$. $\square$

There are several ways to establish harmony between V_2 and the refinement calculi. If a refinement calculus R includes the rule TRANS and enjoys the derivability property $\vdash_{V_2} u \text{ Sat } S \Leftrightarrow \vdash_R S \sqsubseteq u$, then it is trivially harmonic. Including TRANS in R_1 and maintaining the derivability property

at each step also yields a nice presentation: it requires more work to show that R_1 is harmonic, but the treatment of R_4 becomes simpler. This simplification also applies to the existing presentation: R_3 includes TRANS; it is harmonic and satisfies the derivability property by Theorem 5.9. Adding C- i and JOIN- i does not break the derivability property, and hence preserves harmony.

5.2 Applications

5.2.1 Lambda Calculus

The top left corner of Figure 5.1 contains a type system λ_1 for the lambda calculus. By considering pairs of the form (typing context, type) as specifications, it is possible to apply freefinement and obtain a refinement calculus for (extended) lambda terms in the spirit of Denney [20]. The inputs to freefinement are as follows:

1. $\mathbb{K} = \text{Var} \cup \{\lambda x. _ \mid x \in \text{Var}\} \cup \{_ \rightarrow _ \}$
Note that $\mathbb{K}$ defines the language T of lambda terms:

$$e ::= x \mid \lambda x. e \mid e \ e'$$

Here and in the following, x ranges over the set of variables Var , and e ranges over T .

2. $\mathbb{S} = \{[\Gamma; \tau] \mid \Gamma \in \text{Context} \wedge \tau \in \text{Type}\}$, where Context is the set of typing contexts and Type is the set of types that contains the type constructor $_ \rightarrow _$. The intended representation of a typing context Γ is a list of variable names paired with types. Variables may appear more than once in Γ , and variable lookup uses the rightmost occurrence. In the following, σ and τ range over Type , and Γ ranges over Context .
3. $\models_{V_1} \text{Sat } _$ is defined by:
 - $\models_{V_1} x \text{ Sat } [\Gamma; \tau] \Leftrightarrow x : \tau \in \Gamma$
 - $\models_{V_1} \lambda x. e \text{ Sat } [\Gamma; \tau] \Leftrightarrow \exists \sigma, \tau' \cdot \tau = \sigma \rightarrow \tau' \wedge \models_{V_1} e \text{ Sat } [\Gamma, x : \sigma; \tau']$
 - $\models_{V_1} e \ e' \text{ Sat } [\Gamma; \tau] \Leftrightarrow \exists \sigma \cdot \models_{V_1} e \text{ Sat } [\Gamma; \sigma \rightarrow \tau] \wedge \models_{V_1} e' \text{ Sat } [\Gamma; \sigma]$
4. V_1 , shown in the top right corner of Figure 5.1, is obtained from λ_1 by replacing $\Gamma \vdash e : \tau$ with $e \text{ Sat } [\Gamma; \tau]$. The rules VAR, ABS and APP are all of the form A_1 with $n = 0, 1$ and 2 respectively. For example, in the case of ABS, $\text{Pred}(C, S_1, S)$ is defined as
 $\exists x, \Gamma, \sigma, \tau \cdot C = \lambda x. _ \wedge S_1 = [\Gamma, x : \sigma; \tau] \wedge S = [\Gamma; \sigma \rightarrow \tau]$.

λ_1	V_1
$\text{VAR} \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau}$ $\text{ABS} \frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash \lambda x. e : \sigma \rightarrow \tau}$ $\text{APP} \frac{\Gamma \vdash e : \sigma \rightarrow \tau \quad \Gamma \vdash e' : \sigma}{\Gamma \vdash e e' : \tau}$	$\text{VAR} \frac{}{x \text{ Sat } [\Gamma; \tau]}$ provided $x : \tau \in \Gamma$. $\text{ABS} \frac{e \text{ Sat } [\Gamma, x : \sigma; \tau]}{\lambda x. e \text{ Sat } [\Gamma; \sigma \rightarrow \tau]}$ $\text{APP} \frac{e \text{ Sat } [\Gamma; \sigma \rightarrow \tau] \quad e' \text{ Sat } [\Gamma; \sigma]}{e e' \text{ Sat } [\Gamma; \tau]}$
λ_2	R_5
$\text{VAR} \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau}$ $\text{ABS} \frac{\Gamma, x : \sigma \vdash f : \tau}{\Gamma \vdash \lambda x. f : \sigma \rightarrow \tau}$ $\text{APP} \frac{\Gamma \vdash f : \sigma \rightarrow \tau \quad \Gamma \vdash f' : \sigma}{\Gamma \vdash f f' : \tau}$ $\text{SPEC} \frac{}{\Gamma \vdash [\Gamma; \sigma] : \sigma}$	$\text{VAR} \frac{}{[\Gamma; \tau] \sqsubseteq x}$ provided $x : \tau \in \Gamma$. $\text{ABS} \frac{}{[\Gamma; \sigma \rightarrow \tau] \sqsubseteq \lambda x. [\Gamma, x : \sigma; \tau]}$ $\text{APP} \frac{}{[\Gamma; \tau] \sqsubseteq [\Gamma; \sigma \rightarrow \tau] \quad [\Gamma; \sigma]}$ $\text{SPEC} \frac{}{[\Gamma; \sigma] \sqsubseteq [\Gamma; \sigma]}$ $\text{TRANS} \frac{f_1 \sqsubseteq f_2 \quad f_2 \sqsubseteq f_3}{f_1 \sqsubseteq f_3}$ $\text{MONO} \frac{f \sqsubseteq f'}{g[f] \sqsubseteq g[f']}$

Figure 5.1: Freefinement and a typed lambda calculus.

Since V_1 does not contain rules of the form B_1 where $m > 1$, freefinement does not add join terms to the lambda calculus. The system λ_2 in Figure 5.1 is V_2 where $f \text{ Sat } [\Gamma; \tau]$ is written instead as $\Gamma \vdash f : \tau$. The system R_5 , shown in the bottom right of Figure 5.1, is the final harmonic refinement calculus that freefinement produces.

Here is an example top-down typing derivation with R_5 :

$$\begin{aligned}
& [\Gamma; (\sigma \rightarrow \tau) \rightarrow (\sigma \rightarrow \tau)] \\
\sqsubseteq & \quad \text{“ABS”} \\
& \lambda x. [\Gamma, x : \sigma \rightarrow \tau; \sigma \rightarrow \tau] \\
\sqsubseteq & \quad \text{“MONO with ABS”} \\
& \lambda x. \lambda y. [\Gamma, x : \sigma \rightarrow \tau, y : \sigma; \tau] \\
\sqsubseteq & \quad \text{“MONO with APP”} \\
& \lambda x. \lambda y. ([\Gamma, x : \sigma \rightarrow \tau, y : \sigma; \sigma \rightarrow \tau] \ [\Gamma, x : \sigma \rightarrow \tau, y : \sigma; \sigma]) \\
\sqsubseteq & \quad \text{“Twice MONO with VAR”} \\
& \lambda x. \lambda y. (x \ y)
\end{aligned}$$

Since R_5 is harmonic and V_2 is a sound and conservative extension of V_1 , it holds that $\vdash_{\lambda_1} \Gamma \vdash \lambda x. \lambda y. (x \ y) : (\sigma \rightarrow \tau) \rightarrow (\sigma \rightarrow \tau)$.

One might wish to extend R_5 using knowledge particular to lambda calculus typing. It is simple to show that V_1 is complete, so

$$\vdash_{\lambda_1} \Gamma \vdash e : \tau \Leftrightarrow \vdash_{V_1} e \text{ Sat } [\Gamma; \tau] \Leftrightarrow \models_{V_1} e \text{ Sat } [\Gamma; \tau]$$

Furthermore, by Theorems 5.17 and 5.16.2,

$$\models_{V_1} e \text{ Sat } [\Gamma; \tau] \Leftrightarrow \models_{V_2} e \text{ Sat } [\Gamma; \tau]$$

and because V_2 is a sound and conservative extension of V_1 ,

$$\vdash_{V_1} e \text{ Sat } [\Gamma; \tau] \Leftrightarrow \vdash_{V_2} e \text{ Sat } [\Gamma; \tau]$$

Consider the property of *preservation*:

Definition 5.9. A relation $\rightsquigarrow \subseteq T \times T$ satisfies preservation $\stackrel{\text{def}}{=}$ $\forall \Gamma, \tau, e, e' \cdot \text{if } \vdash_{\lambda_1} \Gamma \vdash e : \tau \text{ and } e \rightsquigarrow e', \text{ then } \vdash_{\lambda_1} \Gamma \vdash e' : \tau$.

Theorem 5.18. If $\rightsquigarrow$ satisfies preservation, then:

5.18.1. If $e \rightsquigarrow e'$, then $\models e \sqsubseteq e'$.

5.18.2. If $\vdash_{V_2} e \text{ Sat } [\Gamma; \tau]$ and $e \rightsquigarrow e'$, then $\vdash_{V_2} e' \text{ Sat } [\Gamma; \tau]$.

Proof. The proof of 5.18.2 is trivial. For 5.18.1:

$$\begin{aligned}
& \forall \Gamma, \tau, e, e' \cdot \vdash_{\lambda_1} \Gamma \vdash e : \tau \wedge e \rightsquigarrow e' \Rightarrow \vdash_{\lambda_1} \Gamma \vdash e' : \tau \\
\Leftrightarrow & \{\text{predicate logic}\} \\
& \forall e, e' \cdot e \rightsquigarrow e' \Rightarrow (\forall \Gamma, \tau \cdot \vdash_{\lambda_1} \Gamma \vdash e : \tau \Rightarrow \vdash_{\lambda_1} \Gamma \vdash e' : \tau) \\
\Leftrightarrow & \{\text{equivalent terms}\} \\
& \forall e, e' \cdot e \rightsquigarrow e' \Rightarrow (\forall S \in \mathbb{S} \cdot \models_{v_2} e \text{ Sat } S \Rightarrow \models_{v_2} e' \text{ Sat } S) \\
\Leftrightarrow & \{\text{Lemma 5.8}\} \\
& \forall e, e' \cdot e \rightsquigarrow e' \Rightarrow \models e \sqsubseteq e' \quad \square
\end{aligned}$$

So any relation that satisfies preservation contains only sound refinements that satisfy Harmony 1, and can augment R_5 to yield a sound and harmonic refinement system. Examples of relations that satisfy preservation include:

- The α -conversion relation.
- The β -reduction relation.
- The η -contraction relation. So $\lambda x. (e \ x) \sqsubseteq e$, provided x does not appear free in e .
- The relation $\leq$ on closed terms, where $e \leq e'$ exactly when e has fewer types than e' .

Here is a small example that uses the η -contraction extension:

$$\begin{aligned}
& \lambda x. \lambda y. \lambda z. ((x \ y) \ z) \\
\sqsubseteq & \{\text{MONO with } \eta\text{-contraction}\} \\
& \lambda x. \lambda y. (x \ y) \\
\sqsubseteq & \{\text{MONO with } \eta\text{-contraction}\} \\
& \lambda x. x
\end{aligned}$$

5.2.2 Hoare Logic

Figure 5.2 contains system H , a Hoare logic for simple imperative programs. P is a precondition, Q a postcondition, and c a command in the Hoare triple $\{P\}c\{Q\}$, and $\models_H \{P\}c\{Q\}$ is the usual partial correctness interpretation of $\{P\}c\{Q\}$. By interpreting a specification as a pre-post pair, the rules of H do not fit the rule forms A_1 and B_1 , since the proviso of $AUXVARELIM$ inspects the command c to determine the variables that it writes and reads. However, if specifications also keep track of written and read variables, then it becomes possible to apply freefinement to obtain a refinement calculus in the spirit of Morgan [50]. Here are the inputs:

1. There are constructors for assignments, sequential composition, conditionals and loops:

H	
AUXVARELIM	$\frac{\{P\}c\{Q\}}{\{\exists v \cdot P\}c\{\exists v \cdot Q\}}$ provided $v \notin \text{writes}(c) \cup \text{reads}(c)$.
CONSEQUENCE	$\frac{\{P'\}c\{Q'\}}{\{P\}c\{Q\}}$ provided $P \Rightarrow P'$ and $Q' \Rightarrow Q$.
CONSTANCY	$\frac{\{P\}c\{Q\}}{\{P \wedge R\}c\{Q \wedge R\}}$ provided $FV(R) \cap \text{writes}(c) = \emptyset$.
DISJ	$\frac{\{P\}c\{Q\} \quad \{P'\}c\{Q'\}}{\{P \vee P'\}c\{Q \vee Q'\}}$
VARASSIGN	$\frac{}{\{P[e/x]\}x := e\{P\}}$
SEQCOMP	$\frac{\{P\}c\{Q\} \quad \{Q\}c'\{R\}}{\{P\}c;c'\{R\}}$
COND	$\frac{\{P \wedge b\}c\{Q\} \quad \{P \wedge \neg b\}c'\{Q\}}{\{P\}\mathbf{if} \ b \ \mathbf{then} \ c \ \mathbf{else} \ c'\{Q\}}$
LOOP	$\frac{\{I \wedge b\}c\{I\}}{\{I\}\mathbf{while} \ b \ \mathbf{do} \ c\{I \wedge \neg b\}}$

Figure 5.2: Freefinement and Hoare logic: the system H.

$$\begin{array}{c}
V_1 \\
\\
\text{AUXVARELIM} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}{c \text{ Sat } \bar{x}; \bar{y}: \{\exists v \cdot P, \exists v \cdot Q\}} \\
\text{provided } v \notin \bar{x} \cup \bar{y}. \\
\\
\text{CONSEQUENCE} \frac{c \text{ Sat } \bar{x}'; \bar{y}': \{P', Q'\}}{c \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}} \\
\text{provided } \bar{x} \supseteq \bar{x}' \text{ and } \bar{y} \supseteq \bar{y}' \text{ and } P \Rightarrow P' \text{ and } Q' \Rightarrow Q. \\
\\
\text{CONSTANCY} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}{c \text{ Sat } \bar{x}; \bar{y}: \{P \wedge R, Q \wedge R\}} \\
\text{provided } FV(R) \cap \bar{x} = \emptyset. \\
\\
\text{DISJ} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{P, Q\} \quad c \text{ Sat } \bar{x}; \bar{y}: \{P', Q'\}}{c \text{ Sat } \bar{x}; \bar{y}: \{P \vee P', Q \vee Q'\}} \\
\\
\text{VARASSIGN} \frac{}{x := e \text{ Sat } x; FV(e): \{P[e/x], P\}} \\
\\
\text{SEQCOMP} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{P, Q\} \quad c' \text{ Sat } \bar{x}'; \bar{y}': \{Q, R\}}{c \circ c' \text{ Sat } \bar{x} \cup \bar{x}'; \bar{y} \cup \bar{y}': \{P, R\}} \\
\\
\text{COND} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{P \wedge b, Q\} \quad c' \text{ Sat } \bar{x}'; \bar{y}': \{P \wedge \neg b, Q\}}{\text{if } b \text{ then } c \text{ else } c' \text{ Sat } \bar{x} \cup \bar{x}'; \bar{y} \cup \bar{y}' \cup FV(b): \{P, Q\}} \\
\\
\text{LOOP} \frac{c \text{ Sat } \bar{x}; \bar{y}: \{I \wedge b, I\}}{\text{while } b \text{ do } c \text{ Sat } \bar{x}; \bar{y} \cup FV(b): \{I, I \wedge \neg b\}}
\end{array}$$

Figure 5.3: Freefinement and Hoare logic: the system V_1 .

V_2	
$\text{AUXVARELIM} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}{s \text{ Sat } \bar{x}; \bar{y}: \{\exists v \cdot P, \exists v \cdot Q\}}$ provided $v \notin \bar{x} \cup \bar{y}$.	
$\text{CONSEQUENCE} \frac{s \text{ Sat } \bar{x}'; \bar{y}': \{P', Q'\}}{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}$ provided $\bar{x} \supseteq \bar{x}'$ and $\bar{y} \supseteq \bar{y}'$ and $P \Rightarrow P'$ and $Q' \Rightarrow Q$.	
$\text{CONSTANCY} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}{s \text{ Sat } \bar{x}; \bar{y}: \{P \wedge R, Q \wedge R\}}$ provided $FV(R) \cap \bar{x} = \emptyset$.	
$\text{DISJ} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\} \quad s \text{ Sat } \bar{x}; \bar{y}: \{P', Q'\}}{s \text{ Sat } \bar{x}; \bar{y}: \{P \vee P', Q \vee Q'\}}$	
$\text{VARASSIGN} \frac{}{x := e \text{ Sat } x; FV(e): \{P[e/x], P\}}$	
$\text{SEQCOMP} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\} \quad s' \text{ Sat } \bar{x}'; \bar{y}': \{Q, R\}}{s \circ s' \text{ Sat } \bar{x} \cup \bar{x}'; \bar{y} \cup \bar{y}': \{P, R\}}$	
$\text{COND} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P \wedge b, Q\} \quad s' \text{ Sat } \bar{x}'; \bar{y}': \{P \wedge \neg b, Q\}}{\text{if } b \text{ then } s \text{ else } s' \text{ Sat } \bar{x} \cup \bar{x}'; \bar{y} \cup \bar{y}' \cup FV(b): \{P, Q\}}$	
$\text{LOOP} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{I \wedge b, I\}}{\text{while } b \text{ do } s \text{ Sat } \bar{x}; \bar{y} \cup FV(b): \{I, I \wedge \neg b\}}$	
$\text{SPEC} \frac{}{\bar{x}; \bar{y}: \{P, Q\} \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}$	
$\text{JOIN} \frac{s \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}{\sqcup(\dots, s, \dots) \text{ Sat } \bar{x}; \bar{y}: \{P, Q\}}$	

Figure 5.4: Freefinement and Hoare logic: the system V_2 .

R_6	
AUXVARELIM	$\frac{}{\bar{x}; \bar{y}: \{\exists v \cdot P, \exists v \cdot Q\} \sqsubseteq \bar{x}; \bar{y}: \{P, Q\}}$ provided $v \notin \bar{x} \cup \bar{y}$.
CONSEQUENCE	$\frac{}{\bar{x}; \bar{y}: \{P, Q\} \sqsubseteq \bar{x}'; \bar{y}': \{P', Q'\}}$ provided $\bar{x} \supseteq \bar{x}'$ and $\bar{y} \supseteq \bar{y}'$ and $P \Rightarrow P'$ and $Q \Rightarrow Q'$.
CONSTANCY	$\frac{}{\bar{x}; \bar{y}: \{P \wedge R, Q \wedge R\} \sqsubseteq \bar{x}; \bar{y}: \{P, Q\}}$ provided $FV(R) \cap \bar{x} = \emptyset$.
DISJ	$\frac{}{\bar{x}; \bar{y}: \{P \vee P', Q \vee Q'\} \sqsubseteq \bigsqcup(\bar{x}; \bar{y}: \{P, Q\}, \bar{x}; \bar{y}: \{P', Q'\})}$
VARASSIGN	$\frac{}{x; FV(e): \{P[e/x], P\} \sqsubseteq x := e}$
SEQCOMP	$\frac{}{\bar{x} \cup \bar{x}'; \bar{x} \cup \bar{x}': \{P, R\} \sqsubseteq \bar{x}; \bar{y}: \{P, Q\} \ ; \ \bar{x}'; \bar{y}': \{Q, R\}}$
COND	$\frac{}{\bar{x} \cup \bar{x}'; \bar{y} \cup \bar{y}' \cup FV(b): \{P, Q\} \sqsubseteq \text{if } b \text{ then } \bar{x}; \bar{y}: \{P \wedge b, Q\} \text{ else } \bar{x}'; \bar{y}': \{P \wedge \neg b, Q\}}$
LOOP	$\frac{}{\bar{x}; \bar{y} \cup FV(b): \{I, I \wedge \neg b\} \sqsubseteq \text{while } b \text{ do } \bar{x}; \bar{y}: \{I \wedge b, I\}}$
SPEC	$\frac{}{\bar{x}; \bar{y}: \{P, Q\} \sqsubseteq \bar{x}; \bar{y}: \{P, Q\}}$
JOIN	$\frac{}{\bar{x}; \bar{y}: \{P, Q\} \sqsubseteq \bigsqcup(\dots, \bar{x}; \bar{y}: \{P, Q\}, \dots)}$
TRANS	$\frac{s_1 \sqsubseteq s_2 \quad s_2 \sqsubseteq s_3}{s_1 \sqsubseteq s_3}$
MONO	$\frac{s \sqsubseteq s'}{t[s] \sqsubseteq t[s']}$
UNJOIN	$\frac{}{\bigsqcup(s, \dots, s) \sqsubseteq s}$

Figure 5.5: Freefinement and Hoare logic: the system R_6 .

$$\begin{aligned} \mathbb{K} = & \{x := e \mid x \in \text{Var} \wedge e \in \text{IntExp}\} \\ & \cup \{-; -\} \\ & \cup \{\mathbf{if} \ b \ \mathbf{then} \ - \ \mathbf{else} \ - \mid b \in \text{BoolExp}\} \\ & \cup \{\mathbf{while} \ b \ \mathbf{do} \ - \mid b \in \text{BoolExp}\} \end{aligned}$$

2. A specification consists of two sets of variables and two assertions, written in a notation resembling Morgan's specification statement [49]:

$$\mathbb{S} = \{\bar{x}; \bar{y} : \{P, Q\} \mid \bar{x}, \bar{y} \in \mathcal{P}(\text{Var}) \wedge P, Q \in \text{Assertion}\}$$

3. In the specification $\bar{x}; \bar{y} : \{P, Q\}$, the $\bar{x}$ and $\bar{y}$ are upper bounds on the sets of variables written and read by the command respectively, the P is a precondition and the Q a postcondition:

$$\models_{V_1} c \text{ Sat } \bar{x}; \bar{y} : \{P, Q\} \stackrel{\text{def}}{=} \text{writes}(c) \subseteq \bar{x} \wedge \text{reads}(c) \subseteq \bar{y} \wedge \models_H \{P\}c\{Q\}$$

4. V_1 , shown in Figure 5.3, has the following relationship with H :

$$\models_{V_1} c \text{ Sat } \bar{x}; \bar{y} : \{P, Q\} \Leftrightarrow \text{writes}(c) \subseteq \bar{x} \wedge \text{reads}(c) \subseteq \bar{y} \wedge \models_H \{P\}c\{Q\}$$

Note that:

- The non-structural rules of H have counterparts in V_1 that embody the definitions of *writes* and *reads*. For example, the conclusion of COND reflects that $\text{writes}(\mathbf{if} \ b \ \mathbf{then} \ c \ \mathbf{else} \ c') \stackrel{\text{def}}{=} \text{writes}(c) \cup \text{writes}(c')$ and $\text{reads}(\mathbf{if} \ b \ \mathbf{then} \ c \ \mathbf{else} \ c') \stackrel{\text{def}}{=} \text{reads}(c) \cup \text{reads}(c') \cup FV(b)$.
- The structural rules of H that inspect c for its write and/or read sets have counterparts in V_1 that consult the specification instead. See for example the proviso of AUXVARELIM.
- CONSEQUENCE in V_1 allows the enlargement of write and read sets. This loosening of the bounds is useful in refinement developments, because then the resulting code is not forced to write and read all the variables that were originally available for writing and reading.

The V_1 -counterparts of the structural rules of H are all of the form B_1 . For example, $m = 1$ in the case of CONSTANCY, and $m = 2$ for DISJ. The other rules are of the form A_1 . For example, $n = 2$ in the case of COND, and $n = 1$ for LOOP.

The systems V_2 and R_6 that free refinement produces appear in Figures 5.4 and 5.5. R_6 yields several derived rules that may be useful in practical refinement developments. For example, the rule:

$$\text{DERIVEDVARASSIGN} \frac{}{\bar{x}; \bar{y}: \{P, Q\} \sqsubseteq z := e}$$

provided $z \in \bar{x}$ and $FV(e) \subseteq \bar{y}$ and $P \Rightarrow Q[e/z]$.

can replace VARASSIGN, and is similar to the assignment law of Morgan (Law 1.3 on p. 8 of [50]). Likewise, the derived rule:

$$\text{FOLLOWINGVARASSIGN} \frac{}{\bar{x}; \bar{y}: \{P, Q\} \sqsubseteq \bar{x}; \bar{y}: \{P, Q[e/z]\} \wp z := e}$$

provided $z \in \bar{x}$ and $FV(e) \subseteq \bar{y}$.

is similar to the *following assignment* law of Morgan (Law 3.5 on p. 32 of [50]).

Here is an example showing that R_6 can derive a correct factorial algorithm starting with its specification:

$$\begin{aligned} & y, z; x, y, z: \{\text{true}, y = x!\} \\ \sqsubseteq & \quad \text{"SEQCOMP"} \\ & y, z; \emptyset: \{\text{true}, y = 1 \wedge z = 0\} \wp y, z; x, y, z: \{y = 1 \wedge z = 0, y = x!\} \end{aligned}$$

The first spec statement is refined as follows:

$$\begin{aligned} & y, z; \emptyset: \{\text{true}, y = 1 \wedge z = 0\} \\ \sqsubseteq & \quad \text{"SEQCOMP"} \\ & y; \emptyset: \{\text{true}, y = 1\} \wp z; \emptyset: \{y = 1, y = 1 \wedge z = 0\} \\ \sqsubseteq & \quad \text{"Twice MONO with CONSEQUENCE"} \\ & y; \emptyset: \{1 = 1, y = 1\} \wp z; \emptyset: \{y = 1 \wedge 0 = 0, y = 1 \wedge z = 0\} \\ \sqsubseteq & \quad \text{"Twice MONO with VARASSIGN"} \\ & y := 1 \wp z := 0 \end{aligned}$$

And for the second spec statement:

$$\begin{aligned} & y, z; x, y, z: \{y = 1 \wedge z = 0, y = x!\} \\ \sqsubseteq & \quad \text{"CONSEQUENCE"} \\ & y, z; x, y, z: \{y = z!, y = z! \wedge \neg z \neq x\} \\ \sqsubseteq & \quad \text{"LOOP"} \\ & \text{while } z \neq x \text{ do } y, z; y, z: \{y = z! \wedge z \neq x, y = z!\} \\ \sqsubseteq & \quad \text{"MONO with SEQCOMP"} \\ & \text{while } z \neq x \text{ do } z; z: \{y = z! \wedge z \neq x, y \cdot z = z!\} \wp y; y, z: \{y \cdot z = z!, y = z!\} \\ \sqsubseteq & \quad \text{"MONO with VARASSIGN"} \\ & \text{while } z \neq x \text{ do } z; z: \{y = z! \wedge z \neq x, y \cdot z = z!\} \wp y := y \cdot z \\ \sqsubseteq & \quad \text{"MONO with CONSEQUENCE"} \\ & \text{while } z \neq x \text{ do } z; z: \{y \cdot (z+1) = (z+1)!, y \cdot z = z!\} \wp y := y \cdot z \end{aligned}$$

$\sqsubseteq$ “MONO with VARASSIGN”
while $z \neq x$ **do** $z := z+1 \ ; y := y \cdot z$

Since $\vdash_{R_6} y, z; x, y, z: \{\text{true}, y = x!\} \sqsubseteq y := 1 \ ; z := 0 \ ; \text{while } z \neq x \text{ do } z := z+1 \ ; y := y \cdot z$, it is the case that $\vdash_{V_1} y := 1 \ ; z := 0 \ ; \text{while } z \neq x \text{ do } z := z+1 \ ; y := y \cdot z \text{ Sat } y, z; x, y, z: \{\text{true}, y = x!\}$ and hence also $\vdash_H \{\text{true}\} y := 1 \ ; z := 0 \ ; \text{while } z \neq x \text{ do } z := z+1 \ ; y := y \cdot z \{y = x!\}$.

Here is another example of using R_6 ; it involves join statements. The statement $\sqcup(\bar{x}; \bar{y}: \{P_1, Q_1\}, \bar{x}; \bar{y}: \{P_2, Q_2\})$ is the join of the specification statements $\bar{x}; \bar{y}: \{P_1, Q_1\}$ and $\bar{x}; \bar{y}: \{P_2, Q_2\}$. Expressing it as a spec statement is simple because

$$\sqcup(\bar{x}; \bar{y}: \{P_1, Q_1\}, \bar{x}; \bar{y}: \{P_2, Q_2\}) \equiv \bar{x}; \bar{y}: \{P_1, Q_1\} \text{ also } \{P_2, Q_2\}$$

where the definition of $\{P_1, Q_1\} \text{ also } \{P_2, Q_2\}$, taken from [58], is: $\{(P_1 \wedge z=1) \vee (P_2 \wedge z \neq 1), (Q_1 \wedge z=1) \vee (Q_2 \wedge z \neq 1)\}$ where z is fresh. R_6 can derive both directions of refinement. Firstly:

$\sqcup(\bar{x}; \bar{y}: \{P_1, Q_1\}, \bar{x}; \bar{y}: \{P_2, Q_2\})$
 $\sqsubseteq$ “Twice MONO with CONSEQUENCE”
 $\sqcup(\bar{x}; \bar{y}: \{\exists z. (P_1 \wedge z=1 \vee P_2 \wedge z \neq 1) \wedge z=1, \exists z. (Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1) \wedge z=1\},$
 $\bar{x}; \bar{y}: \{\exists z. (P_1 \wedge z=1 \vee P_2 \wedge z \neq 1) \wedge z \neq 1, \exists z. (Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1) \wedge z \neq 1\})$
 $\sqsubseteq$ “Twice MONO with AUXVARELIM”
 $\sqcup(\bar{x}; \bar{y}: \{(P_1 \wedge z=1 \vee P_2 \wedge z \neq 1) \wedge z=1, (Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1) \wedge z=1\},$
 $\bar{x}; \bar{y}: \{(P_1 \wedge z=1 \vee P_2 \wedge z \neq 1) \wedge z \neq 1, (Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1) \wedge z \neq 1\})$
 $\sqsubseteq$ “Twice MONO with CONSTANCY”
 $\sqcup(\bar{x}; \bar{y}: \{P_1 \wedge z=1 \vee P_2 \wedge z \neq 1, Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1\},$
 $\bar{x}; \bar{y}: \{P_1 \wedge z=1 \vee P_2 \wedge z \neq 1, Q_1 \wedge z=1 \vee Q_2 \wedge z \neq 1\})$
 $\sqsubseteq$ “UNJOIN”
 $\bar{x}; \bar{y}: \{P_1, Q_1\} \text{ also } \{P_2, Q_2\}$

Secondly:

$\bar{x}; \bar{y}: \{(P_1 \wedge z=1) \vee (P_2 \wedge z \neq 1), (Q_1 \wedge z=1) \vee (Q_2 \wedge z \neq 1)\}$
 $\sqsubseteq$ “DISJ”
 $\sqcup(\bar{x}; \bar{y}: \{P_1 \wedge z=1, Q_1 \wedge z=1\}, \bar{x}; \bar{y}: \{P_2 \wedge z \neq 1, Q_2 \wedge z \neq 1\})$
 $\sqsubseteq$ “Twice MONO with CONSTANCY”
 $\sqcup(\bar{x}; \bar{y}: \{P_1, Q_1\}, \bar{x}; \bar{y}: \{P_2, Q_2\})$

Leino and Manohar [43] mention several uses of the join of spec-like statements.

5.2.3 Discussion

The type system λ_1 considered above is very simple. Freefinement also applies to System F and other more sophisticated type systems.

Although λ_1 had only rules of the form A_1 , typing rules of the form B_1 are quite common – examples include rules for subtyping and intersection types:

$$\text{SUB} \frac{\Gamma \vdash e : \tau}{\Gamma \vdash e : \tau'} \quad \text{provided } \tau <: \tau'. \quad \text{INTER} \frac{\Gamma \vdash e : \tau \quad \Gamma \vdash e : \tau'}{\Gamma \vdash e : \tau \wedge \tau'}$$

There is no golden recipe for adapting proof systems to make them amenable to freefinement. However, enriching specifications and/or terms might help. The Hoare logic example used enriched specifications to keep track of write and read sets. Consider again the two problematic rules from before:

$$2 \frac{\text{succ}(n) : \mathbb{N}}{\text{pred}(\text{succ}(n)) : \mathbb{N}} \quad 3 \frac{n : \mathbb{N}}{\text{pred}(n) : \mathbb{N}} \quad \text{provided positive}(n).$$

Rule 2 can be accommodated by choosing $\mathbb{S} = \{‘z’, ‘s’, ‘p’\} \times \{\mathbb{N}\}$. Intuitively, the specification $(‘s’, \mathbb{N})$ tracks the fact that the outermost constructor is ‘succ’. The rule then becomes:

$$2 \frac{n : (‘s’, \mathbb{N})}{\text{pred}(n) : (‘p’, \mathbb{N})}$$

Rule 3 can be accommodated by choosing $\mathbb{S} = \mathbb{N} \times \{\mathbb{N}\}$. Then the sentence $n : (i, \mathbb{N})$ tracks the fact that term n denotes the natural number i . The adapted rule is of the form A_1 with $n = 1$:

$$3 \frac{n : (i, \mathbb{N})}{\text{pred}(n) : (i - 1, \mathbb{N})} \quad \text{provided } i > 0.$$

In some cases it might be useful to enrich the term language. For example, consider the rule of concurrent separation logic [7] that removes auxiliary commands (ghost assignments):

$$\text{AUXILIARY} \frac{\Gamma \vdash \{P\}c\{Q\}}{\Gamma \vdash \{P\}c\backslash a\{Q\}} \quad \text{provided } a \in \text{aux}(c) \text{ and } a \cap (FV(P) \cup FV(Q)) = \emptyset.$$

This rule is not of the form A_1 or B_1 , because it contains a meta-operation in the conclusion. However, if the meta-operation is turned into an explicit constructor (and specifications track auxiliaries), then the rule is of the form B_1 with $m = 1$ and freefinement can handle it.

5.3 Constructing correct OO programs

This section outlines a technique that provides correctness by construction for OO programs. It applies freefinement to the statement verification part of the separation logic proof system in Chapter 2, and combines the result with rules that deal with non-statement constructs.

The judgements for statement verification in Section 2.5 have the form $\Delta; \Gamma \vdash_s \{P\}s\{Q\}$. No rule modifies Δ and Γ , so they are constant parameters of statement verification. Spec statements can consequently exclude Δ and Γ and retain the form $\bar{x}; \bar{y} : \{P, Q\}$ of the previous section. Freefinement also adds join statements to the language, and proceeds as in the Hoare logic example. The resulting refinement judgements have the form $\Delta; \Gamma \vdash s \sqsubseteq s'$. For example, the refinement counterpart of the Frame rule resembles CONSTANCY from before:

$$\frac{FV(R) \cap \bar{x} = \emptyset}{\Delta; \Gamma \vdash \bar{x}; \bar{y} : \{P * R, Q * R\} \sqsubseteq \bar{x}; \bar{y} : \{P, Q\}}$$

The ability to refine statements is useful in the context of OO programming, but correctness by construction also needs to accommodate other language constructs and concepts such as data refinement [49]. To this end, the rest of the section presents a suitably extended language and a development calculus.

5.3.1 Language and proof rules

Figure 5.6 shows the grammar of a slightly modified language. The highlighted productions generate non-executable code, which facilitate the gradual development of methods and statements. The only other difference with the language of Chapter 2 is that method declarations of the form ‘Sig Sd Ss **do** s **end**’ do not explicitly indicate whether they introduce or override the method. However, the earlier convention that all non-constructor methods must be listed in all subclasses makes it simple to distinguish these cases.

The proof rules of Chapter 2 are sufficient to verify code that does not use the highlighted productions. Verifying programs that are partially developed needs additional rules for the new constructs. Familiar rules apply to spec and join statements:

$$\frac{}{\Delta; \Gamma \vdash_s \{P\}\bar{x}; \bar{y} : \{P, Q\}\{Q\}}$$

$$\frac{\Delta; \Gamma \vdash_s \{P\}s\{Q\}}{\Delta; \Gamma \vdash_s \{P\}\sqcup(\dots, s, \dots)\{Q\}}$$

$L ::= \text{class } G \text{ Inh Def Feat } \text{end}$	
$\text{Inh} ::= \text{inherit } H$	
$\text{Def} ::= \text{define } \overline{D}$	
$\text{Feat} ::= \text{feature } \overline{M} \overline{F}$	
$D ::= x.\text{p}_G(\overline{t}:\overline{y}) \text{ as } P$	<i>Define clause</i>
$M ::= \text{Sig}$	<i>Method declaration</i>
Sig Sd	
Sig Sd Ss	
$\text{Sig Sd Ss do } s \text{ end}$	
inherit Sig Sd Ss	
$F ::= f: \text{Type}$	<i>Field declaration</i>
$\text{Sig} ::= m(\text{Args}) \text{ Rt}$	
$\text{Sd} ::= \text{dynamic Spec}$	<i>Dynamic specification</i>
$\text{Ss} ::= \text{static Spec}$	<i>Static specification</i>
$\text{Spec} ::= \{P\}\text{-}\{Q\} \mid \{P\}\text{-}\{Q\} \text{ also Spec}$	<i>Specification</i>
$s ::= \overline{x}; \overline{y}:\{P, Q\}$	<i>Spec statement</i>
$\sqcup(s_1, \dots, s_n)$	<i>Join statement</i>
skip	
$x := e$	<i>Assignment</i>
$s ; s$	<i>Sequential composition</i>
if e then s else s end	<i>Conditional</i>
while e do s end	<i>Loop</i>
$x: \text{Type. } s$	<i>Local variable block</i>
$x := \text{new } G$	<i>Object allocation</i>
$x := y.f$	<i>Field lookup</i>
$x.f := e$	<i>Field update</i>
$x := y.m(\overline{e}) \mid y.m(\overline{e})$	<i>Dynamically dispatched call</i>
$x := y.G::m(\overline{e}) \mid y.G::m(\overline{e})$	<i>Direct method call</i>
$e ::= x \mid e + e \mid e = e \mid \text{Void} \mid 0 \mid 1 \mid 2 \mid \dots$	<i>Expression</i>
$\text{Type} ::= \text{INT} \mid \text{BOOL} \mid G$	
$\text{Args} ::= \overline{x: \text{Type}}$	<i>Formal arguments</i>
$\text{Rt} ::= \epsilon \mid : \text{Type}$	<i>Return type</i>

Figure 5.6: The extended grammar.

A method signature is always verified:

$$\frac{}{\Delta; \Gamma \vdash_m \text{Sig in G parent H}}$$

In order to reduce duplication and improve readability, the presentation will frequently use holes or placeholders in the premises of rules. These holes can then be filled in different ways to obtain different rules. For example, the rule for verifying a method signature with a dynamic specification has a hole, $[\text{Hole}_{\text{B.s.}}]$, that captures the Behavioral subtyping proof obligation:

$$\frac{[\text{Hole}_{\text{B.s.}}]}{\Delta; \Gamma \vdash_m \text{Sig Sd in G parent H}}$$

$[\text{Hole}_{\text{B.s.}}]$ can be filled in two ways:

$$\text{H.m} \notin \text{dom}(\Gamma)$$

$$\begin{aligned} \text{Sd} &= \mathbf{dynamic} \text{ Spec}_G \\ \Gamma(\text{H.m}) &= (\bar{u}, \text{Spec}_H) \\ \Delta \vdash \text{Spec}_G &\implies \text{Spec}_H \end{aligned}$$

The rule for signatures with dynamic and static specifications also uses $[\text{Hole}_{\text{B.s.}}]$, and adds a proof obligation for Dynamic dispatch:

$$\frac{\begin{array}{c} [\text{Hole}_{\text{B.s.}}] \\ [\text{Hole}_{\text{D.d.}}] \end{array}}{\Delta; \Gamma \vdash_m \text{Sig Sd Ss in G parent H}}$$

$[\text{Hole}_{\text{D.d.}}]$ is a placeholder for:

$$\begin{aligned} \text{Ss} &= \mathbf{static} \text{ Spec}_S \\ \text{Sd} &= \mathbf{dynamic} \text{ Spec}_D \\ \Delta \vdash \text{Spec}_S &\xrightarrow{\mathbf{Current} : G} \text{Spec}_D \end{aligned}$$

The following definition is convenient for referring to a collection of well-formed and verified classes:

$$\vdash \bar{L} \stackrel{\text{def}}{=} \forall L \in \bar{L}. \text{apf}(L); \text{specs}(\bar{L}) \vdash_c L$$

5.3.2 Development calculus

The calculus for developing correct OO programs uses judgements of the form $\bar{L} \rightsquigarrow \bar{L}'$. The symbol $\rightsquigarrow$ can be read as ‘evolves to’ or ‘becomes’. To be useful, the calculus must satisfy two properties:

- Soundness: If $\vdash \bar{L}$ and $\bar{L} \rightsquigarrow \bar{L}'$, then $\vdash \bar{L}'$.
- Completeness: If $\vdash \bar{L}$ then $\epsilon \rightsquigarrow^* \bar{L}$.

The soundness property guarantees that the rules will maintain correctness and not introduce errors. The completeness property ensures that any correct program can be developed from scratch with the calculus. In practical terms, this means that the calculus supports the top-down development of every correct program. However, several rules will break with the top-down style and enable more flexible evolution. The rest of this section summarises the calculus.

Adding a class

Any new class needs a fresh name. A new class that does not inherit from any user-defined class inherits no method signatures:

$$\frac{G \notin \text{class-names}(\bar{L}) \cup \{\text{ANY}\}}{\bar{L} \rightsquigarrow \bar{L} + \text{class } G \text{ define feature end}}$$

Otherwise, it inherits from its parent the signatures of all methods that are not constructors:

$$\frac{\begin{array}{l} G \notin \text{class-names}(\bar{L}) \cup \{\text{ANY}\} \\ \bar{L} = \bar{L}_1 + \text{class } H \text{ Inh Def feature } \bar{M} \bar{F} \text{ end} \\ \bar{\text{Sig}} = \text{non-constructor-signatures}(\bar{M}) \end{array}}{\bar{L} \rightsquigarrow \bar{L} + \text{class } G \text{ inherit } H \text{ define feature } \bar{\text{Sig}} \text{ end}}$$

Adding an apf entry

Apf entries support data refinement: they describe how a class represents the abstract concepts/structures expressed by apf predicates. Provided that the definition is well-formed, the rule for adding an apf entry is simple:

$$\frac{\begin{array}{l} \bar{L} = \bar{L}_1 + \text{class } G \text{ Inh define } \bar{D} \text{ Feat end} \\ \textcolor{red}{p}_G \notin \text{defined-apf-entries}(\bar{D}) \\ \bar{D}' = \bar{D} + x.\textcolor{red}{p}_G(\bar{t}:\bar{y}) \text{ as } P \end{array}}{\bar{L} \rightsquigarrow \bar{L}_1 + \text{class } G \text{ Inh define } \bar{D}' \text{ Feat end}}$$

Adding a field

A new field must have a fresh name and a legal type:

$$\begin{array}{l}
 \bar{L} = \bar{L}_1 + \text{class } G \text{ Inh Def feature } \bar{M} \bar{F} \text{ end} \\
 \text{classes} = \text{ancestors}(G, \bar{L}) + G + \text{descendants}(G, \bar{L}) \\
 f \notin \text{field-names}(\text{classes}, \bar{L}) \\
 \text{Type} \in \text{legal-types}(\bar{L}) \\
 \bar{F}' = \bar{F} + f: \text{Type} \\
 \hline
 \bar{L} \rightsquigarrow \bar{L}_1 + \text{class } G \text{ Inh Def feature } \bar{M} \bar{F}' \text{ end}
 \end{array}$$

Adding a constructor signature

A constructor signature has the same name as its enclosing class and no return type.

$$\begin{array}{l}
 \bar{L} = \bar{L}_1 + \text{class } G \text{ Inh Def feature } \bar{M} \bar{F} \text{ end} \\
 \text{Sig} = G(\text{Args}) \\
 \text{valid-signature}(\text{Sig}, G, \bar{L}) \\
 \bar{M}' = \bar{M} + \text{Sig} \\
 \hline
 \bar{L} \rightsquigarrow \bar{L}_1 + \text{class } G \text{ Inh Def feature } \bar{M}' \bar{F} \text{ end}
 \end{array}$$

The predicate $\text{valid-signature}(\text{Sig}, G, \bar{L})$ holds iff

1. $\forall G \in \bar{G}. \text{name}(\text{Sig}) \notin \text{signature-names}(G, \bar{L})$
2. If Sig has a return type, then it must be in $\text{legal-types}(\bar{L})$
3. The argument names in Sig are all distinct
4. The argument types of Sig are all in $\text{legal-types}(\bar{L})$

Adding an ordinary method signature

A new ordinary method signature gets added to a class and all its descendants, which can later decide to inherit or override the eventual implementation.

$$\begin{array}{l}
 \bar{L} = \bar{L}_1 + \bar{L}_2 \\
 \text{class-names}(\bar{L}_2) = G + \text{descendants}(G, \bar{L}) \\
 \text{name}(\text{Sig}) \notin \text{CLASS-NAMESPACE} \\
 \text{valid-signature}(\text{Sig}, \text{class-names}(\bar{L}_2), \bar{L}) \\
 \bar{L}_2 = \biguplus_{i \in 1..n} \text{class } H_i \text{ Inh}_i \text{ Def}_i \text{ feature } \bar{M}_i \bar{F}_i \text{ end} \\
 \bar{L}'_2 = \biguplus_{i \in 1..n} \text{class } H_i \text{ Inh}_i \text{ Def}_i \text{ feature } \bar{M}_i + \text{Sig } \bar{F}_i \text{ end} \\
 \hline
 \bar{L} \rightsquigarrow \bar{L}_1 + \bar{L}'_2
 \end{array}$$

Method manipulation

All remaining rules manipulate methods, and are based on the following rule:

$$\begin{array}{l}
 \overline{L} = \overline{L_1} + L \\
 L = \mathbf{class} \ G \ \mathbf{inherit} \ H \ \mathbf{Def} \ \mathbf{feature} \ \overline{M} \ \overline{F} \ \mathbf{end} \\
 \Gamma = \mathit{specs}(\overline{L}) \\
 \Delta = \mathit{apf}(L) \\
 \overline{M} = \overline{M_1} + M \\
 [\mathit{Hole}_a] \\
 \overline{M'} = \overline{M_1} + M' \\
 \hline
 \overline{L} \rightsquigarrow \overline{L_1} + \mathbf{class} \ G \ \mathbf{inherit} \ H \ \mathbf{Def} \ \mathbf{feature} \ \overline{M'} \ \overline{F} \ \mathbf{end}
 \end{array}$$

Adding a dynamic specification

The rule for adding a dynamic specification to a method signature fills $[\mathit{Hole}_a]$ as follows:

$$\begin{array}{l}
 M = m(\mathit{Args}) \ \mathit{Rt} \\
 [\mathit{Hole}_{B.s.}] \\
 M' = m(\mathit{Args}) \ \mathit{Rt} \ \mathit{Sd}
 \end{array}$$

Adding a static specification

Static specifications are also related to data refinement, since they can state how an operation on abstract structures/concepts, as mentioned in the dynamic specification, is realised in more concrete terms. The rule for adding a static specification fills $[\mathit{Hole}_a]$ with:

$$\begin{array}{l}
 M = \mathit{Sig} \ \mathit{Sd} \\
 [\mathit{Hole}_{D.d.}] \\
 M' = \mathit{Sig} \ \mathit{Sd} \ \mathit{Ss}
 \end{array}$$

Adding a method body

The calculus supports the refinement of method bodies, so the initial body can well be a single spec statement, as this filling for $[\mathit{Hole}_a]$ shows:

$$\begin{array}{l}
 M = \mathit{Sig} \ \mathit{Sd} \ \mathit{Ss} \\
 \mathit{Sig} = m(\mathit{Args}) \ \mathit{Rt} \\
 \mathit{Ss} = \mathbf{static} \ \{S\} \text{--} \{T\} \\
 [\mathit{Hole}_b] \\
 \overline{y} = \overline{x} + \mathit{names}(\mathit{Args}) \\
 M' = \mathit{Sig} \ \mathit{Sd} \ \mathit{Ss} \ \mathbf{do} \ \overline{x}; \overline{y} : \{S, T\} \ \mathbf{end}
 \end{array}$$

The purpose of $[\text{Hole}_b]$ is to constrain $\bar{x}$, the variable(s) that may be modified, appropriately. It has two fillings:

$$\begin{aligned} \text{Rt} &= \epsilon \\ \bar{x} &= \epsilon \\ \text{Rt} &\neq \epsilon \\ \bar{x} &= \mathbf{Result} \end{aligned}$$

Removing a method body

This gives the opportunity to implement the body in a different way, or to inherit a parent version. The filling for $[\text{Hole}_a]$ is simple:

$$\begin{aligned} M &= \text{Sig Sd Ss } \mathbf{do\ s\ end} \\ M' &= \text{Sig Sd Ss} \end{aligned}$$

This is an example of a rule that relaxes the top-down style of development.

Inheriting a method body

Inheriting a method body imposes the Inheritance proof obligation, as the filling for $[\text{Hole}_a]$ shows:

$$\begin{aligned} M &= \text{Sig Sd Ss} \\ \text{Sig} &= m(\text{Args})\ \text{Rt} \\ \text{Ss} &= \mathbf{static}\ \text{Spec}_G \\ \Gamma(H::m) &= (\bar{u}, \text{Spec}_H) \\ \Delta \vdash \text{Spec}_H &\Longrightarrow \text{Spec}_G \\ M' &= \mathbf{inherit}\ \text{Sig Sd Ss} \end{aligned}$$

Uninheriting a method body

The next filling for $[\text{Hole}_a]$ can be used to override a method's implementation instead of inheriting it:

$$\begin{aligned} M &= \mathbf{inherit}\ \text{Sig Sd Ss} \\ M' &= \text{Sig Sd Ss} \end{aligned}$$

Statement refinement

This is where the result of freefinement is incorporated in the development calculus. It fills $[\text{Hole}_a]$ with:

$$\begin{aligned} M &= \text{Sig Sd Ss } \mathbf{do\ s\ end} \\ \Delta; \Gamma \vdash s &\sqsubseteq s' \\ M' &= \text{Sig Sd Ss } \mathbf{do\ s'\ end} \end{aligned}$$

5.4 Related Work

In his work on refinement for the lambda calculus, Denney [20] treats types as rudimentary specifications and introduces a specification construct $?_\tau$ for each type τ . Conceptually, $?_\tau$ corresponds to $[\Gamma; \tau]$ where the context Γ is left implicit. For example, consider the term $\lambda x : \sigma. ?_\tau$ in the context Γ . The $?_\tau$ inside the term corresponds to $[\Gamma, x : \sigma; \tau]$. Denney also considers richer specifications for lambda terms in his PhD thesis [21]. This results in a more powerful refinement calculus in which specification constructs can contain logical assumptions.

The specification statement $\bar{x} : [P, Q]$ of Morgan [49] is analogous to $\bar{x}; \text{Var} : \{P, Q\}$, since there is no restriction on the variables that the statement may read. However, his specification statement is a total correctness specification, and the accompanying refinement calculus [50] establishes total correctness. Similar refinement calculi for total correctness were proposed by Back [1, 2], Morris [51] and Hehner [31]. The books [50, 2, 31] contain many examples of how correct programs can be constructed from specifications via refinement.

Leino and Manohar [43] consider the join of Morgan's specification statements $\bar{x} : [P_1, Q_1]$ and $\bar{x} : [P_2, Q_2]$, and mention several of its uses. Free-refinement adds explicit constructors for joins, and relies on the ability to join arbitrary terms from U in order to establish harmony.

Several systems in the literature provide refinement and correctness by construction for OO programs [4, 48, 6, 5, 55]. The calculus in Section 5.3 is rather minimal and not as elaborate as some published ones. Nevertheless, several characteristics set it apart from existing proposals:

1. It is faithful to the OO paradigm and popular OO languages. It accommodates references, shared mutable state, inheritance and dynamic dispatch.
2. It uses state-of-the-art mechanisms that facilitate modular and flexible reasoning, such as separation logic and abstract predicate families.
3. It is not postulated in isolation. The calculus has a clear relationship with an accepted system for verifying OO programs in a bottom-up way.

There is a relationship between observational equivalence of terms and the function Specs , because $\models_{V_1} \text{Sat}$ gives rise to a notion of observability from the specification point of view. In particular, two terms t and t' are observationally equivalent in this sense iff $t \sim t'$, where

$$t \sim t' \stackrel{\text{def}}{=} \text{Specs}(\{t\}) = \text{Specs}(\{t'\})$$

It is trivial to check that $\sim$ is an equivalence relation. If $\models_{V_1} \text{Sat } _$ is well-behaved, then $t \sim t' \Leftrightarrow \llbracket t \rrbracket = \llbracket t' \rrbracket$ (i.e. $t \sim t' \Leftrightarrow t \equiv t'$) by Corollary B.1.9 in Appendix B and Theorem 5.16.1.

CHAPTER 6

CONCLUSION

The state of the art in reasoning about software has seen several recent advances. In particular, proof systems that are based on separation logic emerged to provide intuitive, flexible and modular reasoning for OO programs. Their success is due to the synergy of the following insights:

- Separation is a natural way to think about shared mutable data. It leads to intuitive and simple correctness proofs for code that manipulates references.
- Abstract predicate families befit reasoning about object-orientation. They are logical abstractions that encapsulate state and hide representation details. Subclasses can change the representation and still satisfy the inherited specifications of routines.
- Giving two specifications to each routine resolves the tension between saying in detail what the particular body achieves, and saying what the general idea of the routine is that all subclasses must respect. Using separation logic and apfs in these specifications yields elegant proofs for OO programs.

The existing work forms a solid foundation for further investigation and research, which can of course proceed in many directions. This thesis tackled the following shortcomings in previous work:

- The basic apf mechanism can be too restrictive. Code sometimes depends on how a particular class implements and/or relates these abstractions, or properties that a whole class hierarchy will fulfill. This is especially important in the presence of multiple inheritance, where a class can relate and combine several concepts in various ways.

- Separation logic verification excluded the executable contracts that are present in existing OO code. These contracts may be weak and/or perform side-effects. They nevertheless capture useful design information which can be exploited in reasoning.
- It is sometimes easier to write code that is correct in the first place instead of hacking away and trying to verify code afterwards. Supporting this requires calculi for correctness by construction and refinement.

All the proposed solutions fit in nicely with the existing work of Parkinson and others [57, 58]. In particular:

- Export and axiom clauses can specify properties of and relationships between apfs that hold for particular classes and whole class hierarchies. The information is first verified and then made available as ordinary assumptions to reason about the code. Together with other generalisations, export and axiom clauses allow flexible reasoning about multiple inheritance.
- Separation logic can be used to verify that executable contracts will always hold at runtime. Central to such reasoning is the new concept of relative purity, which embraces side-effects in executable contracts.
- Freefinement can automatically provide a refinement calculus that interoperates smoothly with the rules for statement verification. This refinement calculus fits into a larger calculus for correctness by construction, which complements the separation logic system for bottom-up verification.

The contributions of this thesis thus advance the state of the art in the verification of existing OO programs, and in the development of new programs that are correct by design. The ultimate hope is that, in the long term, such techniques will encourage and enable engineers to produce software that is more reliable.

APPENDIX A

SEMANTICS OF THE PROOF SYSTEM

This appendix outlines the semantics and soundness proof of the verification system of chapter 3. Since our system is a conservative extension of Parkinson and Bierman’s system, the presentation closely follows the one in [58]. The most interesting difference is the treatment of export and axiom information in the soundness proof of the program verification rule (Theorem A.4 below).

The semantics of the logical formula is defined in terms of a state σ , an interpretation of predicate symbols $\mathcal{I}$, and an interpretation of logical variables $\mathcal{L}$. The interpretation $\mathcal{I}$ maps predicate names to their definitions, whereas a definition maps a list of arguments to a set of states:

$$\begin{aligned}\mathcal{I} &: \text{Preds} \rightarrow (\text{Vals}^* \rightarrow \mathcal{P}(\Sigma)) \\ \mathcal{L} &: \text{Vars} \rightarrow \text{Vals}\end{aligned}$$

Predicates are defined in the standard way:

$$\sigma, \mathcal{I}, \mathcal{L} \models \text{pred}(\overline{X}) \Leftrightarrow \sigma \in (\mathcal{I}(\text{pred})(\mathcal{L}(\overline{X})))$$

Definition A.1. $\mathcal{I} \models \Delta$ iff $\sigma, \mathcal{I}, \mathcal{L} \models \Delta$ for all σ and $\mathcal{L}$.

Under mild syntactic restrictions on predicate definitions, obeyed in this thesis and detailed in [57], one can show that every set of disjoint predicate definitions is satisfiable:

Lemma A.1. For any set of definitions $W_1, \dots, W_m, D_1, \dots, D_n$ where W_i has form $w_i(\overline{x}_i) = Q_i$ and D_j is listed in class G_j , there exists an interpretation $\mathcal{I}$ such that $\mathcal{I} \models [\bigwedge_{i \in 1..m} \forall \overline{x}_i. w_i(\overline{x}_i) \Leftrightarrow Q_i] \wedge [\bigwedge_{j \in 1..n} \text{apf}_{G_j}(D_j)]$ provided that no two distinct definitions in the set define the same predicate.

The semantics of our proof system's judgements is defined next. We do not define the semantics of $\vdash_e$ and $\vdash_a$ explicitly, since we work with their premises (valid logical formulae whose existence is guaranteed) instead. For triples, the usual partial-correctness semantics for separation logic is used: if the precondition holds in the start state, then 1) the statement will not fault (access unallocated memory, for example), and 2) if the statement terminates, then the postcondition holds in the resulting state.

Definition A.2. $\mathcal{I} \models_n \{P\}_s\{Q\}$ iff whenever $\sigma, \mathcal{I}, \mathcal{L} \models P$ then $\forall m \leq n$.

1. $\sigma, s \longrightarrow^m \text{fault}$ does not hold, and
2. if $\sigma, s \longrightarrow^m \sigma', \text{skip}$ then $\sigma', \mathcal{I}, \mathcal{L} \models Q$

The index n deals with mutual recursion in method definitions. $\mathcal{I} \models_n \Gamma$ means that all methods in Γ meet their specifications when executed for up to n steps.

Definition A.3 (Method verification semantics). If m in G is non-abstract, let s denote its body.

$\mathcal{I}, \Gamma \models_0 G.m \mapsto (\bar{x}, \{P\}_s\{Q\})$ always holds.

$\mathcal{I}, \Gamma \models_{n+1} G.m \mapsto (\bar{x}, \{P\}_s\{Q\})$ iff

$\mathcal{I} \models_n \Gamma \Rightarrow \mathcal{I} \models_{n+1} \{P * \mathbf{Current} : G\}_s\{Q\}$

if G is non-abstract and true otherwise.

$\mathcal{I}, \Gamma \models_0 G::m \mapsto (\bar{x}, \{S\}_s\{T\})$ always holds.

$\mathcal{I}, \Gamma \models_{n+1} G::m \mapsto (\bar{x}, \{S\}_s\{T\})$ iff $\mathcal{I} \models_n \Gamma \Rightarrow \mathcal{I} \models_{n+1} \{S\}_s\{T\}$

$\mathcal{I} \models_n \Gamma$ iff $\forall \text{methodspec} \in \Gamma. \mathcal{I}, \Gamma \models_n \text{methodspec}$

We next define the semantics of the statement judgement.

Definition A.4. $\Delta; \Gamma \models \{P\}_s\{Q\}$ iff for all $\mathcal{I}$ and n , if $\mathcal{I} \models \Delta$ and $\mathcal{I} \models_n \Gamma$, then $\mathcal{I} \models_{n+1} \{P\}_s\{Q\}$

In other words, for all interpretations which satisfy the assumptions Δ , if all methods in Γ meet their specifications for up to n steps, then s meets its specification for up to $n + 1$ steps.

The judgements are sound with respect to their semantics.

Lemma A.2.

1. If $\Delta; \Gamma \vdash_m \dots m \dots$, then $\forall \mathcal{I}. \text{ if } \mathcal{I} \models \Delta \text{ then for all } n \text{ and every spec of } m \text{ we have } \mathcal{I}, \Gamma \models_n \text{spec}$
2. If $\Delta; \Gamma \vdash_s \{P\}_s\{Q\}$ then $\Delta; \Gamma \models \{P\}_s\{Q\}$

Whenever a judgement is derivable under weak assumptions, it can also be derived under stronger ones.

Lemma A.3.

1. If $\Delta; \Gamma \vdash_m \dots m \dots$ and $\Delta' \Rightarrow \Delta$, then $\Delta'; \Gamma \vdash_m \dots m \dots$
2. If $\Delta; \Gamma \vdash_s \{P\}s\{Q\}$ and $\Delta' \Rightarrow \Delta$, then $\Delta'; \Gamma \vdash_s \{P\}s\{Q\}$
3. If $\Delta_{APF}, \Delta_E, \Delta_A; \Gamma \vdash_c L$ and $\Delta' \Rightarrow \Delta_{APF}$,
then $\Delta', \Delta_E, \Delta_A; \Gamma \vdash_c L$

Finally, here is the soundness statement and detailed proof sketch of the program verification rule.

Theorem A.4. If a program and its main body s can be proved with the program verification rule, then $\forall \mathcal{I}, n. \mathcal{I} \models_n \{\text{true}\}s\{\text{true}\}$.

Proof.

1. *The goal.* We have to prove $\forall \mathcal{I}, n. \mathcal{I} \models_n \{\text{true}\}s\{\text{true}\}$, which abbreviates $\forall \mathcal{I}, n.$ whenever $\sigma, \mathcal{I}, \mathcal{L} \models \text{true}$, then $\forall m \leq n. 1) \sigma, s \xrightarrow{m} \text{fault}$ does not hold, and 2) if $\sigma, s \xrightarrow{m} \sigma', \text{skip}$ then $\sigma', \mathcal{I}, \mathcal{L} \models \text{true}$. This can be simplified to $\forall n. \sigma, s \xrightarrow{n} \text{fault}$ does not hold.
2. *Strengthened assumptions.* Let $\Delta_T \stackrel{\text{def}}{=} \bigwedge_{i \in 1..t} \text{apf}(L_i)$, where $L_1 \dots L_t$ are all classes in the program. By Lemma A.3, we can strengthen the assumptions under which all classes and the main body have been verified. For every class L_i , we have $\Delta_T, \Delta_E, \Delta_A; \Gamma \vdash_c L_i$, and $\Delta_T \wedge \Delta_E \wedge \Delta_A; \Gamma \vdash_s \{\text{true}\}s\{\text{true}\}$ also holds for the main body s .
3. *The interpretation $\mathcal{I}'$.* Since $\Delta_T \wedge \Delta_E \wedge \Delta_A; \Gamma \vdash_s \{\text{true}\}s\{\text{true}\}$, Lemma A.2 guarantees $\Delta_T \wedge \Delta_E \wedge \Delta_A; \Gamma \models \{\text{true}\}s\{\text{true}\}$. This abbreviates $\forall \mathcal{I}, n.$ if $\mathcal{I} \models \Delta_T \wedge \Delta_E \wedge \Delta_A$ and $\mathcal{I} \models_n \Gamma$, then $\mathcal{I} \models_{n+1} \{\text{true}\}s\{\text{true}\}$, which can be simplified to $\forall \mathcal{I}, n.$ if $\mathcal{I} \models \Delta_T \wedge \Delta_E \wedge \Delta_A$ and $\mathcal{I} \models_n \Gamma$, then $\forall m \leq n + 1. \sigma, s \xrightarrow{m} \text{fault}$ does not hold. Now if we can find an $\mathcal{I}'$ such that $\mathcal{I}' \models \Delta_T \wedge \Delta_E \wedge \Delta_A$ and $\forall n. \mathcal{I}' \models_n \Gamma$, then we can instantiate $\mathcal{I}$ to $\mathcal{I}'$ in the formula and simplify to obtain $\forall n. \sigma, s \xrightarrow{n} \text{fault}$ does not hold. Therefore $\mathcal{I}'$ serves as a witness that s will never fault, which is exactly our goal.

Let $\mathcal{I}'$ be the interpretation whose existence is guaranteed by Lemma A.1 for all the where and define clauses in the program. Clearly $\mathcal{I}' \models \Delta_T$. We next prove $\mathcal{I}' \models \Delta_E$ and then $\mathcal{I}' \models \Delta_A$.

4. *Satisfiability of Δ_E .* Consider an arbitrary export clause $E = P$ **where** $\{w_1(\overline{x}_1) = Q_1; \dots; w_n(\overline{x}_n) = Q_n\}$ in class L. Since $apf(L) \vdash_e E$, we know $[apf(L) \wedge (\bigwedge_{i \in 1..n} \forall \overline{x}_i \cdot w_i(\overline{x}_i) \Leftrightarrow Q_i)] \Rightarrow P$. The interpretation $\mathcal{I}'$ satisfies the antecedent, so we also have $\mathcal{I}' \models P$. Therefore $\mathcal{I}' \models \Delta_E$, and $\mathcal{I}' \models \Delta_T \wedge \Delta_E$.
5. *Satisfiability of Δ_A .* We prove this by induction. If class G has children $H_1 \dots H_k$, let $level(G) \stackrel{\text{def}}{=} 1 + \max(0, level(H_1), \dots, level(H_k))$. Furthermore, $P(n) \stackrel{\text{def}}{=} \forall G$ in the program such that $level(G) \leq n$ and for all axiom clauses $a: P$ in the listing of G, $(\Delta_T \wedge \Delta_E) \Rightarrow axiominfo(G, a: P)$.

- Base case. Consider an arbitrary class G with $level(G) \leq 1$ and an axiom clause $a: P$ appearing in it. G has no subclasses, and
 - (a) If G is abstract, there are no objects with dynamic type G or a subtype thereof, thus $axiominfo(G, a: P)$ holds vacuously and $\Delta_T \wedge \Delta_E$ implies it.
 - (b) If G is non-abstract, then the only objects whose dynamic type is a subtype of G are direct instances of G. Since $(\Delta_T \wedge \Delta_E \wedge \mathbf{Current} : G) \Rightarrow P$ by the Implication premise, we therefore also know $(\Delta_T \wedge \Delta_E) \Rightarrow axiominfo(G, a: P)$.

Thus $P(1)$ holds.

- Step case. Suppose $P(n)$ holds. Now consider a class G at level $n+1$ with axiom clause $a: P$. Every child H of G must list a, say $a: Q$. By the induction hypothesis we know $(\Delta_T \wedge \Delta_E) \Rightarrow axiominfo(H, a: Q)$, and by the Parent Consistency premise of $a: Q$ we know $(\Delta_T \wedge \Delta_E \wedge Q) \Rightarrow P$. Therefore $(\Delta_T \wedge \Delta_E) \Rightarrow axiominfo(H, a: P)$. We have $(\Delta_T \wedge \Delta_E) \Rightarrow axiominfo(G, a: P)$ if G is abstract, and the same holds if G is non-abstract since the Implication premise of $a: P$ guarantees $(\Delta_T \wedge \Delta_E \wedge \mathbf{Current} : G) \Rightarrow P$. Thus $P(n+1)$ holds.

So $\mathcal{I}' \models \Delta_T \wedge \Delta_E \wedge \Delta_A$.

6. *Wrapping up.* We still have to prove $\forall n \cdot \mathcal{I}' \models_n \Gamma$. Let m be an arbitrary method in the program. Since $\Delta_T \wedge \Delta_E \wedge \Delta_A; \Gamma \vdash_m m$, by Lemma A.2 we know for all n and every spec of m that $\mathcal{I}', \Gamma \models_n \text{spec}$. Thus $\forall n \cdot \forall \text{methodspec} \in \Gamma \cdot \mathcal{I}', \Gamma \models_n \text{methodspec}$, in other words $\forall n \cdot \mathcal{I}' \models_n \Gamma$. $\square$

APPENDIX B

ANTITONE GALOIS CONNECTIONS

Lemma 5.2 established that an antitone Galois connection exists between the functions Specs and Terms:

$$X \subseteq \text{Terms}(Y) \Leftrightarrow Y \subseteq \text{Specs}(X) \quad (*)$$

Theorems derived from this equivalence come in pairs because of the symmetry between Specs and Terms. Here are a few well-known ones together with their proofs:

Corollary B.1.

B.1.1. $X \subseteq \text{Terms}(\text{Specs}(X))$

B.1.2. $Y \subseteq \text{Specs}(\text{Terms}(Y))$

B.1.3. $X \subseteq X' \Rightarrow \text{Specs}(X) \supseteq \text{Specs}(X')$

B.1.4. $Y \subseteq Y' \Rightarrow \text{Terms}(Y) \supseteq \text{Terms}(Y')$

B.1.5. $X \subseteq X' \Rightarrow \text{Terms}(\text{Specs}(X)) \subseteq \text{Terms}(\text{Specs}(X'))$

B.1.6. $Y \subseteq Y' \Rightarrow \text{Specs}(\text{Terms}(Y)) \subseteq \text{Specs}(\text{Terms}(Y'))$

B.1.7. $\text{Specs}(\text{Terms}(\text{Specs}(X))) = \text{Specs}(X)$

B.1.8. $\text{Terms}(\text{Specs}(\text{Terms}(Y))) = \text{Terms}(Y)$

B.1.9. $\text{Specs}(X) \subseteq \text{Specs}(X') \Leftrightarrow \text{Terms}(\text{Specs}(X)) \supseteq \text{Terms}(\text{Specs}(X'))$

B.1.10. $\text{Terms}(Y) \subseteq \text{Terms}(Y') \Leftrightarrow \text{Specs}(\text{Terms}(Y)) \supseteq \text{Specs}(\text{Terms}(Y'))$

$$\text{B.1.11. } \text{Specs}(X \cup X') = \text{Specs}(X) \cap \text{Specs}(X')$$

$$\text{B.1.12. } \text{Terms}(Y \cup Y') = \text{Terms}(Y) \cap \text{Terms}(Y')$$

Proof.

B.1.1. In (*), instantiate Y with $\text{Specs}(X)$.

B.1.3. $X \subseteq \text{"Assumption" } X' \subseteq \text{"B.1.1" } \text{Terms}(\text{Specs}(X'))$. In (*), instantiate Y with $\text{Specs}(X')$.

B.1.5. If $X \subseteq X'$, then $\text{Specs}(X) \supseteq \text{Specs}(X')$ holds by B.1.3. The result follows from B.1.4.

B.1.7. From B.1.1 and B.1.3 follows $\text{Specs}(X) \supseteq \text{Specs}(\text{Terms}(\text{Specs}(X)))$. Instantiating Y with $\text{Specs}(X)$ in B.1.2 yields $\text{Specs}(X) \subseteq \text{Specs}(\text{Terms}(\text{Specs}(X)))$.

B.1.9. $\Rightarrow$ holds by B.1.4. From $\text{Terms}(\text{Specs}(X)) \supseteq \text{Terms}(\text{Specs}(X'))$ and B.1.3, $\text{Specs}(\text{Terms}(\text{Specs}(X))) \subseteq \text{Specs}(\text{Terms}(\text{Specs}(X')))$. $\text{Specs}(X) \subseteq \text{Specs}(X')$ by B.1.7.

B.1.11. Proof by indirect equality. For arbitrary Y :

$$\begin{aligned} & Y \subseteq \text{Specs}(X \cup X') \\ \Leftrightarrow & \quad \{\text{By (*)}\} \\ & X \cup X' \subseteq \text{Terms}(Y) \\ \Leftrightarrow & \quad \{\text{Set theory}\} \\ & X \subseteq \text{Terms}(Y) \quad \wedge \quad X' \subseteq \text{Terms}(Y) \\ \Leftrightarrow & \quad \{\text{By (*)}\} \\ & Y \subseteq \text{Specs}(X) \quad \wedge \quad Y \subseteq \text{Specs}(X') \\ \Leftrightarrow & \quad \{\text{Set theory}\} \\ & Y \subseteq \text{Specs}(X) \cap \text{Specs}(X') \end{aligned}$$

□

BIBLIOGRAPHY

- [1] Ralph-Johan Back. Correctness preserving program refinements: Proof theory and applications. *Mathematical Centre Tracts*, 131, 1980.
- [2] Ralph-Johan Back and Joakim von Wright. *Refinement Calculus: A Systematic Introduction*. Springer-Verlag, 1998. Graduate Texts in Computer Science.
- [3] Mike Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: An overview. In *CASSIS '05*, volume 3362 of *LNCS*, pages 49–69. Springer, 2005.
- [4] Paul L. Bergstein. Object-preserving class transformations. In *Conference proceedings on Object-oriented programming systems, languages, and applications*, OOPSLA '91, pages 299–313, New York, NY, USA, 1991. ACM.
- [5] Paulo Borba, Augusto Sampaio, Ana Cavalcanti, and Márcio Cornélio. Algebraic reasoning for object-oriented programming. *Science of Computer Programming*, 52(1-3):53–100, 2004.
- [6] Paulo Borba, Augusto Sampaio, and Márcio Cornélio. A refinement algebra for object-oriented programming. In Luca Cardelli, editor, *ECOOP '03: Object-Oriented Programming*, volume 2743 of *Lecture Notes in Computer Science*, pages 457–482. Springer Berlin Heidelberg, 2003.
- [7] Stephen Brookes. A semantics for concurrent separation logic. *Theor. Comput. Sci.*, 375:227–270, April 2007.
- [8] James Brotherston, Richard Bornat, and Cristiano Calcagno. Cyclic proofs of program termination in separation logic. In *POPL '08: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 101–112, New York, NY, USA, 2008. ACM.

- [9] Lilian Burdy, Yoonsik Cheon, David R. Cok, Michael D. Ernst, Joseph R. Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer (STTT)*, 7(3):212–232, 2005.
- [10] Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. In *POPL ’09: Proceedings of the 36th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 289–300, New York, NY, USA, 2009. ACM.
- [11] Luca Cardelli. A semantics of multiple inheritance. *Inf. Comput.*, 76(2-3):138–164, 1988.
- [12] Yoonsik Cheon, Gary Leavens, Murali Sitaraman, and Stephen Edwards. Model variables: cleanly supporting abstraction in design by contract: Research articles. *Softw. Pract. Exper.*, 35(6):583–599, 2005.
- [13] Yoonsik Cheon and Gary T. Leavens. A runtime assertion checker for the Java modeling language (JML). In *Proceedings of the international conference on Software engineering research and practice (SERP ’02), Las Vegas*, pages 322–328. CSREA Press, 2002.
- [14] Wei-Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. Enhancing modular OO verification with separation logic. In *POPL ’08: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 87–99, New York, NY, USA, 2008. ACM.
- [15] Ilinca Ciupa, Andreas Leitner, Manuel Oriol, and Bertrand Meyer. Experimental assessment of random testing for object-oriented software. In *ISSTA ’07: Proceedings of the 2007 international symposium on Software testing and analysis*, pages 84–94, New York, NY, USA, 2007. ACM.
- [16] D. Clarke, S. Drossopoulou, P. Müller, J. Noble, and T. Wrigstad. Aliasing, confinement, and ownership in object-oriented programming (IWACO). In *Object-Oriented Technology. ECOOP 2008 Workshop Reader*, volume 5475 of *LNCS*, pages 30–41. Springer, 2008.
- [17] Á. Darvas and K. Rustan M. Leino. Practical reasoning about invocations and implementations of pure methods. In *FASE ’07*, volume 4422 of *LNCS*, pages 336–351. Springer, 2007.

- [18] Á. Darvas and P. Müller. Reasoning about method calls in interface specifications. *Journal of Object Technology (JOT)*, 5(5):59–85, June 2006.
- [19] Ádám Darvas and Peter Müller. Faithful mapping of model classes to mathematical structures. *IET Software*, 2(6):477–499, 2008.
- [20] Ewen Denney. Simply-typed underdeterminism. *Journal of Computer Science and Technology*, 13:491–508, 1998.
- [21] Ewen Denney. A theory of program refinement. Technical Report ECS-LFCS-99-412, University of Edinburgh, 1999.
- [22] Edsger W. Dijkstra. The humble programmer. *Commun. ACM*, 15(10):859–866, October 1972.
- [23] Dino Distefano and Matthew J. Parkinson J. jStar: towards practical verification for Java. In *OOPSLA '08: Proceedings of the 23rd ACM SIGPLAN conference on Object-oriented programming systems languages and applications*, pages 213–226, New York, NY, USA, 2008. ACM.
- [24] Johan Dovland, Einar Broch Johnsen, Olaf Owe, and Martin Steffen. Incremental reasoning for multiple inheritance. In *IFM '09*, pages 215–230, Berlin, Heidelberg, 2009. Springer.
- [25] Sophia Drossopoulou, Adrian Francalanza, Peter Müller, and Alexander J. Summers. A unified framework for verification techniques for object invariants. In Jan Vitek, editor, *ECOOP*, volume 5142 of *LNCS*, pages 412–437. Springer, 2008.
- [26] ECMA International. *Standard ECMA-367. Eiffel: Analysis, Design and Programming Language*. 2nd edition, June 2006.
- [27] Margaret A. Ellis and Bjarne Stroustrup. *The annotated C++ reference manual*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990.
- [28] EVE. The Eiffel Verification Environment. <http://se.inf.ethz.ch/research/eve/>.
- [29] Gobosoft. The Gobo Eiffel Structure Library. <http://www.gobosoft.com/eiffel/gobo/structure/index.html>.

- [30] Bhargav S. Gulavani, Supratik Chakraborty, Ganesan Ramalingam, and Aditya V. Nori. Bottom-up shape analysis. In *SAS '09*, volume 5673 of *LNCS*, pages 188–204. Springer, 2009.
- [31] Eric C. R. Hehner. *A practical theory of programming*. Springer-Verlag New York, Inc., New York, NY, USA, 1993.
- [32] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12:576–580, October 1969.
- [33] Bart Jacobs and Frank Piessens. Inspector methods for state abstraction. *Journal of Object Technology*, 6(5):55–75, June 2007.
- [34] Bart Jacobs and Frank Piessens. The VeriFast program verifier. Technical Report CW-520, Katholieke Universiteit Leuven, August 2008.
- [35] Bart Jacobs, Jan Smans, and Frank Piessens. Verifying the composite pattern using separation logic. *SAVCBS Composite pattern challenge track*, 2008.
- [36] I. T. Kassios. The dynamic frames theory. *Formal Aspects of Computing*, 2010. To appear.
- [37] Ioannis T. Kassios. Dynamic frames: Support for framing, dependencies and sharing without restrictions. In Jayadev Misra, Tobias Nipkow, and Emil Sekerinski, editors, *FM*, volume 4085 of *LNCS*, pages 268–283. Springer, 2006.
- [38] Neelakantan R. Krishnaswami, Lars Birkedal, Jonathan Aldrich, and John C. Reynolds. Idealized ML and Its Separation Logic. Draft at <http://www.cs.cmu.edu/~neelk/idealized-ml-draft.pdf>, 2006.
- [39] G. T. Leavens, K. R. M. Leino, and P. Müller. Specification and verification challenges for sequential object-oriented programs. *Formal Aspects of Computing*, 19(2):159–189, 2007.
- [40] Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: a behavioral interface specification language for Java. *SIGSOFT Softw. Eng. Notes*, 31(3):1–38, 2006.
- [41] Gary T. Leavens and Peter Müller. Information hiding and visibility in interface specifications. In *ICSE*, pages 385–395. IEEE Computer Society, 2007.

- [42] K. R. M. Leino and P. Müller. Verification of equivalent-results methods. In *ESOP '08*, volume 4960 of *LNCS*, pages 307–321. Springer, 2008.
- [43] K. Rustan M. Leino and Rajit Manohar. Joining specification statements. *Theor. Comput. Sci.*, 216(1-2):375–394, 1999.
- [44] K. Rustan M. Leino and Peter Müller. Object invariants in dynamic contexts. In Martin Odersky, editor, *ECOOP*, volume 3086 of *LNCS*, pages 491–516. Springer, 2004.
- [45] K Rustan M Leino and Wolfram Schulte. A verifying compiler for a multi-threaded object-oriented language. *Software System Reliability and Security*, 9:351–416, 2007.
- [46] Chenguang Luo and Shengchao Qin. Separation logic for multiple inheritance. *Electr. Notes Theor. Comput. Sci.*, 212:27–40, 2008.
- [47] Bertrand Meyer, Ilinca Ciupa, Andreas Leitner, and Lisa Ling Liu. Automatic testing of object-oriented software. In *SOFSEM '07*, volume 4362 of *LNCS*, pages 114–129. Springer, 2007.
- [48] Anna Mikhajlova and Emil Sekerinski. Class refinement and interface refinement in object-oriented programs. In John Fitzgerald, Cliff B. Jones, and Peter Lucas, editors, *FME '97: Industrial Applications and Strengthened Foundations of Formal Methods*, volume 1313 of *Lecture Notes in Computer Science*, pages 82–101. Springer Berlin Heidelberg, 1997.
- [49] Carroll Morgan. The specification statement. *ACM Trans. Program. Lang. Syst.*, 10:403–419, July 1988.
- [50] Carroll Morgan. *Programming from specifications (2nd ed.)*. Prentice Hall International (UK) Ltd., Hertfordshire, UK, 1994.
- [51] Joseph M. Morris. A theoretical basis for stepwise refinement and the programming calculus. *Sci. Comput. Program.*, 9:287–306, December 1987.
- [52] P. Müller, A. Poetzsch-Heffter, and G. T. Leavens. Modular invariants for layered object structures. *Science of Computer Programming*, 62(3):253–286, 2006.
- [53] David A. Naumann. Observational purity and encapsulation. *Theor. Comput. Sci.*, 376(3):205–224, 2007.

- [54] Peter W. O'Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *CSL '01*, volume 2142 of *LNCS*, pages 1–19. Springer, 2001.
- [55] Richard F. Paige and Jonathan S. Ostroff. ERC – an object-oriented refinement calculus for Eiffel. *Form. Asp. Comput.*, 16(1):51–79, April 2004.
- [56] Matthew Parkinson and Gavin Bierman. Separation logic and abstraction. In *POPL '05: Proceedings of the 32nd annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 247–258, New York, NY, USA, 2005. ACM.
- [57] Matthew J. Parkinson. Local reasoning for Java. PhD thesis. Technical Report UCAM-CL-TR-654, University of Cambridge, Computer Laboratory, November 2005.
- [58] Matthew J. Parkinson and Gavin M. Bierman. Separation logic, abstraction and inheritance. In *POPL '08: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 75–86, New York, NY, USA, 2008. ACM.
- [59] Mohammad Raza and Philippa Gardner. Footprints in local reasoning. *Logical Methods in Computer Science*, 5(2:4):1–27, 2009.
- [60] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS '02: Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, Washington, DC, USA, 2002. IEEE Computer Society.
- [61] A. Rudich, Á. Darvas, and P. Müller. Checking well-formedness of pure-method specifications. In *FM '08*, volume 5014 of *LNCS*, pages 68–83. Springer, 2008.
- [62] Alexandru Salcianu and Martin C. Rinard. Purity and side effect analysis for Java programs. In *VMCAI '05*, volume 3385 of *LNCS*, pages 199–215, 2005.
- [63] Bernd Schoeller, Tobias Widmer, and Bertrand Meyer. Making specifications complete through models. In *Architecting Systems with Trustworthy Components*, volume 3938 of *LNCS*, pages 48–70. Springer, 2006.
- [64] A. J. Summers, S. Drossopoulou, and P. Müller. The need for flexible object invariants. In *IWACO '09: International Workshop on Aliasing*,

- Confinement and Ownership in Object-Oriented Programming*, pages 1–9, New York, NY, USA, 2009. ACM.
- [65] Stephan van Staden and Cristiano Calcagno. Reasoning about multiple related abstractions with MultiStar. In *Proceedings of the ACM international conference on Object oriented programming systems languages and applications*, OOPSLA '10, pages 504–519, New York, NY, USA, 2010. ACM.
- [66] Stephan van Staden, Cristiano Calcagno, and Bertrand Meyer. Verifying executable object-oriented specifications with separation logic. In Theo D'Hondt, editor, *ECOOP*, volume 6183 of *Lecture Notes in Computer Science*, pages 151–174. Springer, 2010.
- [67] Stephan van Staden, Cristiano Calcagno, and Bertrand Meyer. Free-finement. In *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '12, pages 7–18, New York, NY, USA, 2012. ACM.

Publications arising from the work of this thesis are [65], [66], and [67].