

Number 654



**UNIVERSITY OF
CAMBRIDGE**

Computer Laboratory

Local reasoning for Java

Matthew J. Parkinson

November 2005

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2005 Matthew J. Parkinson

This technical report is based on a dissertation submitted August 2005 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Churchill College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/TechReports/>

ISSN 1476-2986

Abstract

This thesis develops the local reasoning approach of separation logic for common forms of modularity such as abstract datatypes and objects. In particular, this thesis focuses on the modularity found in the Java programming language.

We begin by developing a formal semantics for a core imperative subset of Java, *Middleweight Java* (MJ), and then adapt separation logic to reason about this subset. However, a naïve adaptation of separation logic is unable to reason about encapsulation or inheritance: it provides no support for modularity.

First, we address the issue of encapsulation with the novel concept of an *abstract predicate*, which is the logical analogue of an abstract datatype. We demonstrate how this method can encapsulate state, and provide a mechanism for *ownership transfer*: the ability to transfer state safely between a module and its client. We also show how abstract predicates can be used to express the calling protocol of a class.

However, the encapsulation provided by abstract predicates is too restrictive for some applications. In particular, it cannot reason about multiple datatypes that have shared read-access to state, for example list iterators. To compensate, we alter the underlying model to allow the logic to express properties about read-only references to state. Additionally, we provide a model that allows both sharing and disjointness to be expressed directly in the logic.

Finally, we address the second modularity issue: inheritance. We do this by extending the concept of abstract predicates to *abstract predicate families*. This extension allows a predicate to have multiple definitions that are indexed by class, which allows subclasses to have a different internal representation while remaining behavioural subtypes. We demonstrate the usefulness of this concept by verifying a use of the visitor design pattern.

Acknowledgments

I should like to begin by thanking my supervisors, Gavin Bierman and Andrew Pitts. Their careful stewardship has guided me through this thesis and developed my understanding of research immensely. I would also like to thank my colleagues in the Computer Laboratory for providing a supportive and friendly environment to work in. Especially Alisdair Wren for his attempts to teach me grammar, and for his assiduous proof reading abilities.

The *East London Massive* (Josh Berdine, Richard Bornat, Cristiano Calcagno, Dino Distefano, Ivana, Mijajlović, and Peter O’Hearn) also deserve many thanks for providing exciting discussions. Particular thanks to Peter O’Hearn for his comments that have help focus my research.

I should like to thank the following people for feedback and discussions: Moritz Becker, Nick Benton, Sophia Drossoupolou, Manuel Fähndrich, Philippa Gardner, Tony Hoare, Andrew Kennedy, Alan Lawrence, Ronald Middlekoop, David Naumann, Uday Reddy, John Reynolds, Chris Ross, Claudio Russo, Peter Sewell, Mark Shinwell, Matthew Smith, Sam Staton, Gareth Stoye, Hayo Thielecke, David Walker, Hongseok Yang and Nobuko Yoshida. I should also like to thank all the people whom I have forgotten to thank, as well as the anonymous referees for POPL 2005 and WOOD 2003.

Finally, I should like to thank my family and friends for their love and friendship.

Contents

1	Introduction	9
1.1	Background	10
1.2	Content and contribution	16
1.3	Related work	17
2	Middleweight Java	19
2.1	Syntax	19
2.2	Types	21
2.3	Operational semantics	27
2.4	Well-typed configuration	34
2.5	Type soundness	35
2.6	Related Work	41
3	A logic for Java	43
3.1	Inner MJ	43
3.2	Storage model	44
3.3	Assertion language	46
3.4	Java rules	48
3.5	Dynamic dispatch	52
3.6	Example: Tree copy	54
3.7	Semantics	56
3.8	Encapsulation and inheritance	61
4	Abstract predicates	63
4.1	Proof rules	63
4.2	Examples	65
4.3	Semantics	69
4.4	Discussion	72
5	Read sharing	79
5.1	Counting	80
5.2	Example: List iterators	80
5.3	Disjointness	84
5.4	Abstract semantics	85
5.5	Concrete models	86
5.6	Concluding remarks	89

6 Inheritance	91
6.1 The problem	91
6.2 Abstract predicate families	92
6.3 Examples	94
6.4 Semantics	100
6.5 Reverification for inherited methods	106
7 Conclusion	109
7.1 Open questions	109
Nomenclature	111
Bibliography	115

1

Introduction

Software is unreliable. Currently, operating systems are susceptible to viruses, and commercial software packages are sold with disclaimers not guarantees: software licences do not say the program will function as intended or that it will not damage other data. Instead, they say that the software is not sold for any purpose, it might damage your data and that if it does, there is no entitlement to compensation. In many ways the situation is getting worse as software products grow ever larger. However, the programmer has several weapons of varying power to improve software reliability, in particular: types, testing and verification.

Types Types allow the compiler to guarantee certain errors do not occur in a program; for example, multiplying an integer by a string. Each operation, or function, is specified to work with certain types of value and the compiler guarantees that the program does not perform operations on the wrong type of data. However, types fall short of eliminating all errors. Consider the following:

```
int [] x = new int[30];
for(int n = 1; n <= 30; n++)
    x[n] = 1;
```

This program fragment creates an array of 30 elements and assigns 1 to elements 1 to 30. However, in Java arrays are numbered from 0: the array is defined from 0 to 29, not 1 to 30. Therefore, on the last iteration of the loop, the assignment will fail. Fortunately, Java provides dynamic checks for accesses outside the range of an array, so it will throw an exception. The problem is worse for programming languages without dynamic checks, like C, as they have unpredictable behaviour in the error cases.

Standard type systems do not eliminate this type of error. These properties can be checked with dependent types [87] or static analyses [29], but there are many problems where type checking or static analysis is undecidable or infeasible.

Testing Testing can be used to check more complicated properties. The program code is run many times against different inputs to see if anything “goes wrong”. The inputs are selected to attempt to cover all the possible executions, and then the program’s outputs are compared with the expected results.

To help automate testing and localise errors, *run-time assertions* are used to find these “wrong” states: for example, Microsoft Office has a quarter of a million run-time assertions [44]. Consider the following:

```

if(debug && x.length < 30) throw new Error();
for(int n = 0; n <30 ; n++)
    x[n] = 1;

```

The `if` statement checks that the array is of an appropriate length before proceeding and, if not, throws a descriptive error. This helps track the location of the bug: the code fails quickly. The key weakness of testing is that it is not possible to exhaustively test code: testing cannot guarantee to remove all errors as we cannot consider all possible inputs.

Program verification Finally we turn to program verification: the “holy grail” of software reliability. Rather than checking that a property holds for a number of inputs, program verification allows one to formally prove, for all possible runs of the program, that the property will hold. This allows us to make strong guarantees about the software. Floyd [39] and Hoare [45] pioneered the use of logic for program verification. Hoare’s logic provides pre- and post-conditions for each command, written $\{P\}C\{Q\}$, which is read “if property P holds of the start state and C terminates, then the final state will satisfy Q .” This is called *partial correctness*; total correctness additionally guarantees the program will terminate. Hoare provided axioms and rules to derive pre- and post-conditions for a given program, for example:

$$\{P[E/x]\}x := E\{P\} \qquad \frac{\{P\}C1\{Q\} \quad \{Q\}C2\{R\}}{\{P\}C1;C2\{R\}}$$

With rules of this form properties about a program’s final values can be shown. For example one can prove an implementation of Euclid’s algorithm actually gives the greatest common divisor of its two inputs.

Despite its many advantages program verification is rarely used in practice. One key reason for this lack of use is that current program verification techniques do not scale. References and pointers are a particular impediment for scalability as they allow apparently unrelated pieces of code to affect each other’s behaviour (for a detailed background of this problem see §1.1.2).

However, recently, O’Hearn, Reynolds and Yang [65] have developed an approach called *local reasoning*. Local reasoning deals with pointers and references, and has the potential to scale, but thus far it has only been applied to C-like languages with no support for modularity. This thesis builds a logic for Java by extending local reasoning to account for the modularity found in Java, particularly *encapsulation* and *inheritance*. Encapsulation is where the internal representation of an object is hidden from its clients, leaving the programmer free to change their internal representation without requiring changes to the clients. Inheritance allows a class to be extended with new features. This extended class can then be used in place of the original class—this is a defining characteristic of object-oriented programming languages.

The rest of the introduction begins by discussing the relevant background material. In §1.2 we summarize the content and contribution of this thesis. We conclude the introduction by discussing and comparing related methods for reasoning about Java.

1.1 Background

We now provide a summary of the relevant background work. We assume the reader is familiar with Java, first-order logic and Hoare logic. There are two key difficulties when reasoning about object-oriented programs: reference aliasing and dynamic dispatch. We begin by discussing dynamic dispatch and a method for reasoning about it known as *behavioural subtyping*. We then present the issues associated with reference aliasing and some approaches for reasoning about aliases.

1.1.1 Dynamic dispatch

In object-oriented languages like Java, the static type does not determine the precise class of the runtime values. If a variable has static type C , then the runtime values could be an object of class C or any subtype of C . This definition allows any instance of a subtype to be used in place of its supertype without causing errors.¹ For a class D to be a subclass of C , D must have all the methods and fields of C and they must have the same types. We will refer to the type of a variable or field as the *static type*, and the type of the object it actually refers to as the *dynamic type*. For soundness, the dynamic type of a variable or field is always a subtype of the static type (see §2.4).

Method calls in Java are determined by the dynamic type of the receiver not by its static type. This is known as *dynamic dispatch*. Consider the following:

```
C x; x = new D(); x.m();
```

This will not call C 's m method, but instead will call D 's m method. Dynamic dispatch may seem complicated, but it can be particularly useful. Consider the following:

```
Shape [] shapes;
...
for(Shape shape : shapes) {
    shape.draw();
}
```

We omit the definition of `Shape` and its subclasses `Circle` and `Triangle`. This code uses an array of `Shape` objects and applies `draw()` to each one using Java 5's new `for each` construct. Due to dynamic dispatch, the `draw` method of the dynamic type will be invoked, e.g. the `draw` methods for a `Circle` or a `Triangle` class. The programmer does not have to rewrite this code if new shapes are created.

Now consider reasoning about programs in the presence of dynamic dispatch. Consider the following code fragment:

```
C m1(C x) {
    x.m2();
    return x;
}
```

Assume $C.m2()$ has a specification $\{P_C\}\text{-}\{Q_C\}$. We could give $m1$ the specification $\{P_C \wedge x : C\}\text{-}\{Q_C\}$, but this prevents a call to the method $m1$ where the type of x is a subtype of C . Assume C has a single subtype D , which has the specification $\{P_D\}\text{-}\{Q_D\}$ for method $m2$. We could validate the method with the following specification:

$$\{x \neq \text{null} \wedge P_C \wedge P_D\}\text{-}\{Q_C \vee Q_D\}$$

What happens if we introduce a new class? Re-verifying $m1$ for every new class is clearly impractical for any realistic program. For example, the `equal()` method in the `java.lang.Object` is overridden approximately 300 times in the Java 1.4 libraries. We cannot reason about 300 classes for a single call; we need to reduce the complexity.

In the programmer's informal reasoning, the programmer makes a simplifying assumption that the methods of each subclass behave, at some level, the same as the supertype's method. If we are to reason *formally* about dynamic dispatch, we need to capture this assumption. In particular, we need to consider a new subtype relation that is appropriate for dynamic dispatch: *behavioural subtyping*. In the rest of this subsection we will describe different definitions of behavioural subtyping, and in §3.5 we present a new notion of behavioural subtyping.

¹In Java the ideas of subtype, and subclass are conflated. Likewise in this thesis we will not separate the concepts.




Figure 1.1: Visualization of behavioural specifications and subtyping.

Before we can define behavioural subtyping, we have to define behavioural types. America [3] considered the behaviour of a class to be specified by its methods' specifications, and the class's invariants. A method's specification is really a set of start states,² P_C , called its pre-condition and a set of end states, Q_C , called its post-condition, such that if the method starts in a state in P_C and the method terminates, then the terminating state will be in Q_C . We illustrate the idea of a specification in Figure 1.1 (A). This shows the states in P_C are mapped to the states in Q_C . The arrows are used to emphasize the mapping we are considering. A class's invariant is a property of a class instance that is true after every method call. Liskov and Wing [55] extend behavioural types to also consider a history property: a relationship between any state and a later state that the class or method must always preserve; for example a counter only goes up.

For a class to be a behavioural subtype of its superclass, it must preserve all of the parent's invariants, and its methods' specifications must be compatible. Compatibility means that the parent method's pre-condition must imply the subclass method's pre-condition and the subclass method's post-condition must imply the parent method's post-condition. The set of states A implies the set of states B, if and only if every state in A is also in B. We illustrate this in Figure 1.1 (B): P_C is the parent's pre-condition, P_D is the subclass's pre-condition, Q_C is the parent's post-condition and Q_D is the subclass's post-condition. In the diagram we represent the implications by containment of one region within another: P_C implies P_D as P_D contains P_C , likewise Q_D implies Q_C . Hence, we can see the mapping in (B) is a behavioural subtype of the mapping in (A).

To allow the internal representation to be altered, America [3] uses a transfer function from the internal values of the subclass to the parent's internal values. This function can allow the internal field names to be altered in the subtype. This function is similar to Hoare's data abstraction function [46] that maps from the abstract values of a datatype to their actual representation.

²We can interpret a formula as the set of states for which the formula is true.

Note: In this thesis we only use properties on method specifications to define behavioural subtyping. Invariants and transfer functions can be captured by the abstraction mechanism we add to the logic.

To reduce the burden of discovering and verifying compatible method specifications Dhara and Leavens [30, 31] proposed *specification inheritance*, which allows one to automatically generate specifications that ensure a class is a behavioural subtype of its parent. For each method, we must show that it satisfies all of its supertypes' specifications. For example if C is the only supertype of D , and C 's method m has specification $\{P_C\} \dashv \{Q_C\}$ then D 's specification is of the form $\{P_C \vee P_D\} \dashv \{Q_C \wedge Q_D\}$. It is easy to see this satisfies the requirements on method specifications. We illustrate specification inheritance in Figure 1.1 (C). Again we can see this mapping is a behavioural subtype of (A). This approach generates specifications that are behavioural subtypes even when the underlying specifications are not. Findler and Felleisen [37] show that generating run-time tests from these forms of specification can produce misleading error messages.

Poetzsch-Heffter and Müller [72] have proposed a different form of specification inheritance: each method must satisfy every specification of its supertypes' corresponding methods, for example if D has a method m with body \bar{s} and has a single supertype C with specification $\{P_C\} \dashv \{Q_C\}$, then we must prove both $\vdash \{P_D\} \bar{s} \{Q_D\}$ and $\vdash \{P_C\} \bar{s} \{Q_C\}$. Hence every method body requires multiple verifications.

This is more general than Dhara and Leavens's specification inheritance as it allows \bar{s} to terminate in a state not satisfying Q_C as long as it does not start in a state satisfying P_C . In Figure 1.1 (D) we illustrate a method that is compatible at the behavioural level, but does not satisfy the standard pair of implications [3, 55]. We consider D 's method to satisfy the intersection of two specifications: one maps P_C to Q and the other maps P_D to Q_D . Q_D does not have to be contained in Q_C , but Q must be contained in Q_C . This generalises Dhara and Leavens specification inheritance, and moves beyond the standard pair of implications.

In §3.5 we define a new notion *specification compatibility* that captures Poetzsch-Heffter and Müller's notion of specification inheritance without the need for multiple verifications of each method.

1.1.2 Aliasing

We have so far discussed methods for reasoning about dynamic dispatch. The second key stumbling block to reasoning about object-oriented languages, like Java, is reference aliasing. Consider the following code:

```
...
x.f = 3;
y.f = 4;
ASSERT(x.f != y.f)
```

The assertion might initially seem to hold, but in languages with references, or pointers, multiple names can refer to the same entity. This is *aliasing*. The assertion only holds if x and y reference different objects.

Aliasing makes reasoning about programs considerably harder. It breaks Hoare's original assignment rule [45]: $\{P[E/l]\} \ 1 = E; \ \{P\}$. Consider an unsound use of this rule:

$$\{x.f = 3 \wedge 4 = 4 \wedge x = y\} \ y.f = 4; \ \{x.f = 3 \wedge y.f = 4 \wedge x = y\}$$

This is an instance of the assignment rule because $x.f = 4 \wedge y.f = 4 \wedge x = y[4/y.f]$ equals $x.f = 3 \wedge 4 = 4 \wedge x = y$. For any sensible semantics the post-condition is *false* as it implies $3 = 4$.

One solution to the problem of assignment with aliasing is Morris’s component substitution [59]. He modifies substitution to operate on components.

$$x.a[E/y.a] \stackrel{\text{def}}{=} \text{if } x = y \text{ then } E \text{ else } x.a$$

Every time we make a possibly aliased assignment, we insert a conditional test for the cases where it was aliased and where it was not. Alternatively one can see the components as an array, for example seeing $x.a$ as an array access $a[x]$. This then allows the Hoare logic rules for array assignment to be used rather than introducing the conditional test. These methods explicitly demonstrate the global nature of the store. Assignment to a particular component can potentially affect any other assertion about that component.

In Hoare logic without aliased state one can write local specifications, and then infer global specifications using the Hoare logic rule of invariance [4], or the specification logic rule of constancy [77]:

$$\frac{\{P\}\bar{s}\{Q\}}{\{P \wedge R\}\bar{s}\{Q \wedge R\}} \text{ provided } \bar{s} \text{ does not modify any free variables of } R$$

Without aliasing the side-condition can be verified by a simple syntactic check of \bar{s} . When aliasing is present the check is no longer simply syntactic and more complex conditions about aliasing and interference must be used.

This problem impacts on behavioural types. The simplest solution is to require methods to be annotated with the details of how they alter the state. For example [43] the annotation “`modifies this.f1, x.f2;`” specifies that the method only alters the receiver’s `f1` field, and `x`’s `f2` field, no other state is modified. However, in the presence of behavioural subtyping these clauses are too precise: subtypes usually modify more state than their parents modify (they are, after all, extensions), but for soundness it can only modify the same or less. This is known as the *extended state problem*.

This problem was addressed by Leino [54]. He proposed the use of *data groups*, which allows the `modifies` clauses to be given in terms of abstract groups of fields. The membership of a group is then specified by the class. Hence it is possible for a subclass to modify more state than its parent, but it cannot modify data from any more data groups.

Whilst data groups solve some problems, they cannot directly deal with recursive data structures as the shape of the heap affects the values that are modified. Consider a method that modifies a list. What is the `modifies` clause? For a three element list it might be `x.tl, x.tl.tl, x.tl.tl.tl`, but clearly this approach does not scale. (One might imagine providing a regular expression for the `modifies` clause, but this is perhaps too complicated.)

To extend `modifies` clauses to recursive data structures, some form of ownership has been proposed [32, 60, 61, 79]. An object or field is said to own other objects or fields. If a field is specified as modified, then that field and all of the other state it owns could be modified. We will briefly describe some type systems that enforce ownership properties but refer the interested reader to Clarke’s thesis [22] for a comprehensive survey of different type systems for reasoning about ownership and aliasing.

Ownership types [15, 22, 23, 24] allow the programmer to enforce one object’s ownership of another. In Figure 1.2, we demonstrate a typical ownership diagram: the two node objects are owned by the list object, and the list and data objects are owned by a distinguished “world” object. The objects owned by a particular object are viewed as its representation: ownership types prevent external references into this representation, and hence enforce encapsulation of its representation. Universes [32] take a similar approach to ownership types. They only prevent the writing to the representation rather than any access. This allows universes to handle list

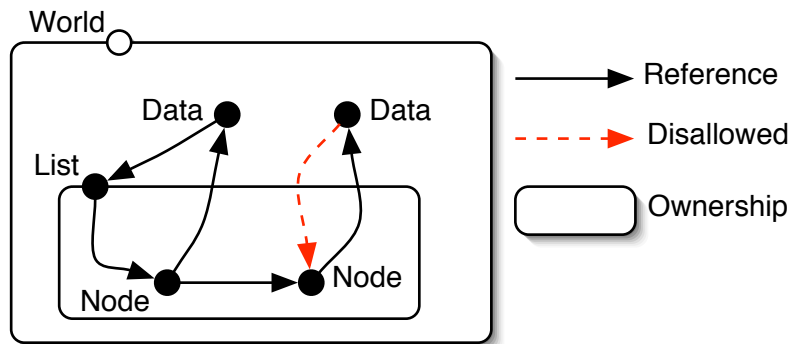


Figure 1.2: Example of ownership types [23]

iterators that are difficult to handle in ownership types. However, universes are less expressive in representing nested ownership.

So far we have reviewed some of the difficulties with reasoning about aliased state and described briefly some solutions. Now we will discuss *local reasoning*, which comes from observations of Burstall [20]. He noted that if you explicitly consider disjointness in the logic then you can simplify your reasoning. He presented a series of examples on lists and trees, and presented a predicate to assert two lists were acyclic and had disjoint domains. Bornat [12] has shown how to extend this style of reasoning to a more general recursive datatype setting.

Based on Burstall’s observations about disjointness, Reynolds [76], and Ishtiaq and O’Hearn [48] took a different approach. Rather than adding a predicate for disjointness, they added a new logical connective for disjointness: a second form of conjunction $P * Q$, which means both P and Q hold, but in disjoint parts of the heap. Ishtiaq and O’Hearn [48] showed this new connective was the multiplicative conjunction from the “Logic of Bunched Implications” by O’Hearn and Pym [64].

This programming logic is now known as *separation logic*. It has a subtly different semantics from standard Hoare logic as the triples are *tight specifications*: that is for $\{P\}\bar{s}\{Q\}$ to hold, P must describe all the heap \bar{s} needs during execution. This leads to the *frame rule*:

$$\frac{\{P\}\bar{s}\{Q\}}{\{P * R\}\bar{s}\{Q * R\}} \text{ provided } \bar{s} \text{ does not modify the free variables of } R.^3$$

Although this rule has a side-condition, it is simple and syntactic as the issues with aliasing are dealt with using the $*$ connective. There is no need for type systems to check uniqueness or ownership. Additionally, if one precludes global variables then methods will always have an empty modifies set, removing the need for `modifies` clauses: the specifications contain all the information about the state that is modified.

The frame rule allows *local reasoning*: the code can be specified based solely on the state it uses. Clients can then extend the specification to the global setting that accounts for all of their additional state. This kind of reasoning is essential for open programs: programs where the client cannot verify their requirements against the implementation, but must simply rely on the specification provided.

³Bornat [13] has shown how to treat the framing of variables with $*$, hence removing the side-condition.

1.2 Content and contribution

We will now provide an overview of the thesis's content and highlight its contribution.

Chapter 2 We begin the thesis by identifying Middleweight Java (MJ), a core imperative fragment of Java. We formalize the type system for this fragment and give a novel operational semantics, which we prove is type sound.

Contribution: The key contributions of this chapter are the development of the semantics and type system and the proof of soundness.

Chapter 3 This chapter extends local reasoning to the Java programming language and serves as an introduction to separation logic. A rule is given for dynamic dispatch based on a new formulation of behavioural subtyping named *specification inheritance*. It describes how this formulation is more flexible than standard behavioural subtyping. We illustrate the logic's use with a simple tree copy program and prove the logic is sound with respect to the semantics in Chapter 2. We conclude by showing that the logic allows local reasoning, but does not provide modular reasoning in the sense of encapsulation or inheritance.

Contribution: This chapter has two key contributions: a separation logic for Java and a new definition of behavioural subtyping named *specification compatibility*.

Chapter 4 In this chapter we address the issues of encapsulation and abstract datatypes, by developing a novel concept of an *abstract predicate*. Abstract predicates are the logical equivalent of an abstract datatype. We show the encapsulation they provide allows *ownership transfer* [66]: the ability for state to be safely transferred between the client and the module, or class. In particular, we show how to reason about the object pool pattern [42]. We prove the extension is sound and conclude the chapter with a discussion of related work, and a demonstration of how to reason about a C-style memory manager.

Contribution: The key contribution of this chapter is the concept of an *abstract predicate*. Additionally, we demonstrate that abstract predicates provide more flexible modularity than the hypothetical frame rule [66].

Chapter 5 This chapter shows that the encapsulation provided by abstract predicates in separation logic is too strict to allow many common programming idioms, including list iterators. We present a new model of separation logic that allows read sharing. In the standard presentation of separation logic, only a single datatype (or, in the concurrent setting, a thread) can have access to a location in memory. In this chapter, we extend separation logic to allow multiple datatypes to own read references to the same location and demonstrate this work with a list iterator example. However, the initial presentation does not allow the disjointness of datastructures to be properly expressed. We present a model of named permissions and show that this model is capable of representing disjointness in datatypes while allowing read sharing.

Contribution: This chapter has two key contributions: demonstrating the use of read-only permissions with abstract predicates and presenting a model of separation logic that can express both read sharing and disjointness.

Chapter 6 This chapter highlights the issues of reasoning about inheritance in separation logic. We extend the concept of an abstract predicate to an *abstract predicate family*. The extension allows a predicate to have one of many different definitions based on its argument's type. We demonstrate the utility of abstract predicate families with two examples: Cell/Recell [1] a simple

example of subtyping; and the visitor design pattern [40]. We prove this extension is sound, and conclude by discussing future work.

Contribution: The key contribution of this chapter is the *abstract predicate family* and the demonstration of its utility when reasoning about object-oriented programs.

Chapter 7 This chapter concludes the thesis.

1.3 Related work

To conclude this chapter, we discuss other logics for reasoning about object-oriented languages, and compare them to the logic developed in this thesis. We will not consider object-based languages [1] and associated logics [2] for two reasons: firstly, we want our verification methods to be close to programmers, who generally use class-based languages; and secondly, the class hierarchy actually imposes a structure which helps with specification. Reus [74] has given a denotational comparison between specifications for class- and object-based languages and shows that the domains are simpler for class-based languages.

There are several Hoare-like logics embedded into theorem provers. Von Oheimb and Nipkow [84] embed a logic into Isabelle. The assertion language is part of Isabelle, and the assertions use an explicit state parameter. They do not support any form of encapsulation, inheritance or frame-like property. Poetzsch-Heffter and Müller [71] present a different logic for reasoning about Java-like languages. Again, they use an explicit store in the logic. However, they give axioms for manipulating the store [72]. Müller’s thesis [60] addresses the issues of modularity by adding abstract fields with dependency information and universe types [32]. He provides `modifies` clauses to provide a frame property. The `modifies` clauses are given in terms of the abstract fields and the fields they depend on. The universe types [32] allow object invariants to depend on other objects.

The syntax of Müller’s assertion language is closely related to Leavens *et al.*’s JML [53], which attempts to unify many approaches to tool support by giving them a common syntax. There are now a reasonable collection of tools for JML [19]: some compile dynamic checks into the code; some provide test cases; and some do static verification. JML does not provide any help with dealing with aliasing. Müller *et al.* [61] apply the ideas from Müller’s thesis on integrating universes to JML to provide a modular frame property, but this is not part of the language yet. Many of the tools use their own methods, and syntax, for dealing with aliasing. One notable example of a JML tool is ESC/Java [29]. Interestingly it is not designed to be sound or complete, but simply to find common errors.

In a similar vein to JML, `Spec#` of Barnett, Leino, and Schulte [6] allows assertions to be statically verified. Any assertions that cannot be statically verified will be converted into run-time checks. Many difficult features, such as invariants, are modelled using additional logical fields: fields only required for the proof. These allow the invariant to be broken while the field is false. The invariant is, in effect, `this.inv ⇒ I`: while the field is true, *I* must hold, otherwise it may or may not hold. More complex schemes can be used where the invariant is a stack of invariants corresponding to each of the superclasses’ invariants [5]. This work encodes the ideas of object ownership into fields that represent object ownership. Recently, there has been a proposal to extend this methodology to a concurrent setting [49].

A different, automated method has been taken by Fähndrich and DeLine [36]. They revisit some work on compiler detected errors known as `typestates` [81], and have developed an object-oriented extension. They use `typestates` as extended invariants: an invariant that has many different, visible states. `Typestates` can be viewed as abstract predicates on an object’s state, in a very similar way to the work developed in Chapters 4 and 6. Fähndrich and DeLine use

typestates to check for errors in calling protocols. Their logic is considerably less expressive than the logic developed in this thesis, but their aim was automatic checking, not correctness proofs. They use the logic for an automatic protocol checker Fugue [28].

There have been a few works based on the component substitution ideas of Morris [59]. In particular, De Boer and Pierik [26, 27, 69, 70] have extended this to an object-oriented setting and introduce a form of substitution for dynamic allocation of objects, $[_{new/u}]$. Recently Berger, Honda, and Yoshida [7] have also developed a logic based on Morris's work [59]. They add modal operators about locations, which are used to provide structural reasoning about observations. The work differs fundamentally from separation logic as it does not have explicit concepts of resource. Their logic can, however, deal with higher order functions.

Finally, Middelkoop's recent Master's thesis [56, 57] also adapts separation logic to reason about Java. However, his work takes many different design decisions to this thesis, and does not attempt to solve any of the issues surrounding modularity.

The logics outlined above either do not provide support for modularity and framing, or if they do, then the extensions are separate from the logic. For example, the logical rules might use information from the type systems for ownership or uniqueness. However, this tight integration can make it hard to express some dependencies. Consider the following code:

```
void m() {
  if(x == null) {
    <code modifying (1)>
  } else {
    <code modifying (2)>
  }
}
```

Given a standard `modifies` clause we cannot express the dependency between the state and what is modified: the method modifies the union of (1) and (2). To then get an accurate specification we would then have to add assertions that say (1) is unchanged when `x` is not `null`, and (2) is not changed otherwise. However, we have lost much of the advantage of `modifies` clauses as the specifications must provide additional framing properties.

By using separation logic we automatically get notions of framing and ownership *in the logic*. This provides a single unified logical view of these concepts rather than having to extend our language with many different ideas. We believe this leads to a clearer logic and to simpler proof techniques.

2

Middleweight Java

In this chapter we define Middleweight Java, MJ, our proposal for an imperative core calculus for Java. MJ contains the core object-oriented and imperative features of Java, in particular block structured local state, mutable fields, inheritance and method overriding. MJ does not contain interfaces, super calls, static fields, reflection, concurrency and inner classes.

It is important to note that MJ is an entirely valid subset of Java: all MJ programs are literally executable Java programs. An alternative would be to allow extra-language features; for example, Classic Java uses annotations and let bindings which are not valid Java syntax in the operational semantics [38]. This is important as we want programmers to reason directly about this code, without having to consider a separate language.

This chapter is structured as follows. We present MJ's syntax in §2.1, its type system in §2.2 and its single-step operational semantics in §2.3. In §2.4 we provide definitions for typing configurations, and in §2.5 we present a proof of correctness of the type system, and conclude in §2.6 with a discussion of related Java semantics.

2.1 Syntax

The syntax for MJ programs is given in Figure 2.1. An MJ program is a collection of class definitions plus a sequence of statements, \bar{s} , to be evaluated. This sequence corresponds to the body of the `main` method in a Java program [41, §12.1.4].

For example, Figure 2.2 presents some typical MJ class definitions. This code defines two classes: `Cell` which is a subclass of the `Object` class, and `Recell` which is a subclass of `Cell`. The `Cell` class has one field, `contents`. The constructor simply assigns to the field the value of the parameter. The `Cell` class defines a method `set`, which sets the field to a given value.

`Recell` objects inherit the `contents` field from their superclass, and also have another field, `undo`. The constructor first calls the superclass's constructor (which will assign the `contents` field) and then sets the `undo` field to the `null` object reference. The `Recell` class definition overrides the `set` method of its superclass.

As with Featherweight Java [47], FJ, we insist on a certain amount of syntactic regularity in class definitions, although this is really just to make the definitions compact. We insist that all class definitions (1) include a supertype (we assume a distinguished class `Object`); (2) include a constructor (for simplicity we only allow a single constructor per class); (3) have a call to `super` as the first statement in a constructor method; (4) have a `return` as the last statement of every method definition except for `void` methods and constructor definitions (this

Program

$$\pi ::= cd_1 \dots cd_n; \bar{s}$$

Class definition

$$cd ::= \text{class } C \text{ extends } C\{\overline{fd} \text{ } \overline{cnd} \text{ } \overline{md}\}$$

Field definition

$$fd ::= C f;$$

Constructor definition

$$cnd ::= C(\overline{C} \overline{x})\{\text{super}(\overline{e}); \overline{s}\}$$

Method definition

$$md ::= \tau m(\overline{C} \overline{x})\{\overline{s}\}$$

Return type

$$\tau ::= C \mid \text{void}$$

Expression

$e ::= x$	Variable
null	Null
$e.f$	Field access
$(C)e$	Cast
pe	Promotable expression

Promotable expression

$pe ::= e.m(\overline{e})$	Method invocation
new $C(\overline{e})$	Object creation

Statement

$s ::= ;$	No-op
$pe;$	Promoted expression
if $(e == e)\{\overline{s}\}$ else $\{\overline{s}\}$	Conditional
$e.f = e;$	Field assignment
$C x;$	Local variable declaration
$x = e;$	Variable assignment
return $e;$	Return
$\{\overline{s}\}$	Block

Figure 2.1: Syntax for MJ programs

<pre>class Cell extends Object{ Object contents; Cell (Object start){ super(); this.contents = start; } void set(Object update){ this.contents = update; } }</pre>	<pre>class Recell extends Cell{ Object undo; Recell (Object start){ super(start); this.undo = null; } void set(Object update){ this.undo = this.contents; this.contents = update; } }}</pre>
--	--

Figure 2.2: Example MJ code: Cell and Recell

constraint is enforced by the type system); and (5) write out field accesses explicitly, even when the receiver is `this`.

In what follows, again for compactness, we assume that MJ programs are well-formed, that is we insist that (1) they do not have duplicate definitions for classes, fields and methods (for simplicity we do not allow overloaded methods—as overloading is determined statically, overloaded methods can be simulated faithfully in MJ); (2) fields are not redefined in subclasses (we do not allow shadowing of fields); and (3) there are no cyclic class definitions (for example `A extends B` and `B extends A`). We do not formalise these straightforward conditions.

Returning to the definition in Figure 2.1: a class definition contains a collection of field and method definitions, and a single constructor definition. A field is defined by a type and a name. Methods are defined as a return type, a method name, an ordered list of arguments, where an argument is a variable type and name, and a body. A constructor is defined by the class name, an ordered list of arguments and a body. There are a number of well-formedness conditions for a collection of class definitions which are formalised in the next section.

The rest of Figure 2.1 defines MJ expressions and statements. We assume a number of metavariables: f ranges over field names, m over method names, and x over variables. We assume that the set of variables includes a distinguished variable, `this`, which is not permitted to occur as the name of an argument to a method or on the left of an assignment. In what follows, we shall find it convenient to write \bar{e} to denote the possibly empty sequence e_1, \dots, e_n (and similarly for \bar{C} , \bar{x} , etc.). We write \bar{s} to denote the sequence $s_1 \dots s_n$ with no commas (and similarly for \bar{fd} and \bar{md}). We abbreviate operations on pairs of sequences in the obvious way, thus for example we write $\bar{C} \bar{x}$ for the sequence $C_1 x_1, \dots, C_n x_n$ where n is the length of \bar{C} and \bar{x} .

The reader will note MJ has two classes of expressions: the class of ‘promotable expressions’, pe , defines expressions that can be promoted to statements by postfixing a semicolon ‘;’; the other, e , defines the other expression forms. This slightly awkward division is imposed upon us by our desire to make MJ a valid fragment of Java. Java [41, §14.8] only allows particular expression forms to be promoted to statements by postfixing a semicolon. This leads to some rather strange syntactic surprises: for example, `x++`; is a valid statement, but `(x++)`; is not!

MJ includes the essential imperative features of Java. Thus we have fields, which can be both accessed and assigned to, as well as variables, which can be locally declared and assigned. As with Java, MJ supports block-structure; consider the following valid MJ code fragment.

```
if(var1 == var2) { ; }
else {
  Object temp;
  temp = var1;
  var1 = var2;
  var2 = temp;
}
```

This code compares two variables, `var1` and `var2`. If they are not equal then it creates a new locally scoped variable, `temp`, and uses this to swap the values of the two variables. At the end of the block, `temp` will no longer be in scope and will be removed from the variable stack.

2.2 Types

As with FJ, for simplicity, MJ does not have primitive types, sometimes called base types. Thus all well-typed expressions are of a class type, C . All well-typed statements are of type `void`, except for the statement form `return e`; which has the type of e , i.e. a class type. We use τ to range over valid statement types. The type of a method is a pair, written $\bar{C} \rightarrow \tau$, where \bar{C} is a sequence of argument class types and τ is the return type (if a method does not return anything, its return type is `void`). We use μ to range over method types.

Expression types

C a valid class name, including a distinguished class `Object`

Statement types

$\tau ::= \text{void} \mid C$

Method types

$\mu ::= C_1, \dots, C_n \rightarrow \tau$

Java, and MJ, class definitions contain both typing information and code. This typing information is extracted and used to typecheck the code. Thus before presenting the typechecking rules we need to specify how this typing information is induced by MJ code.

This typing information consists of two parts: The subclassing relation, and a class table (which stores the types associated with classes). First let us consider the subclassing relation. Recall that in MJ we restrict class declarations so that they must give the name of class that they are extending, even if this class is the `Object` class. We forbid the `Object` class being defined in the program.

A well-formed program, π , then induces an immediate subclassing relation, which we write \prec_1 .¹ We define the *subclassing relation*, \prec , as the reflexive, transitive closure of this immediate subclassing relation. This can be defined formally as follows.

TR-IMMEDIATE	$\frac{\text{class } C_1 \text{ extends } C_2\{\dots\} \in \pi}{C_1 \prec_1 C_2}$
TR-TRANSITIVE	$\frac{C_1 \prec C_2 \quad C_2 \prec C_3}{C_1 \prec C_3}$
TR-EXTENDS	$\frac{C_1 \prec_1 C_2}{C_1 \prec C_2}$
TR-REFLEXIVE	$C \prec C$

A well-formed program π also induces a class table, δ_π . As the program is fixed, we will simply write this as δ . A class table, δ , is actually a triple, $(\delta_M, \delta_C, \delta_F)$, which provides typing information about the methods, constructors, and fields, respectively. δ_M is a partial map from a class name to a partial map from a method name to that method's type. Thus $\delta_M(C)(m)$ is intended to denote the type of method m in class C . δ_C is a partial map from a class name to the type of that class's constructor's arguments. δ_F is a partial map from a class name to a map from a field name to a type. Thus $\delta_F(C)(f)$ is intended to denote the type of f in class C . The details of how a well-formed program π induces a class table δ are given below.

Notation: δ_M^Δ and δ_F^Δ define the types of the methods and fields that each class defines, respectively. δ_M and δ_F are the types of the methods and fields a class defines and inherits, respectively.

¹This relation is also a partial function that uniquely defines all non-`Object` classes' immediate superclass.

Immediate definitions

$$\begin{aligned}
\delta_C(C) &\stackrel{\text{def}}{=} \overline{C} && \text{where } \text{cnd} = C(\overline{C}\overline{x})\{\dots\} \\
\delta_M^\Delta(C) &\stackrel{\text{def}}{=} \{m \mapsto (\overline{C} \rightarrow C'') \mid C''m(\overline{C}\overline{x})\{\dots\} \in \overline{md}\} \\
\delta_F^\Delta(C) &\stackrel{\text{def}}{=} \{f \mapsto C'' \mid C''f; \in \overline{fd}\}
\end{aligned}$$

where class C extends $C'\{\overline{fd} \text{ cnd } \overline{md}\} \in \pi$

Inherited definitions

$$\begin{aligned}
\delta_M(C)(m) &\stackrel{\text{def}}{=} \begin{cases} \delta_M^\Delta(C)(m) & \delta_M^\Delta(C)(m) \text{ defined} \\ \delta_M(C')(m) & \text{otherwise} \end{cases} \\
\delta_F(C)(f) &\stackrel{\text{def}}{=} \begin{cases} \delta_F^\Delta(C)(m) & \delta_F^\Delta(C)(m) \text{ defined} \\ \delta_F(C')(m) & \text{otherwise} \end{cases}
\end{aligned}$$

where $C \prec_1 C'$

Note: The constructors are not inherited.

Until now we have assumed that the class definitions are well-formed. Now let us define formally well-formedness of class definitions. First we will find it useful to define when a type (either a method type or statement type) is well-formed with respect to a class table. This is written as a judgement $\delta \vdash \mu \text{ ok}$ and $\delta \vdash \tau \text{ ok}$ which mean that μ and τ are a valid type given the class table δ , respectively. (We define $\text{dom}(\delta)$ to be the domain of δ_F , δ_M or δ_C , which are all equal.)

$$\begin{array}{l}
\text{T-C}_{\text{TYPE}} \quad \frac{C \in \text{dom}(\delta)}{\delta \vdash C \text{ ok}} \\
\text{T-V}_{\text{TYPE}} \quad \delta \vdash \text{void ok} \\
\text{T-M}_{\text{TYPE}} \quad \frac{\delta \vdash C_1 \text{ ok} \quad \dots \quad \delta \vdash C_n \text{ ok} \quad \delta \vdash \tau \text{ ok}}{\delta \vdash C_1, \dots, C_n \rightarrow \tau \text{ ok}}
\end{array}$$

Now we define formally the judgement for well-formed class tables, which is written $\vdash \delta \text{ ok}$. This essentially checks two things: firstly that all types are valid given the classes defined in the class table, and secondly that if a method is overridden then it is done so at the *same* type. The judgements are given as follows

T-FIELDSOK	$\frac{\forall f \in \text{dom}(\delta_F^\Delta(C)). (\delta \vdash \delta_F^\Delta(C)(f) \text{ ok } C \prec_1 C' \delta_F(C')(f) \text{ undefined})}{\delta \vdash_C \delta_F \text{ ok}}$
T-CONSOK	$\frac{\forall C' \in \delta_C(C). \delta \vdash C' \text{ ok}}{\delta \vdash_C \delta_C \text{ ok}}$
T-METHOK	$\frac{\delta \vdash \mu \text{ ok } \delta_m(C)(m) = \mu \quad C \prec_1 C' \quad m \in \text{dom}(\delta_m(C')) \Rightarrow \delta_m(C')(m) = \mu}{\delta \vdash C.m \text{ ok}}$
T-METHSOK	$\frac{\forall m \in \text{dom}(\delta_M(C)). \delta \vdash C.m \text{ ok}}{\delta \vdash_C \delta_M \text{ ok}}$
T-CLASSOK	$\frac{\forall C \in \text{dom}(\delta). (\delta \vdash_C \delta_F \text{ ok } \delta \vdash_C \delta_M \text{ ok } \delta \vdash_C \delta_C \text{ ok } C \prec_1 C' \delta \vdash C' \text{ ok})}{\vdash \delta \text{ ok}}$

A typing environment, γ , is a map from program variables to expression types. We write $\gamma, x: C$ to denote the map γ disjointly extended so that x maps to C . We write $\delta \vdash \gamma \text{ ok}$ to mean that the types in the image of the typing environment are well-formed with respect to the class table, δ .

We can now define the typing rules for expressions and promotable expressions. A typing judgement for a given MJ expression, e , is written $\delta; \gamma \vdash e: C$ where δ is a class table and γ is a typing environment. These rules are given in Figure 2.3. Our treatment of casting is the same as in FJ [47]—thus we include a case for ‘stupid’ casting, where the target class is completely unrelated to the subject class. This is needed to handle the case where an expression without stupid casts may reduce to one containing a stupid cast. Consider the following expression, taken from [47], where classes `A` and `B` are both defined to be subclasses of the `Object` class, but are otherwise unrelated.

(A) `(Object)new B()`

According to the Java Language Specification [41, §15.16], this is well-typed, but consider its operational behaviour: a `B` object is created and is dynamically upcast to `Object` (which has no dynamic effect). At this point we wish to downcast the `B` object to `A`—a ‘stupid’ cast! Thus if we wish to maintain a subject reduction theorem (that the result of a single step reduction of a well-typed program is also a well-typed program) we need to include the `TE-STUPIDCAST` rule. For the same reason, we also require two rules for typing an `if` statement. Java [41, §15.21.2] enforces statically that when comparing objects, one object should be a subclass of the other. However this is not preserved by the dynamics. Consider again the unrelated classes `A` and `B`. The following code fragment is well-typed but at runtime will end up comparing objects of unrelated classes.

`if ((Object)new A() == (Object)new B()) { ... } else { ...};`

Definition 2.2.1. A *valid* Java/MJ program is one that does not have occurrences of `TE-STUPIDCAST` or `TS-STUPIDIF` in its typing derivation.

A typing judgement for a sequence of statements, \bar{s} , is written $\delta; \gamma \vdash \bar{s}: \tau$ where δ is a class table and γ is a typing environment. We begin by presenting rules for a singleton sequence. As we noted earlier, statements in Java can have a non-`void` type (see rule `TS-RETURN`). The rules for forming these typing judgements are given in Figure 2.4.

TE-VAR	$\frac{\delta \vdash \gamma \text{ ok} \quad \vdash \delta \text{ ok}}{\delta; \gamma, x: C \vdash x: C}$	
TE-NULL	$\frac{\delta \vdash C \text{ ok} \quad \delta \vdash \gamma \text{ ok} \quad \vdash \delta \text{ ok}}{\delta; \gamma \vdash \text{null}: C}$	
TE-FIELDACCESS	$\frac{\delta; \gamma \vdash e: C_2 \quad \delta_{\text{F}}(C_2)(f) = C_1}{\delta; \gamma \vdash e.f: C_1}$	
TE-UPCAST	$\frac{\delta; \gamma \vdash e: C_2 \quad \delta \vdash C_1 \text{ ok}}{\delta; \gamma \vdash (C_1)e: C_1}$	provided $C_2 \prec C_1$
TE-DOWNCAST	$\frac{\delta; \gamma \vdash e: C_2 \quad \delta \vdash C_1 \text{ ok}}{\delta; \gamma \vdash (C_1)e: C_1}$	provided $C_1 \prec C_2$
TE-STUPIDCAST	$\frac{\delta; \gamma \vdash e: C_2 \quad \delta \vdash C_1 \text{ ok}}{\delta; \gamma \vdash (C_1)e: C_1}$	provided $C_2 \not\prec C_1 \wedge C_1 \not\prec C_2$
TE-METHOD	$\frac{\delta; \gamma \vdash e: C' \quad \delta; \gamma \vdash e_1: C_1 \quad \dots \quad \delta; \gamma \vdash e_n: C_n}{\delta; \gamma \vdash e.m(e_1, \dots, e_n): C}$	provided $\delta_{\text{M}}(C')(m) = C'_1, \dots, C'_n \rightarrow C$ $\wedge C_1 \prec C'_1 \dots C_n \prec C'_n$
TE-NEW	$\frac{\delta; \gamma \vdash e_1: C'_1 \quad \dots \quad \delta; \gamma \vdash e_n: C'_n}{\delta; \gamma \vdash \text{new } C(e_1, \dots, e_n): C}$	provided $\delta_{\text{C}}(C) = C_1, \dots, C_n$ $\wedge C'_1 \prec C_1 \dots C'_n \prec C_n$

Figure 2.3: Typing rules for expressions

Java allows variable declarations to occur at any point in a block. To handle this we introduce two typing rules for sequencing: TS-INTRO for the case where the first statement in the sequence is a variable declaration, and TS-SEQ for the other cases. An alternative approach would be to force each statement to return the new typing environment. We feel that our presentation is simpler.

Finally in Figure 2.4 we define the typing of the `super` call in the constructor of a class. A call to the empty constructor in a class that directly extends `Object` is always valid, and otherwise it must be a valid call to the constructor of the parent class. The expressions evaluated before this call are not allowed to reference `this` [41, §8.8.5.1]. This is because it is not a validly initialised object until the parent's constructor has been called.

Note: We could remove the subtyping from the individual rules and add a single subtype rule. However this would allow implicit casting where Java does not.

Before we give the typing rules involving methods we first define some useful auxiliary functions for accessing the bodies of constructors and methods.

TS-NOOP	$\frac{\vdash \delta \text{ ok} \quad \delta \vdash \gamma \text{ ok}}{\delta; \gamma \vdash ; : \text{void}}$	
TS-PE	$\frac{\delta; \gamma \vdash pe : \tau}{\delta; \gamma \vdash pe; : \text{void}}$	
TS-IF	$\frac{\delta; \gamma \vdash \overline{s_1} : \text{void} \quad \delta; \gamma \vdash \overline{s_2} : \text{void} \quad \delta; \gamma \vdash e_1 : C' \quad \delta; \gamma \vdash e_2 : C''}{\delta; \gamma \vdash \text{if } (e_1 == e_2) \{ \overline{s_1} \} \text{ else } \{ \overline{s_2} \} : \text{void}}$	provided $C'' \prec C' \vee C' \prec C''$
TS-STUPIDIF	$\frac{\delta; \gamma \vdash \overline{s_1} : \text{void} \quad \delta; \gamma \vdash \overline{s_2} : \text{void} \quad \delta; \gamma \vdash e_1 : C' \quad \delta; \gamma \vdash e_2 : C''}{\delta; \gamma \vdash \text{if } (e_1 == e_2) \{ \overline{s_1} \} \text{ else } \{ \overline{s_2} \} : \text{void}}$	provided $C'' \not\prec C' \wedge C' \not\prec C''$
TS-FIELDWRITE	$\frac{\delta; \gamma \vdash e_1 : C_1 \quad \delta; \gamma \vdash e_2 : C_2 \quad \delta_F(C_1)(f) = C_3}{\delta; \gamma \vdash e_1.f = e_2; : \text{void}}$	provided $C_2 \prec C_3$
TS-VARWRITE	$\frac{\delta; \gamma \vdash x : C \quad \delta; \gamma \vdash e : C'}{\delta; \gamma \vdash x = e; : \text{void}}$	provided $C' \prec C \wedge x \neq \text{this}$
TS-RETURN	$\frac{\delta; \gamma \vdash e : C}{\delta; \gamma \vdash \text{return } e; : C}$	
TS-BLOCK	$\frac{\delta; \gamma \vdash s_1 \dots s_n : \text{void}}{\delta; \gamma \vdash \{ s_1 \dots s_n \} : \text{void}}$	
TS-INTRO	$\frac{\delta; \gamma, x : C \vdash s_1 \dots s_n : \tau}{\delta; \gamma \vdash C x; s_1 \dots s_n : \tau}$	provided $x \notin \text{dom}(\gamma)$
TS-SEQ	$\frac{\delta; \gamma \vdash s_1 : \text{void} \quad \delta; \gamma \vdash s_2 \dots s_n : \tau}{\delta; \gamma \vdash s_1 s_2 \dots s_n : \tau}$	provided $s_1 \neq C x$;
T-COBJECT	$\frac{C \prec_1 \text{Object}}{\delta; \gamma, \text{this} : C \vdash \text{super}() : \text{void}}$	
T-CSUPER	$\frac{\delta; \gamma' \vdash \overline{e} : \overline{C'}}{\delta; \gamma \vdash \text{super}(\overline{e}); : \text{void}}$	provided $\gamma(\text{this}) = C \wedge \gamma = \gamma' \uplus \{ \text{this} : C \} \wedge \delta_C(C') = \overline{C}$ $\wedge C \prec_1 C' \wedge \overline{C'} \prec \overline{C}$

Figure 2.4: Typing rules for statements

T-MDEFN	$\frac{\delta; \{\text{this} : C, \bar{x} : \bar{C}\} \vdash \bar{s} : \tau}{\delta \vdash \text{mbody}(C, m) \text{ ok}}$ <p style="text-align: center; margin: 0;">provided $\text{mbody}(C, m) = (\bar{x}, \bar{s}) \wedge \delta_{\text{M}}(C)(m) = (\bar{C} \rightarrow \tau)$</p>
T-MBODYS	$\frac{\forall m \in \delta_{\text{M}}(C). \delta \vdash \text{mbody}(C, m) \text{ ok}}{\delta \vdash C \text{ mok}}$
T-CDEFN	$\frac{\delta; \text{this} : C, \bar{x} : \bar{C} \vdash \text{super}(\bar{e}); : \text{void} \quad \delta; \text{this} : C, \bar{x} : \bar{C} \vdash \bar{s} : \text{void}}{\delta \vdash C \text{ cok}}$ <p style="text-align: center; margin: 0;">provided $\delta_{\text{C}}(C) = \bar{C} \wedge \text{cnbody}(C) = (\bar{x}, \text{super}(\bar{e}); \bar{s})$</p>
T-PROGDEF	$\frac{\forall C \in \delta. (\delta \vdash C \text{ cok} \quad \delta \vdash C \text{ mok})}{\delta \vdash \pi \text{ ok}}$

Figure 2.5: Typing MJ programs

Method Body	$\text{mbody}(C, m) \stackrel{\text{def}}{=} \begin{cases} (\bar{x}, \bar{s}) & md_i = C''m(\bar{C}\bar{x})\{\bar{s}\} \\ \text{mbody}(C', m) & m \notin md_1 \dots md_n \end{cases}$ <p style="text-align: center; margin: 0;">where class C extends $C'\{\bar{f}\bar{d} \text{ cnd } md_1 \dots md_n\} \in \pi$</p>
Constructor Body	$\text{cnbody}(C) \stackrel{\text{def}}{=} (\bar{x}, \bar{s})$ <p style="text-align: center; margin: 0;">where class C extends $C'\{\bar{f}\bar{d} C(\bar{C}''\bar{x})\{\bar{s}\} \bar{m}\bar{d}\} \in \pi$</p>

We can now formalise the notion of a program being well-typed with respect to its class table. This is denoted by the judgement $\delta \vdash \pi \text{ ok}$. Informally this involves checking that each method body and constructor body is well-typed and that the type deduced matches that contained in δ . We introduce two new judgement forms: $\delta \vdash C \text{ mok}$ denotes that the methods of class C are well-typed, and $\delta \vdash C \text{ cok}$ denotes that the constructor method of class C is well-typed. The rules for forming these judgements are given in Figure 2.5.

2.3 Operational semantics

We define the operational semantics of MJ in terms of transitions between *configurations*. A configuration is a four-tuple, containing the following information:

1. **Heap:** A pair of finite partial functions, the first maps object identifiers, *oids*, to class names, and the second maps pairs of oids and fields to values (oids and null).²
2. **Variable Stack:** This essentially maps variable names to *oids*. To handle static block-structured scoping it is implemented as a list of lists of partial functions from variables to values. (This is explained in more detail later.) We use \circ to denote stack cons.
3. **Term:** The program (a closed frame) to be evaluated.
4. **Frame stack:** This is essentially the program context in which the term is currently being evaluated.

This is defined formally in Figure 2.6. CF is a closed frame (i.e. with no hole) and OF is an open frame (i.e. requires an expression to be substituted in for the hole). Additionally we have

²This is different to the original design of MJ, in [9]. The changes are to streamline the presentation in the later chapters, see §3.2.

<p>Configuration $config ::= (H, VS, CF, FS)$ $\quad \quad \quad (H, VS, EF, FS)$</p> <p>Frame stack $FS ::= F \circ FS \mid []$</p> <p>Frame $F ::= CF \mid OF$</p> <p>Closed frame $CF ::= \bar{s} \mid \text{return } e; \mid \{ \}$ $\quad \quad \quad e \mid \text{super}(\bar{e});$</p> <p>Open frame $OF ::= \text{if } (\bullet == e) \{ \bar{s} \} \text{ else } \{ \bar{s} \}$ $\quad \quad \quad \text{if } (v == \bullet) \{ \bar{s} \} \text{ else } \{ \bar{s} \}$ $\quad \quad \quad \bullet.f \mid \bullet.f = e;$ $\quad \quad \quad v.f = \bullet; \mid (C)\bullet$ $\quad \quad \quad v.m(\bar{v}, \bullet, \bar{e}) \mid \bullet.m(\bar{e})$ $\quad \quad \quad \text{new } C(\bar{v}, \bullet, \bar{e})$ $\quad \quad \quad \text{super}(\bar{v}, \bullet, \bar{e});$ $\quad \quad \quad x = \bullet; \mid \text{return } \bullet;$</p>	<p>Exception frame $EF ::= \mathbf{NPE} \mid \mathbf{CCE}$</p> <p>Values $v ::= \text{null} \mid o$</p> <p>Variable stack $VS ::= MS \circ VS \mid []$</p> <p>Method scope $MS ::= BS \circ MS \mid []$</p> <p>Block scope $BS ::= \text{is a finite partial function from variables}$ $\quad \quad \quad \text{to pairs of expression types and values}$</p> <p>Heap $H ::= (H_v, H_t)$ $H_t ::= \text{is a finite partial function from } oids \text{ to}$ $\quad \quad \quad \text{class names}$ $H_v ::= \text{is a finite partial function from } oids$ $\quad \quad \quad \text{paired with field names to values}$</p>
---	---

Figure 2.6: Syntax for the MJ operational semantics

frames that correspond to exceptional states, described in §2.3.2. We use $H(o)$ as a shorthand for applying the second function in H to o , and $H(o, f)$ for applying the first to (o, f) . We use the obvious shorthands for dom and function update, for example $o \in dom(H)$ means o is in the domain of H 's first function.

In the following example we demonstrate how the variable scopes correctly model the block structure scoping of Java. The left hand side gives the source code and the right the contents of the variable stack.

Note: We could use only a single level of list, but this complicates other issues, as we can define a variable to have different types in two different non-nested scopes. Also we could collapse the third and fourth elements of a configuration and consider only non-empty stacks: we would simply be replacing a comma with a cons.

<pre> 1: B m(A a) { 2: B r; 3: if (this.x == a) { 4: r = a; 5: } else { 6: A t; 7: t = a.m1(); 8: r = a.m2(t, t); 9: } 10: return r; 11: }</pre>	<pre> ← VS ← ({a ↦ (A, v1)} ∘ []) ∘ VS ← ({r ↦ (B, null), a ↦ (A, v1)} ∘ []) ∘ VS ← ({ } ∘ {r ↦ (B, null), a ↦ (A, v1)} ∘ []) ∘ VS ← ({t ↦ (A, null)} ∘ {r ↦ (B, null), a ↦ (A, v1)} ∘ []) ∘ VS ← ({t ↦ (A, v3)} ∘ {r ↦ (B, null), a ↦ (A, v2)} ∘ []) ∘ VS ← ({t ↦ (A, v3)} ∘ {r ↦ (B, v4), a ↦ (A, v2)} ∘ []) ∘ VS ← ({r ↦ (B, v4), a ↦ (A, v2)} ∘ []) ∘ VS ← VS</pre>
---	---

Before calling this method, let us assume we have a variable scope, VS . A method call should not affect any variables in the current scopes, so we create a new method scope, $\{a \mapsto (A, v_1)\} \circ []$, on entry to the method. This scope consists of a single block scope that points the argument a at the value v_1 with a type annotation of A . On line 2 the block scope is extended to contain the new variable r . On line 5 we assumed that $this.x \neq a$ and enter into a new block. This has the effect of adding a new empty block scope, $\{ \}$, into the current method scope. On line 6 this new scope is extended to contain the variable t . Notice how it is added to the current

top scope, as was the variable τ . Updates can, however, occur to block scopes anywhere in the current method scope. This can be seen on line 8 where τ is updated in the enclosing scope. On line 9 the block scope is disposed of, and hence τ cannot be accessed by the statement on line 10.

We find it useful to define two operations on a variable stack, VS , in addition to the usual list operations. The first, $VS(x)$, evaluates a variable, x , in the top method scope, MS (in the top list of block scopes). This is a partial function and is only defined if the variable name is in the scope. The second, $VS[x \mapsto v]$, updates the top method scope, MS , with the value v for the variable x . Again this is a partial function and is undefined if the variable is not in the scope.

$$\begin{aligned} ((BS \circ MS) \circ VS)(x) &\stackrel{\text{def}}{=} \begin{cases} BS(x) & x \in \text{dom}(BS) \\ (MS \circ VS)(x) & \text{otherwise} \end{cases} \\ ([\circ VS)(x) &\stackrel{\text{def}}{=} [(x) \stackrel{\text{def}}{=} \text{undefined} \\ ((BS \circ MS) \circ VS)[x \mapsto v] &\stackrel{\text{def}}{=} \begin{cases} (BS[x \mapsto (v, C)] \circ MS) \circ VS & BS(x) = (-, C) \\ BS \circ ((MS \circ VS)[x \mapsto v]) & \text{otherwise} \end{cases} \\ ([\circ VS)[x \mapsto v] &\stackrel{\text{def}}{=} [[x \mapsto v] \stackrel{\text{def}}{=} \text{undefined} \end{aligned}$$

where we treat

$$\begin{aligned} BS \circ (MS \circ VS) &= (BS \circ MS) \circ VS \\ BS \circ \text{undefined} &= \text{undefined} \end{aligned}$$

2.3.1 Reductions

This section defines the transition rules that correspond to meaningful computation steps. In spite of the complexity of MJ, there are only seventeen rules, one for each syntactic constructor.

We begin by giving the transition rules for accessing, assigning values to, and declaring variables. Notice that the side condition in the E-VARWRITE rule ensures that we can only write to variables declared in the current method scope. The E-VARINTRO rule follows Java's restriction that a variable declaration *cannot* hide an earlier declaration within the current method scope.³ (Note also how the rule defines the binding for the new variable in the current *block* scope.)

$$\begin{array}{ll} \text{E-VARACCESS} & (H, VS, x, FS) \rightarrow (H, VS, v, FS) \quad \text{provided } VS(x) = v \\ \text{E-VARWRITE} & (H, VS, x = v; , FS) \rightarrow (H, VS', ; , FS) \quad \text{provided } VS' = VS[x \mapsto v] \\ \text{E-VARINTRO} & (H, (BS \circ MS) \circ VS, C x; , FS) \rightarrow (H, (BS' \circ MS) \circ VS, ; , FS) \\ & \quad \text{provided } x \notin \text{dom}(BS \circ MS) \wedge BS' = BS[x \mapsto (\text{null}, C)] \end{array}$$

Now we consider the rules for constructing and removing scopes. The first rule E-BLOCKINTRO introduces a new block scope, and leaves a 'marker' token, $\{\}$, on the frame stack. The second E-BLOCKELIM removes the token and the top block scope. The final rule E-RETURN leaves the scope of a method, by removing the top scope, MS .

³This sort of variable hiding is, in contrast, common in functional languages such as SML.

E-BLOCKINTRO	$(H, MS \circ VS, \{\bar{s}\}, FS) \rightarrow (H, (\{\} \circ MS) \circ VS, \bar{s}, (\{\}) \circ FS)$
E-BLOCKELIM	$(H, (BS \circ MS) \circ VS, \{\}, FS) \rightarrow (H, MS \circ VS, ;, FS)$
E-RETURN	$(H, MS \circ VS, \text{return } v;, FS) \rightarrow (H, VS, v, FS)$

Next we give the transition rules for the conditional expression. One should note that the resulting term of the transition is a block.

E-IF1	$(H, VS, (\text{if } (v_1 == v_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};), FS) \rightarrow (H, VS, \{\bar{s}_1\}, FS)$ provided $v_1 = v_2$
E-IF2	$(H, VS, (\text{if } (v_1 == v_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};), FS) \rightarrow (H, VS, \{\bar{s}_2\}, FS)$ provided $v_1 \neq v_2$

Next we consider the rules dealing with objects. First let us define the transition rule dealing with field access and assignment, as they are reasonably straightforward.

E-FIELDACCESS	$(H, VS, o.f, FS) \rightarrow (H, VS, v, FS)$ provided $(o, f) \in \text{dom}(H) \wedge H(o, f) = v$
E-FIELDWRITE	$(H, VS, o.f = v;, FS) \rightarrow (H', VS, ;, FS)$ provided $(o, f) \in \text{dom}(H) \wedge H' = H[(o, f) \mapsto v]$

Now we consider the rules dealing with casting of objects. Rule E-CAST simply ensures that the cast is valid (if it is not, the program should enter an error state—this is covered in §2.3.2). Rule E-NULLCAST simply ignores any cast of a `null` object reference.

E-CAST	$(H, VS, ((C_1)o), FS) \rightarrow (H, VS, o, FS)$ provided $H(o) = C_2 \wedge C_2 \prec C_1$
E-NULLCAST	$(H, VS, ((C)\text{null}), FS) \rightarrow (H, VS, \text{null}, FS)$

Let us consider the transition rule involving the creation of objects. The E-NEW rule creates a fresh *oid*, o , and extends the heap with the new heap object consisting of the class C and sets all the fields to `null`. As we are executing a new method (the constructor), a new method scope is created and added on to the variable stack. This method scope initially consists of just one block scope, that consists of bindings for the method parameters, and also a binding for the `this` identifier. The method body \bar{s} is then the next term to be executed, but importantly the continuation `return o;` is placed on the frame stack. This is because the result of this statement is the *oid* of the object, and the method scope is removed.

E-NEW	$(H, VS, \text{new } C(\bar{v}), FS) \rightarrow (H', (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS)$ provided $H' = H[o \mapsto C][(o, \bar{f}) \mapsto \text{null}] \wedge \delta_C(C) = \bar{C}$ $\wedge \text{cnbody}(C) = (\bar{x}, \bar{s}) \wedge o \notin \text{dom}(H)$ $\wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\}$
-------	--

Next we consider the transition rule for the `super` statement, which occurs inside constructor methods. This rule is basically the same as the previous, without the construction of the new object in the heap. It inspects the stack to determine the type of `this`, and hence which constructor body the `super` call refers to.

$$\begin{array}{l}
\text{E-SUPER} \quad (H, VS, \text{super}(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS) \\
\quad \text{provided } VS(\text{this}) = (o, C) \quad \wedge C \prec_1 C' \\
\quad \wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\} \wedge \delta_C(C) = \bar{C} \\
\quad \wedge \text{cbody}(C') = (\bar{x}, \bar{s})
\end{array}$$

Next we can give the transition rule for method invocation. Invocation is relatively straightforward: note how a new method scope is created, consisting of just the bindings for the method parameters and the `this` identifier. The reader should note the method body is selected by the dynamic, or run-time, type of the object, not its static type. We require two rules as a method returning a `void` type requires an addition to the stack to clear the new method scope once the method has completed. Recall that in E-METHOD rule, the last statement in the sequence \bar{s} will be a `return` if the method is well typed.

$$\begin{array}{l}
\text{E-METHOD} \quad (H, VS, o.m(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } e;) \circ FS) \\
\quad \text{provided } \text{mbody}(C, m) = (\bar{x}, \bar{s} \text{ return } e;) \quad \wedge H(o) = C \\
\quad \wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\} \wedge \delta_M(C)(m) = \bar{C} \rightarrow C' \\
\text{E-METHODVOID} \quad (H, VS, o.m(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS) \\
\quad \text{provided } H(o) = C \quad \wedge \delta_M(C)(m) = \bar{C} \rightarrow \text{void} \\
\quad \wedge \text{mbody}(C, m) = (\bar{x}, \bar{s}) \wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\}
\end{array}$$

Finally we have two reductions rules for fully reduced terms. The first rule deals with completed statements, and the second for evaluated expressions.

$$\begin{array}{l}
\text{E-SKIP} \quad (H, VS, ;, F \circ FS) \rightarrow (H, VS, F, FS) \\
\text{E-SUB} \quad (H, VS, v, F \circ FS) \rightarrow (H, VS, F[v], FS)
\end{array}$$

To assist the reader, all the reduction rules are repeated in full in Figure 2.7. There are a number of other reduction rules, which simply decompose terms. These rules essentially embody the order of evaluation, and are given in Figure 2.8.

2.3.2 Error states

A number of expressions will lead us to a predictable error state. These are errors that are allowed at run-time as they are dynamically checked for by the Java Virtual Machine. Java's type system is not capable of removing these errors statically. The two errors that can be generated, in MJ, are `NullPointerException`, written here **NPE** [41, §15.11.1, §15.12.4.4], and `ClassCastException` [41, §5.5], **CCE**.

$$\begin{array}{l}
\text{E-NULDFIELD} \quad (H, VS, \text{null}.f, FS) \rightarrow (H, VS, \mathbf{NPE}, FS) \\
\text{E-NULFWRITE} \quad (H, VS, \text{null}.f = v, FS) \rightarrow (H, VS, \mathbf{NPE}, FS) \\
\text{E-NULFMETHOD} \quad (H, VS, \text{null}.m(v_1, \dots, v_n), FS) \rightarrow (H, VS, \mathbf{NPE}, FS) \\
\text{E-INVCAST} \quad (H, VS, (C)o, FS) \rightarrow (H, VS, \mathbf{CCE}, FS) \\
\quad \text{provided } H(o) = (C', \mathbb{F}) \\
\quad \wedge C' \not\prec C
\end{array}$$

Definition 2.3.1 (Terminal configuration). A configuration is said to be *completed* if it is of the form $(H, VS, v, [])$ or $(H, VS, ;, [])$, and is an *exception* if it is of the form (H, VS, EF, FS) . A configuration is said to be *terminal* if it is completed or an exception. We call a configuration *stuck* if it cannot reduce, that is *config* is stuck iff $\neg \exists \text{config}'. \text{config} \rightarrow \text{config}'$.

E-VARACCESS	$(H, VS, x, FS) \rightarrow (H, VS, v, FS)$	provided $VS(x) = v$
E-VARWRITE	$(H, VS, x = v;, FS) \rightarrow (H, VS', ;, FS)$	provided $VS' = VS[x \mapsto v]$
E-VARINTRO	$(H, (BS \circ MS) \circ VS, C x;, FS) \rightarrow (H, (BS' \circ MS) \circ VS, ;, FS)$	provided $x \notin \text{dom}(BS \circ MS) \wedge BS' = BS[x \mapsto (\text{null}, C)]$
E-BLOCKINTRO	$(H, MS \circ VS, \{\bar{s}\}, FS) \rightarrow (H, (\{\} \circ MS) \circ VS, \bar{s}, (\{\} \circ FS)$	
E-BLOCKELIM	$(H, (BS \circ MS) \circ VS, \{\}, FS) \rightarrow (H, MS \circ VS, ;, FS)$	
E-RETURN	$(H, MS \circ VS, \text{return } v;, FS) \rightarrow (H, VS, v, FS)$	
E-IF1	$(H, VS, (\text{if } (v_1 == v_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};), FS) \rightarrow (H, VS, \{\bar{s}_1\}, FS)$	provided $v_1 = v_2$
E-IF2	$(H, VS, (\text{if } (v_1 == v_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};), FS) \rightarrow (H, VS, \{\bar{s}_2\}, FS)$	provided $v_1 \neq v_2$
E-FIELDACCESS	$(H, VS, o.f, FS) \rightarrow (H, VS, v, FS)$	provided $(o, f) \in \text{dom}(H) \wedge H(o, f) = v$
E-FIELDWRITE	$(H, VS, o.f = v;, FS) \rightarrow (H', VS, ;, FS)$	provided $(o, f) \in \text{dom}(H) \wedge H' = H[(o, f) \mapsto v]$
E-CAST	$(H, VS, ((C_1)o), FS) \rightarrow (H, VS, o, FS)$	provided $H(o) = C_2 \wedge C_2 \prec C_1$
E-NULLCAST	$(H, VS, ((C)\text{null}), FS) \rightarrow (H, VS, \text{null}, FS)$	
E-NEW	$(H, VS, \text{new } C(\bar{v}), FS) \rightarrow (H', (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS)$	provided $H' = H[o \mapsto C][(o, \bar{f}) \mapsto \text{null}] \wedge \delta_C(C) = \bar{C}$ $\wedge \text{cnbody}(C) = (\bar{x}, \bar{s}) \wedge o \notin \text{dom}(H)$ $\wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\}$
E-SUPER	$(H, VS, \text{super}(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS)$	provided $VS(\text{this}) = (o, C) \wedge C \prec_1 C'$ $\wedge BS = \{\text{this} \mapsto (o, C'), \bar{x} \mapsto (\bar{v}, \bar{C})\} \wedge \delta_C(C) = \bar{C}$ $\wedge \text{cnbody}(C') = (\bar{x}, \bar{s})$
E-METHOD	$(H, VS, o.m(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } e;) \circ FS)$	provided $\text{mbody}(C, m) = (\bar{x}, \bar{s} \text{ return } e;) \wedge H(o) = C$ $\wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\} \wedge \delta_M(C)(m) = \bar{C} \rightarrow C'$
E-METHODVOID	$(H, VS, o.m(\bar{v}), FS) \rightarrow (H, (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS)$	provided $H(o) = C \wedge \delta_M(C)(m) = \bar{C} \rightarrow \text{void}$ $\wedge \text{mbody}(C, m) = (\bar{x}, \bar{s}) \wedge BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\}$
E-SKIP	$(H, VS, ;, F \circ FS) \rightarrow (H, VS, F, FS)$	
E-SUB	$(H, VS, v, F \circ FS) \rightarrow (H, VS, F[v], FS)$	

Figure 2.7: MJ reduction rules

EC-SEQ	$(H, VS, s_1 s_2 \dots s_n, FS) \rightarrow (H, VS, s_1, (s_2 \dots s_n) \circ FS)$
EC-RETURN	$(H, MS \circ VS, \text{return } e;, FS) \rightarrow (H, MS \circ VS, e, (\text{return } \bullet;) \circ FS)$
EC-EXPSTATE	$(H, VS, e;, FS) \rightarrow (H, VS, e, FS)$
EC-IF1	$(H, VS, \text{if } (e_1 == e_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};, FS) \rightarrow (H, VS, e_1, (\text{if } (\bullet == e_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};) \circ FS)$
EC-IF2	$(H, VS, \text{if } (v_1 == e_2)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};, FS) \rightarrow (H, VS, e_2, (\text{if } (v_1 == \bullet)\{\bar{s}_1\} \text{ else } \{\bar{s}_2\};) \circ FS)$
EC-FIELDACCESS	$(H, VS, e.f, FS) \rightarrow (H, VS, e, (\bullet.f) \circ FS)$
EC-CAST	$(H, VS, (C)e, FS) \rightarrow (H, VS, e, ((C)\bullet) \circ FS)$
EC-FIELDWRITE1	$(H, VS, e_1.f = e_2;, FS) \rightarrow (H, VS, e_1, (\bullet.f = e_2;) \circ FS)$
EC-FIELDWRITE2	$(H, VS, v_1.f = e_2;, FS) \rightarrow (H, VS, e_2, (v_1.f = \bullet;) \circ FS)$
EC-VARWRITE	$(H, VS, x = e;, FS) \rightarrow (H, VS, e, (x = \bullet;) \circ FS)$
EC-NEW	$(H, VS, \text{new } C(\bar{v}, e, \bar{e}), FS) \rightarrow (H, VS, e, (\text{new } C(\bar{v}, \bullet, \bar{e})) \circ FS)$
EC-SUPER	$(H, VS, \text{super}(\bar{v}, e, \bar{e}), FS) \rightarrow (H, VS, e, (\text{super}(\bar{v}, \bullet, \bar{e})) \circ FS)$
EC-METHOD1	$(H, VS, e.m(\bar{e}), FS) \rightarrow (H, e, (\bullet.m(\bar{e})) \circ FS)$
EC-METHOD2	$(H, VS, v.m(\bar{v}, e, \bar{e}), FS) \rightarrow (H, VS, e, (v.m(\bar{v}, \bullet, \bar{e})) \circ FS)$

Figure 2.8: MJ syntax decomposition reduction rules

2.3.3 Example execution

To help the reader understand our operational semantics, in this section we will consider a simple code fragment and see how the operational semantics captures its dynamic behaviour.

Consider the following MJ code whose effect is to swap the contents of two variables `var1` and `var2`, using a temporary variable `t`.

```
if(var1 == var2) {
  ;
} else {
  Object t;
  t = var1;
  var1 = var2;
  var2 = t;
}
```

We consider its execution in a configuration where MS maps `var1` to a value, say $v1$, and `var2` to a value, say $v2$.

$$\begin{aligned}
& (H, MS \circ VS, \text{if } (\text{var1} == \text{var2})\{;\} \text{ else } \{\dots\}, FS) \\
\rightarrow & (H, MS \circ VS, \text{var1}, (\text{if } (\bullet == \text{var2})\{;\} \text{ else } \{\dots\}) \circ FS) \\
\rightarrow & (H, MS \circ VS, v1, (\text{if } (\bullet == \text{var2})\{;\} \text{ else } \{\dots\}) \circ FS) \\
\rightarrow & (H, MS \circ VS, (\text{if } (v1 == \text{var2})\{;\} \text{ else } \{\dots\}), FS) \\
\rightarrow & (H, MS \circ VS, \text{var2}, (\text{if } (v1 == \bullet)\{;\} \text{ else } \{\dots\}) \circ FS) \\
\rightarrow & (H, MS \circ VS, v2, (\text{if } (v1 == \bullet)\{;\} \text{ else } \{\dots\}) \circ FS) \\
\rightarrow & (H, MS \circ VS, (\text{if } (v1 == v2)\{;\} \text{ else } \{\dots\}), FS)
\end{aligned}$$

At this point there are two possibilities: we will consider the case where $v1 \neq v2$.

Note: For compactness, we use \circ for `Object`.

$$\begin{aligned}
&\rightarrow (H, (\{\}) \circ MS \circ VS, \circ t; \dots, (\{\}) \circ FS) \\
&\rightarrow (H, (\{\}) \circ MS \circ VS, \circ t; (t = \text{var1}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, ;, (t = \text{var1}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, (t = \text{var1}; \dots), (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, t = \text{var1}; (\text{var1} = \text{var2}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, \text{var1}, (t = \bullet;) \circ (\text{var1} = \text{var2}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, v1, (t = \bullet;) \circ (\text{var1} = \text{var2}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (\text{null}, 0)\} \circ MS \circ VS, t = v1; (\text{var1} = \text{var2}; \dots) \circ (\{\}) \circ FS)
\end{aligned}$$

At this point we update the variable stack; note how the update does not change the type.

$$\begin{aligned}
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, ;, (\text{var1} = \text{var2}; \dots) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, \text{var1} = \text{var2}; \dots, (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, \text{var1} = \text{var2}; (\text{var2} = t;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, \text{var2}, \text{var1} = \bullet; \circ (\text{var2} = t;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, v2, \text{var1} = \bullet; \circ (\text{var2} = t;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS \circ VS, \text{var1} = v2; (\text{var2} = t;) \circ (\{\}) \circ FS)
\end{aligned}$$

Let MS' be a variable scope which is the same as MS except that var1 is mapped to $v2$ instead of $v1$.

$$\begin{aligned}
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS' \circ VS, ;, (\text{var2} = t;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS' \circ VS, \text{var2} = t; (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS' \circ VS, t, (\text{var2} = \bullet;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS' \circ VS, v1, (\text{var2} = \bullet;) \circ (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS' \circ VS, \text{var2} = v1; (\{\}) \circ FS)
\end{aligned}$$

Let MS'' be a variable scope which is the same as MS' except that var2 is mapped to $v1$.

$$\begin{aligned}
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS'' \circ VS, ;, (\{\}) \circ FS) \\
&\rightarrow (H, (\{t \mapsto (v1, 0)\} \circ MS'' \circ VS, \{\}, FS) \\
&\rightarrow (H, MS'' \circ VS, ;, FS)
\end{aligned}$$

At this point the execution of the `if` statement has been completed and its temporary variable, `temp` has been removed from the scope. The variable stack has had the values of `var1` and `var2` correctly swapped.

2.4 Well-typed configuration

To prove type soundness for MJ, we need to extend our typing rules from expressions and statements to configurations. We write $\delta \vdash (H, VS, CF, FS) : \tau$ to mean (H, VS, CF, FS) is well-typed with respect to δ and will result in a value of type τ (or a valid error state). We break this into three properties, written: $\delta \vdash H \text{ ok}$, $\delta, H \vdash VS \text{ ok}$ and $\delta, H, VS \vdash CF \circ FS : \text{void} \rightarrow \tau$.

The first property, written $\delta \vdash H \text{ ok}$, ensures that every field points to a valid object or `null`, and all the types mentioned in the heap are in δ .

TH-OBJECTTYPED	$\frac{H(o) = C \quad C \prec \tau \quad \delta \vdash C \text{ ok}}{\delta, H \vdash o : \tau}$
TH-NULLTYPED	$\frac{\delta \vdash C \text{ ok}}{\delta, H \vdash \text{null} : C}$
TH-OBJECTOK	$\frac{H(o) = (C) \quad \forall f \in \text{dom}(\delta_{\mathbb{F}}(C)). \quad \delta, H \vdash H(o, f) : \delta_{\mathbb{F}}(C)(f)}{\delta, H \vdash o \text{ ok}}$
TH-HEAPOK	$\frac{\forall o \in \text{dom}(H). \quad \delta, H \vdash o \text{ ok}}{\delta \vdash H \text{ ok}}$

The second property, written $\delta, H \vdash VS \text{ ok}$, constrains every variable to be either a valid object identifier, or `null`.

TV-EMPTY	$\delta, H \vdash [] \text{ ok}$
TV-MSEMPY	$\delta, H \vdash VS \text{ ok}$
	$\delta, H \vdash [] \circ VS \text{ ok}$
TV-STACK	$\delta, H \vdash MS \circ VS \text{ ok} \quad \forall (v, C) \in \text{cod}(BS). \delta, H \vdash v : C$
	$\delta, H \vdash (BS \circ MS) \circ VS \text{ ok}$

We generalise the final property to type arbitrary framestacks, written $\delta, H, VS \vdash FS : \tau \rightarrow \tau'$. It types each frame in the context formed from the heap and variable stack. This requires us to define a collapsing of the context to form a typing environment. We must extend the typing environment to map, in addition to variables, object identifiers, o , and holes, \bullet , to values. The collapsing function is defined as follows.

$\text{heaptypes}((H_t, H_v))$	$\stackrel{\text{def}}{=} H_t$
$\text{stacktype}((BS \circ MS))$	$\stackrel{\text{def}}{=} \{x \mapsto C \mid (x \mapsto C, v) \in BS\} \uplus \text{stacktype}(MS)$
$\text{stacktype}([] \circ VS)$	$\stackrel{\text{def}}{=} \emptyset$
$\text{context}(H, MS \circ VS)$	$\stackrel{\text{def}}{=} \text{heaptypes}(H) \uplus \text{stacktype}(MS)$

The syntax of expressions in framestacks is extended to contain both object identifiers and holes. Hence we require two additional expression typing rules.

TE-OID	$o : C \in \gamma \quad \delta \vdash \gamma \text{ ok} \quad \vdash \delta \text{ ok}$
	$\delta; \gamma \vdash o : C$
TE-HOLE	$\bullet : C \in \gamma \quad \delta \vdash \gamma \text{ ok} \quad \vdash \delta \text{ ok}$
	$\delta; \gamma \vdash \bullet : C$

We can now define frame stack typing as shown in Figure 2.9. We have the obvious typing for an empty stack. We require special typing rules for the frames that alter variable scoping: TF-STACKBLOCK, TF-STACKMETHOD, TF-STACKMETHOD2 and TF-STACKINTRO. We also require a rule, TF-SEQUENCE, for unrolling sequences because the sequence can contain items to alter scoping. We then require two rules for typing the rest of the frames; one for frames that require an argument, called TF-STACKOPEN, and one for frames that do not, called TF-STACKCLOSED.

2.5 Type soundness

Our main technical contribution in this chapter is a proof of type soundness of the MJ type system. In order to prove correctness we first prove two useful propositions in the style of Wright and Felleisen[86]. The first proposition states that any well typed non-terminal configuration can make a reduction step.

Proposition 2.5.1 (Progress). *If config is not terminal and $\delta \vdash \text{config} : \tau$ then $\exists \text{config}' . \text{config} \rightarrow \text{config}'$.*

Proof. By case analysis of CF . By considering all well typed configurations, we can show how each term can reduce.

TF-STACKEMPTY	$\delta, H, (BS \circ []) \circ [] \vdash [] : \tau \rightarrow \tau$
TF-ERROR	$\delta, H, VS \vdash EF \circ FS : \tau \rightarrow \tau'$
TF-STACKBLOCK	$\frac{\delta, H, MS \circ VS \vdash FS : \text{void} \rightarrow \tau}{\delta, H, (BS \circ MS) \circ VS \vdash (\{\}) \circ FS : \tau' \rightarrow \tau}$
TF-STACKMETHOD	$\frac{\delta, H, VS \vdash FS : \tau \rightarrow \tau'}{\delta, H, (BS \circ []) \circ VS \vdash (\text{return } \bullet;) \circ FS : \tau \rightarrow \tau'}$
TF-STACKMETHOD2	$\frac{\delta; \text{context}(H, MS \circ VS) \vdash e : \tau \quad H, VS \vdash FS : \tau \rightarrow \tau'}{\delta, H, MS \circ VS \vdash (\text{return } e;) \circ FS : \tau'' \rightarrow \tau'}$
TF-STACKINTRO	$\frac{\delta, H, (BS[x \mapsto (\text{null}, C)] \circ MS) \circ VS \vdash FS : \text{void} \rightarrow \tau}{H, (BS \circ MS) \circ VS \vdash (C \ x;) \circ FS : \tau' \rightarrow \tau}$ provided $x \notin \text{dom}(BS \circ MS)$
TF-SEQUENCE	$\frac{\delta, H, VS \vdash (s_1) \circ (s_2 \dots s_n) \circ FS : \tau \rightarrow \tau'}{\delta, H, VS \vdash (s_1 s_2 \dots s_n) \circ FS : \tau \rightarrow \tau'}$
TF-STACKOPEN	$\frac{\delta; \text{context}(H, VS), \bullet : C \vdash OF : \tau \quad H, VS \vdash FS : \tau \rightarrow \tau'}{\delta, H, VS \vdash OF \circ FS : C \rightarrow \tau'}$ provided $OF \neq (\text{return } \bullet;)$
TF-STACKCLOSED	$\frac{\delta; \text{context}(H, VS) \vdash CF : \tau \quad H, VS \vdash FS : \tau \rightarrow \tau'}{\delta, H, VS \vdash CF \circ FS : \tau'' \rightarrow \tau'}$ provided $CF \neq (\text{return } e;)$ \wedge $CF \neq (\{\})$ \wedge $CF \neq s_1 \dots s_n \wedge n > 1$ \wedge $CF \neq C \ x$

Figure 2.9: Frame stack typing

- Case:** $CF = (\text{return } e;)$ As it is well typed we know $VS = MS \circ VS'$. This has two possible cases of reducing. Either $e = v$ and can reduce by E-RETURN or $e \neq v$ and can reduce by EC-RETURN.
- Case:** $CF = (\{\})$ As this is well typed, we know $VS = (BS \circ MS) \circ VS'$ and hence this can reduce by E-BLOCK.
- Case:** $CF = C \ x;$ As it is well typed, we know $VS = (BS \circ MS) \circ VS'$ and also $x \notin \text{context}(VS)$. From the definition of context we can see that this gives $x \notin \text{dom}(BS \circ MS)$ and hence it can reduce by E-VARINTRO.
- Case:** $CF = x$ We know that x must be in $\text{context}(H, VS)$, and as x cannot be in H it must come from VS and more precisely in MS , where $VS = MS \circ VS'$, hence it can reduce by E-VARACCESS.
- Case:** $CF = o$ Either $FS \neq []$ and reduces by E-SUB, or $FS = []$ which is a terminal framestack.
- Case:** $CF = \text{null}$ Either $FS \neq []$ and reduces by E-SUB, or $FS = []$ which is a terminal framestack.
- Case:** $CF = e.f$ This can be broken into three cases. Either $e = \text{null}$ and reduces by E-NULDFIELD; $e = o$ and reduces by E-FIELDACCESS; or $e \neq v$ and reduces by EC-FIELDACCESS.
- Case:** $CF = e'.m(e_1, \dots, e_n)$ If $e' \neq v$ then EC-METHOD1 will apply. If any of e_1, \dots, e_n are not values then EC-METHOD2 will be applied. Otherwise we have the case where $CF = v'.m(v_1, \dots, v_n)$ in which case E-METHOD applies if $v' = o$ or E-NULDMETHOD if $v' = \text{null}$.
- Case:** $CF = \text{new } C(e_1, \dots, e_n)$ reduces by EC-NEW when $e_i \neq v$ or E-NEW otherwise.
- Case:** $CF = (C)e$ Either $e \neq v$ and reduces by EC-CAST; $e = o$ and reduces by either E-CAST or

- E-INVCAST depending on the type of o ; or $e = \text{null}$ and reduces by E-NULLCAST.
- Case:** $CF = ;$; Either $FS \neq []$ and reduces by E-SKIP; or $FS = []$ which is a terminal state.
- Case:** $CF = s_1 \dots s_n$ reduces by EC-SEQ.
- Case:** $CF = \mathbf{if} (e_1 == e_2) \{ \bar{s}_1 \} \mathbf{else} \{ \bar{s}_2 \}$. Either $e_1 \neq v_1$ and reduces by EC-IF1; $e_2 \neq v_2 \wedge e_1 = v_1$ and reduces by EC-IF2; or $e_1 = v_1 \wedge e_2 = v_2$ can reduce by either E-IF1 and E-IF2 depending on the test .
- Case:** $CF = x = e$ Either $e = v$ and can reduce using E-VARWRITE if x is in the environment. The typing rule tells us that $\text{context}(H, VS) \vdash x : C$, hence it is in the environment. Otherwise $e \neq v$ which can reduce by EC-VARWRITE.
- Case:** $CF = e.f = e'$ Either $e \neq v$ and reduces using EC-FIELDWRITE1; $e = v \wedge e' \neq v'$ reduces using EC-FIELDWRITE2; $e = o \wedge e' = v'$ and reduces using E-FIELDWRITE; or $e = \text{null} \wedge e' = v'$ reduces using E-NULLWRITE.
- Case:** $CF = e;$ reduces by EC-EXPSTATE.
- Case:** $CF = \mathbf{super}(e_1, \dots, e_n)$ If all the expressions are values then this can reduce by E-SUPER, otherwise it can reduce by EC-SUPER.

□

Next we find it useful to prove the following few lemmas. The first states that subtyping on frame stacks is *covariant*.

Lemma 2.5.2 (Covariant subtyping of frame stack). *if $\delta, H, VS \vdash FS : \tau_1 \rightarrow \tau_2$ and $\tau_3 \prec \tau_1$ then $\exists \tau_4. H, VS \vdash FS : \tau_3 \rightarrow \tau_4$ and $\tau_4 \prec \tau_2$.*

Proof. By induction on the size of FS . The base case, $FS = []$, holds trivially as TF-STACKEMPTY is covariant. For the inductive step we must show that if FS is covariant, then so is $F \circ FS$. If F is closed then trivially holds as it ignores the argument. We only need consider open frames. First consider $\text{return } \bullet$; as this affects the typing environment. Holds because it is covariantly typed if the remainder of the frame stack is covariant. The remaining open frames require us to prove:

$$\delta; \gamma, \bullet : \tau \vdash OF : \tau_1 \wedge \tau_2 \prec \tau \Rightarrow \exists \tau_3. \delta; \gamma, \bullet : \tau_2 \vdash OF : \tau_3 \wedge \tau_3 \prec \tau_1$$

We proceed by case analysis on OF .

Case: $OF = \mathbf{if} (\bullet == e) \{ \bar{s}_1 \} \mathbf{else} \{ \bar{s}_2 \}$; Holds, because e, \bar{s}_1 and \bar{s}_2 do not contain \bullet they are unaffected by the change in type and using TS-IF or TS-STUPIDIF the statement is well-typed if τ_2 is a valid type.

Case: $OF = \mathbf{if} (v == \bullet) \{ \bar{s}_1 \} \mathbf{else} \{ \bar{s}_2 \}$; Similar to previous case.

Case: $OF = \bullet.f$ Holds as $\tau_2 \prec \tau$ and field types cannot be overridden, hence we know $\delta_F(\tau_2)(f) = \delta_F(\tau)(f)$.

Case: $OF = \bullet.f = e$; Similar to previous case.

Case: $OF = \bullet.m(e_1, \dots, e_n)$ Similar to previous case, except we use the fact method types cannot be overridden.

Case: $OF = x = \bullet$; We know from the assumptions that

$$\frac{\overline{\delta; \gamma, \bullet : \tau \vdash x : C} \quad \overline{\delta; \gamma, \bullet : \tau \vdash \bullet : \tau} \quad \overline{\tau \prec C}}{\delta; \gamma, \bullet : \tau \vdash x = \bullet : \text{void}}$$

As the sub-typing relation is transitive and $\tau_2 \prec \tau$, we can see this is well typed for $\bullet : \tau_2$, and the result type is the same.

Case: $OF = v.f = \bullet$; Similar to previous case.

Case: $OF = v.m(v_1, \dots, v_{i-1}, \bullet, e_{i+1}, \dots, e_n)$ Similar to previous case.

Case: $OF = \mathbf{new} C(v_1, \dots, v_{i-1}, \bullet, e_{i+1}, \dots, e_n)$ Similar to previous case.

Case: $OF = \mathbf{super}(v_1, \dots, v_{i-1}, \bullet, e_{i+1}, \dots, e_n)$ Similar to previous case.

Case: $OF = (C)\bullet$ The result type of this frame does not depend on the argument type. The three possible typing rules for this case combined to only require \bullet is typeable. Hence this case is trivial. □

Next we prove that adding a new heap object preserves the typing.

Lemma 2.5.3 (Heap extension preserves typing). *If $\delta, H, VS \vdash FS : \tau \rightarrow \tau'$, $o \notin \text{dom}(H)$ and $\delta \vdash C$ and $\delta, H \vdash VS \text{ ok}$ then (1) $\delta, H[o \mapsto C][o, \bar{f} \mapsto \text{null}], VS \vdash FS : \tau \rightarrow \tau'$ and (2) $\delta, H[o \mapsto C][o, \bar{f} \mapsto \text{null}] \vdash VS \text{ ok}$.*

Proof. First we prove (1) by induction on the length of FS . All the rules for typing the frame stack follow by induction, except for TF-OPENFRAME, TF-CLOSEDFRAME and TF-METHOD2. To prove these it suffices to prove

$$\delta; \gamma \vdash F : \tau \wedge \delta \vdash C \Rightarrow \delta; \gamma, o : C \vdash F : \tau$$

where $o \notin \text{dom}(\gamma)$.

We can see $\delta \vdash \gamma \text{ ok} \wedge \delta \vdash C \Rightarrow \delta \vdash \gamma, o : C \text{ ok}$ from the definition of $\gamma \text{ ok}$. We can use this to prove all the axioms of the typing judgement. The rules follow directly by induction except for TS-INTRO. For this rule, we must prove that x and o are different. This is true as they are from different syntactic categories.

To prove (2) we induct over the structure of VS . We require the following

$$\delta, H \vdash v : \tau \Rightarrow \delta, H[o \mapsto C][o, \bar{f} \mapsto \text{null}] \vdash v : \tau$$

which holds directly from the definitions. □

We require two lemmas for dealing with typing sequences when they are added to the frame stack.

Lemma 2.5.4 (Block). *If $\delta, H, MS \circ VS \vdash FS : \text{void} \rightarrow \tau$ and $\delta; \text{context}(H, (BS \circ MS) \circ VS) \vdash \bar{s} : \text{void}$, then $\delta, H, (BS \circ MS) \circ VS \vdash \bar{s} \circ \{ \} \circ FS : \text{void} \rightarrow \tau$*

Lemma 2.5.5 (Return). *If $\delta, H, VS \vdash FS : \tau \rightarrow \tau'$ and $\delta; \text{context}(H, MS \circ VS) \vdash s_1 \dots s_n : \tau$, then $\delta, H, MS \circ VS \vdash s_1 \circ \dots \circ s_n \circ FS : \text{void} \rightarrow \tau'$ where s_n is of the form $\text{return } e$;*

The two lemmas have very similar proofs. We present the proof for the first lemma.

Proof. We prove this by induction on the size of \bar{s} .

Base case: Assume $\delta, H, MS \circ VS \vdash FS : \text{void} \rightarrow \tau$ and $\delta; \text{context}(H, (BS \circ MS) \circ VS) \vdash s : \text{void}$ by TF-CLOSEDFRAME and TF-BLOCK we can derive $\delta, H, (BS \circ MS) \circ VS \vdash s \circ \{ \} \circ FS : \text{void} \rightarrow \tau$ as required.

Inductive case: Assume property holds for \bar{s} and show it holds for $s_1 \bar{s}$. Assume:

$$\delta, H, MS \circ VS \vdash FS : \text{void} \rightarrow \tau \tag{2.1}$$

$$\delta; \text{context}(H, (BS' \circ MS) \circ VS) \vdash s_1 \bar{s} \tag{2.2}$$

Prove:

$$\delta, H, (BS' \circ MS) \circ VS \vdash s_1 \bar{s} \circ \{ \} \circ FS : \text{void} \rightarrow \tau \tag{2.3}$$

We proceed by a case analysis of s_1 .

Case: $s_1 = Cx$; From (2.2) we can deduce $\delta; \text{context}(H, (BS'[x \mapsto (C, \text{null})] \circ MS) \circ VS) \vdash \bar{s} : \text{void}$. Using this and specialising the inductive hypothesis allows us to deduce $\delta, H, (BS'[x \mapsto (C, \text{null})] \circ MS) \circ VS \vdash \bar{s} \circ \{ \} \circ FS : \text{void} \rightarrow \tau$. This is exactly what we require to prove (2.3), which completes this case.

Case: $s_1 \neq Cx$; From (2.2) we can deduce

$$\delta; \text{context}(H, (BS' \circ MS) \circ VS) \vdash \bar{s} : \text{void} \quad (2.4)$$

$$\delta; \text{context}(H, (BS' \circ MS) \circ VS) \vdash s_1 : \text{void} \quad (2.5)$$

Specialising the inductive hypothesis and using (2.4) gives us $\delta, H, (BS' \circ MS) \circ VS \vdash \bar{s} \circ \{ \} \circ FS : \text{void} \rightarrow \tau$, which combined with (2.5) gives (2.3), and completes this case. \square

We can now prove the second important proposition of our reduction system, which states that if a configuration can make a transition, then the resulting configuration is of the appropriate type. (This is sometimes referred to as the “subject reduction” theorem.)

Proposition 2.5.6 (Type preservation). *If $\delta \vdash \text{config} : \tau$ and $\text{config} \rightarrow \text{config}'$ then $\exists \tau'. \delta \vdash \text{config}' : \tau'$ where $\tau' \prec \tau$.*

Proof. By case analysis on the \rightarrow relation. The decomposition rules, in Figure 2.8, and E-SKIP follow by trivially restructuring the proof trees. The remaining rules we prove as follows. Assume

- a. $\delta \vdash H$ ok
- b. $\delta, H \vdash VS$ ok
- c. $\delta, H, VS \vdash CF \circ FS : \text{void} \rightarrow \tau$
- d. $(H, VS, CF, FS) \rightarrow (H', VS', CF', FS')$

We must prove the following

1. $\delta \vdash H'$ ok
2. $\delta, H' \vdash VS'$ ok
3. $\delta, H', VS' \vdash CF' \circ FS' : \text{void} \rightarrow \tau' \wedge \tau' \prec \tau$

Case: E-SUB If $FS = CF \circ FS'$ then holds trivially.

Otherwise $FS = OF \circ FS'$.

1. Holds by assumptions as heap is unchanged.
2. Holds by assumptions as stack is unchanged.
3. Holds due to

$$\begin{aligned} \delta; \text{context}(H, VS), \bullet : \tau' \vdash OF : \tau'' \wedge \delta; \text{context}(H, VS) \vdash v : \tau' \\ \Rightarrow \delta; \text{context}(H, VS) \vdash OF[v/\bullet] : \tau'' \end{aligned}$$

which is shown by induction on OF .

Case: E-RETURN

1. Holds as heap is unchanged.
2. Holds as new stack is a sub stack of the old one.
3. Holds as returned value only depends on the heap.

Case: E-VARACCESS Let $F = x$

1. Holds as heap is unchanged
2. Holds as stack is unchanged
3. As stack is well typed, if x evaluates to v , then v 's type is a subtype of x 's static type. Using covariant lemma rest of framestack is well-typed.

Case: E-VARWRITE Let $F = x = v$;

1. Holds as heap is unchanged.
2. For F to be well-typed v must be a sub-type of x 's static type. So updated stack is well-formed.
3. Holds as new frame stack is contained in the old one, with a skip prepended.

Case: E-VARINTRO

1. Holds as heap is unchanged.
2. New entry in stack is well-typed, so new stack is well-typed.
3. New framestack is typed by rearranging the old ones proof.

Case: E-IF1

1. Holds as heap is unchanged.
2. Holds as stack is unchanged.
3. New framestack is typed, because all the elements were typed by the previous one.

Case: E-IF2 Identical to previous case.

Case: E-BLOCKINTRO

1. Holds as heap is unchanged.
2. Adding an empty block to a well-typed stack results in a well-typed stack.
3. Holds as consequence of lemma 2.5.4.

Case: E-CAST

1. Holds as heap is unchanged.
2. Holds as stack is unchanged.
3. Holds directly by covariant lemma.

Case: E-FIELDWRITE Let $F = o.f = v$;

1. As F is well typed we know that v is a subtype of the field's type the heap update leaves it well typed.
2. Holds as the stack is unchanged.
3. New framestack is typed, because all the elements were typed by the previous one, and the typing information in the heap is unchanged.

Case: E-FIELDACCESS Let $F = o.f$

1. Holds as heap is unchanged.
2. Holds as stack is unchanged.
3. As heap is well-typed, we know $H(H(o, f)) \prec \delta_F(H(o), f)$. Hence the new stack is typed by the covariant lemma.

Case: E-METHOD Let $F = o.m(\bar{v})$

1. Holds as the heap is unchanged.
2. As frame stack was well-typed if \bar{v} has type \bar{C} then $\bar{C} \prec \bar{C}'$ and $\delta_F(H(o), C) = \bar{C}' \rightarrow C$. Therefore new scope is well-typed.

3. As the new scope gives the same scope as the method was type-checked against, we can use lemma 2.5.5 and the covariant lemma to show the new framestack is well-typed. Covariant lemma is required as the `return e;` could be a more specific type than the method signature gives.

Case: E-METHODVOID Same as previous case, with a trivial alteration for the `return`.

Case: E-NEW

1. The heap is extended with an object. This extension has all the required fields, and their values are all null. So the new heap is well-typed.
2. Follows by similar argument to previous case, with the heap extension lemma 2.5.3.
3. Follows by similar argument to previous case, but covariant lemma is not required, and heap extension preserves typing is require 2.5.3.

Case: E-SUPER

1. Holds as heap is unchanged.
2. Follows by similar argument to previous case.
3. Follows by similar argument to previous case.

□

We can now combine the propositions 2.5.1 and 2.5.6 to prove the type safety of the MJ type system.

Theorem 2.5.7 (Type safety). *If $config : \tau$ and $config \rightarrow^* config'$ then $config' : \tau'$ where $\tau' \prec \tau$ and either $config'$ is terminal or not stuck.*

2.6 Related Work

There have been many works on formalizing subsets of Java. Our work is closely related to, and motivated by, Featherweight Java [47]. We have upheld the same philosophy, by keeping MJ a valid subset of Java. However FJ lacks many key features we wish to reason about. It has no concept of state or object identity, which is essential for developing usable specification logics. Our work has extended FJ to contain what we feel are the important imperative features of Java.

Another related calculus is Classic Java [38], which embodies many of the language features of MJ. However, Classic Java is not a valid subset of Java, as it uses type annotations and also uses let-binding to model both sequencing and locally scoped variables. Hence several features of Java, such as its block structured mutable state and unusual syntactic distinction of promotable expressions are not modelled directly. However, Felleisen *et al.*'s motivation is different to ours: they are interested in extending Java with mixins, rather than reasoning about features of Java.

Eisenbach, Drossopoulou, *et al.* have developed type soundness proofs for various subsets of Java [33, 34]. They consider a larger subset of Java than MJ as they model exceptions and arrays, however they do not model block structured scoping. Our aim was to provide an imperative core subset of Java amenable to formal proof, rather than to prove soundness of a large fragment. Syme [82] has formalised Eisenbach, Drossopoulou, *et al.*'s semantics in Isabelle.

There have been other fragments of Java formalised in theorem provers. Nipkow and Oheimb [62] formalize a subset of Java in the Isabelle theorem prover. More recently, Nipkow and Klein have formalized a larger subset of Java, named Jinja [51]. They also model the bytecode and compiler in the theorem prover.

3

A logic for Java

In this chapter we present a programming logic for reasoning about Java. Although the logic described in this chapter is unable to reason about forms of abstraction commonly found in object-oriented languages (namely encapsulation and inheritance) it forms a basis for the logics in the rest of the thesis. Additionally it serves as an introduction to separation logic.

We choose to build our logic on top of the separation logic by O’Hearn, Reynolds and others [76, 48, 65, 75], as it addresses the problems of aliasing (see §1.1.2 for details). Separation logic extends Hoare logic with spatial connectives, which allow assertions to define separation between parts of the heap. This separation provides the key feature of separation logic—*local reasoning*—specifications need only mention the state that they access. This local reasoning approach has proved itself successful when considering some real-world algorithms, including the Schorr-Waite graph marking algorithm [88] and a copying garbage collector [11].

Until now, all the work on separation logic has focused on simple, imperative, low-level languages. In this chapter we present an extension to separation logic to allow reasoning about Java-like languages, that is languages with features that are prevalent in the real-world including encapsulation and inheritance. It is known that separation logic solves the problems caused by aliasing. However a naïve adaptation of this approach falls short of reasoning about encapsulation and inheritance. Solutions to these issues will be presented in Chapters 4 and 6.

The rest of the chapter is structured as follows. First we present a simplification of MJ that is well suited to separation logic. In §3.2 we describe the storage model we will reason about, and then, in §3.3, we give the syntax and semantics of separation logic assertions. In §3.4 we present the Hoare logic rules to reason about Java, and then in §3.5, we present an additional rule to reason about dynamically dispatched method calls. We then give an example of copying a tree in §3.6, and present a full semantic account of the logic and its soundness in §3.7. We conclude, in §3.8, by discussing encapsulation and inheritance.

3.1 Inner MJ

In order to allow clean and elegant separation logic rules, we make a few simplifications to MJ, and call this subset *Inner MJ*. The first and most important simplification is that statements are not allowed multiple indirections: a statement can only perform a single read or write of a field. Additionally, for compactness, we remove `void` methods. The simplifications to the syntax are presented in Figure 3.1; the original syntax is given on page 20.

```

Method definition
md ::= C m(C1 x1, ..., Cn xn){ $\bar{s}$  return x;}
Expressions
e ::= x | null
Statements
s ::= x = y.f; | x = (C)y; | x = new C();
    | x.f = e; | x = y.m( $\bar{e}$ ); | C x;
    | { $\bar{s}$ } |; | if (e == e){ $\bar{s}$ } else { $\bar{s}$ }

```

Figure 3.1: Syntax of Inner MJ

Additionally, we make a semantic restriction on method and constructor bodies that they do not modify their parameters. This reduces the complexity of reasoning about method and constructor calls. This may at first sound restrictive but simply assigning each parameter to a new local variable allows modification. This restriction removes the need to consider the call by value arguments to a method.

Although we have simplified the syntax, we have not removed any computational power. A simple transformation from full MJ to Inner MJ is given in Figure 3.2. We write $(e)_z$ for the transformation on expressions to statements. This replaces e by a series of small statements which evaluate the expression. The final statement stores the result in z . We then define the transformation on statements (s) . This evaluates each expression, stores its result in a variable, and then evaluates the statement for these variables. The transformation on a program simply transforms all the method and constructor bodies. In Lemma 3.7.17 on page 61 we prove this translation preserves program semantics.

3.2 Storage model

Recent work on separation logic has allowed reasoning about pointer datastructures with address arithmetic [65, 75]. This includes complicated doubly linked lists where only the XOR of the forwards and backwards pointers is stored in the heap. This reasoning is powerful and interesting, but it does not represent the programmer's view of memory in Java: Java does *not* provide address arithmetic. Obviously, this model can be used to represent Java's memory by compiling field locations away, but representing the compilation in the logic seems an unnecessary burden on the proof. It would be analogous to reasoning about a high-level language by considering the assembly language to which it is compiled. The original work on separation logic [76, 48] considered the heap to be a partial finite function from locations to unlabelled records:

$$H : \text{Locations} \rightarrow_{fn} \text{Values}^+$$

where Locations correspond to oids from the previous chapter, and Values are oids and null.

Note: In particular many models consider a simplification to just pairs, $\text{Values} \times \text{Values}$, rather than unlabelled records, Values^+

In order to represent objects in this model, we can replace the tuple by a partial finite function from field names to values:

$$H : \text{Locations} \rightarrow_{fn} (\text{Fields} \rightarrow_{fn} \text{Values})$$

This model is similar to those used in object-oriented semantics [9, 33, 38]. It is useful to be able to split an object into its different aspects, e.g. a `ColourPoint` class has a location and a

Expressions

$$\begin{aligned}
\langle x \rangle_z &= z = x; \\
\langle \text{null} \rangle_z &= z = \text{null}; \\
\langle e.f \rangle_z &= \{D\ y; \langle e \rangle_y z = y.f;\} \\
\langle (C)\ e \rangle_z &= \{D\ y; \langle e \rangle_y z = (C)y;\} \\
\langle e.m(e_1, \dots, e_n) \rangle_z &= \{D\ y; D_1\ y_1; \dots; D_n\ y_n; \langle e \rangle_y (\langle e_1 \rangle_{y_1} \dots \langle e_n \rangle_{y_n} \\
&\quad z = y.m(y_1, \dots, y_n));\} \\
\langle \text{new } C(e_1, \dots, e_n) \rangle_z &= \{D_1\ y_1; \dots; D_n\ y_n; \langle e_1 \rangle_{y_1} \dots \langle e_n \rangle_{y_n} \\
&\quad z = \text{new } C(y_1, \dots, y_n);\}
\end{aligned}$$

Statements

$$\begin{aligned}
\langle ; \rangle &= ; \\
\langle e; \rangle &= \{D\ y; \langle e \rangle_y\} \\
\langle \text{if } (e_1 == e_2) \{ \bar{s}_1 \} \text{ else } \{ \bar{s}_2 \} \rangle &= \{D_1\ y_1; D_2\ y_2; \langle e_1 \rangle_{y_1} \langle e_2 \rangle_{y_2} \\
&\quad \text{if } (y_1 == y_2) \{ \langle \bar{s}_1 \rangle \} \text{ else } \{ \langle \bar{s}_2 \rangle \} \} \\
\langle e_1.f = e_2 \rangle &= \{D_1\ y_1; D_2\ y_2; \langle e_1 \rangle_{y_1} \langle e_2 \rangle_{y_2} y_1.f = y_2;\} \\
\langle x = e; \rangle &= \langle e \rangle_x \\
\langle \text{return } e; \rangle &= D\ y; \langle e \rangle_y \text{return } y; \\
\langle \{ \bar{s} \} \rangle &= \{ \langle \bar{s} \rangle \} \\
\langle C\ x; \rangle &= C\ x; \\
\langle s\ \bar{s} \rangle &= \langle s \rangle \langle \bar{s} \rangle \\
\langle \epsilon \rangle &= \epsilon
\end{aligned}$$

provided the static types of e, e_1, \dots, e_n are D, D_1, \dots, D_n respectively, and y, y_1, \dots, y_n are always fresh variables.

Figure 3.2: Transformation of expressions and statements

colour. We want to consider methods that modify the location property without having to reason about the colour property, hence it is important to be able to split an object. Later we will see that this simplifies the logic, as it allows the rules to be given with respect to just the fields used by a command. The frame rule can then be used to extend the specifications to the rest of the object, as well as to other objects. However, in this model it is difficult to define the composition of heaps such that an oid can be in both heaps as long as the two heaps have different fields for that oid.

By uncurrying the partial finite functions, we get a natural definition that allows splitting at the field level.

$$H : \text{Locations} \times \text{Fields} \rightarrow_{fn} \text{Values}$$

Finally, we store each object's type. This could be done using a special field. In order to allow sharing of the immutable type information we distinguish it as in the definition below.¹

¹Java does not allow the explicit disposal of an object and an object can never change type, i.e. the type of an object is immutable, so it does not matter who knows it. If one allows disposal of objects then sharing this information would be somewhat dangerous, and methods described in Chapter 5 should be employed for such sharing.

Definition 3.2.1 (Heap). A heap, H , is a pair of finite partial functions: the first stores the fields, H_v , and the second stores the types of objects, H_t .

$$H : \mathcal{H} \stackrel{\text{def}}{=} (\text{Locations} \times \text{Fields} \rightarrow_{\text{fin}} \text{Values}) \times (\text{Locations} \rightarrow_{\text{fin}} \text{Classes})$$

We write $H \perp H'$ to mean two heaps are disjoint and compatible:

$$(H'_v, H'_t) \perp (H''_v, H''_t) \stackrel{\text{def}}{=} \text{dom}(H'_v) \perp \text{dom}(H''_v) \wedge H'_t = H''_t$$

We write $H * H$ for the composition of two heaps:

$$(H'_v, H'_t) * (H''_v, H''_t) \stackrel{\text{def}}{=} \begin{cases} (H'_v \circ H''_v, H'_t) & (H'_v, H'_t) \cap (H''_v, H''_t) = \emptyset \\ \text{undefined} & \text{otherwise} \end{cases}$$

where \circ is union of disjoint partial functions. Finally we define an order on heaps \subseteq

$$H_1 \subseteq H_2 \stackrel{\text{def}}{=} \exists H_3. H_1 * H_3 = H_2$$

Note: This definition is the same as in the previous chapter. Classes is the set of all class names.

Again we use the same shorthands for accessing the first and second components of the heap: $(H_v, H_t)(o) = H_t(o)$ and $(H_v, H_t)(o, f) = H_v(o, f)$. We use the obvious shorthands for *dom* and function update that can always be disambiguated by context, i.e. $(o, f) \in \text{dom}(H)$ checks that (o, f) is in the domain of the first component of H .

A stack is defined as in the previous chapter. We also use ghost² variables that do not interfere with program execution. Ghost variables are used to express protocols of method calls, so they require a global scope. We cannot simply use the program stack as this cannot define a variable across method calls: it only contains local variables. So we define a ghost stack that maps ghost variables, e.g. X, Y , etc. to values.

A logical state is a triple of the form (VS, H, I) where VS is a program variable stack, H is a heap and I is a ghost variable stack.

3.3 Assertion language

The key difference between the assertion languages of separation logic and Hoare logic is the introduction of the spatial, or multiplicative, conjunction $*$. The spatial conjunction $P * Q$ means the heap can be split into two disjoint parts in which P and Q hold respectively.

There are two flavours of separation logic: classical and intuitionistic [48]. In classical separation logic the spatial conjunction does not admit weakening or contraction, while in intuitionistic separation logic it admits weakening, but not contraction. Also unlike classical separation logic, intuitionistic separation logic does not have an empty heap predicate: *true* is the unit of both $*$ and \wedge .

In this thesis we will use the intuitionistic version. This is because Java is a garbage collected language, hence it does not have an explicit dispose primitive. State is deallocated when it is inaccessible. As intuitionistic separation logic admits weakening, the following holds $P * Q \Rightarrow Q$. This allows state to be ignored, hence we do not need to explicitly dispose state. If the language was not garbage collected, then intuitionistic separation logic would allow memory leaks.

²These variables are sometimes called auxiliary or logical variables.

$$\begin{aligned}
 VS, H, I \models \text{false} &\stackrel{\text{def}}{=} \text{never} \\
 VS, H, I \models \text{true} &\stackrel{\text{def}}{=} \text{always} \\
 VS, H, I \models P \wedge Q &\stackrel{\text{def}}{=} VS, H, I \models P \text{ and } VS, H, I \models Q \\
 VS, H, I \models P \vee Q &\stackrel{\text{def}}{=} VS, H, I \models P \text{ or } VS, H, I \models Q \\
 VS, H, I \models P \Rightarrow Q &\stackrel{\text{def}}{=} \forall H' \supseteq H. \text{ if } VS, H', I \models P \text{ then } VS, H', I \models Q \\
 VS, H, I \models \forall x.P &\stackrel{\text{def}}{=} \forall v \in \mathbf{Values}. VS[x \mapsto v], H, I \models P \\
 VS, H, I \models \exists x.P &\stackrel{\text{def}}{=} \exists v \in \mathbf{Values}. VS[x \mapsto v], H, I \models P \\
 VS, H, I \models e.f \mapsto e' &\stackrel{\text{def}}{=} H(\llbracket e \rrbracket_{VS, I}, f) = \llbracket e' \rrbracket_{VS, I} \\
 VS, H, I \models e : C &\stackrel{\text{def}}{=} H(\llbracket e \rrbracket_{VS, I}) = C \\
 VS, H, I \models e = e' &\stackrel{\text{def}}{=} \llbracket e \rrbracket_{VS, I} = \llbracket e' \rrbracket_{VS, I} \\
 VS, H, I \models P * Q &\stackrel{\text{def}}{=} \exists H_1, H_2. H_1 * H_2 = H, \quad VS, H_1, I \models P \text{ and } VS, H_2, I \models Q \\
 VS, H, I \models P \multimap Q &\stackrel{\text{def}}{=} \forall H_1. \text{ if } H_1 \perp H \text{ and } VS, H_1, I \models P \text{ then } VS, H_1 * H, I \models Q
 \end{aligned}$$

where $\llbracket X \rrbracket_{VS, I} = I(X)$, $\llbracket \text{null} \rrbracket_{VS, I} = \text{null}$ and $\llbracket x \rrbracket_{VS, I} = VS(x)$.

Figure 3.3: Semantics of the assertion language

However, Reynolds noted that intuitionistic separation logic does not correctly model garbage collected languages [76] as assertions can be made about inaccessible state. Consider the following code fragment:

```
x = new C(); x = null;
```

The constructed object is not accessible. In the programming logic, given the usual semantics [48], it is possible to prove it exists in the heap after the assignment of `null`. If we model the garbage collector operationally, then we can prove the object may have been garbage collected. Therefore our assertion language would be unsound. Calcagno and O’Hearn [21] present a solution to this problem where they alter the semantics of the existential quantifier to correctly model garbage collected languages. Fortunately, we can avoid these complications as we do not model the garbage collector operationally, and hence it is sound to remain in “vanilla” intuitionistic separation logic.

Formulae in our assertion language are given by the following grammar

$$\begin{aligned}
 P, Q ::= & \text{true} \mid \text{false} \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid \forall x.P \mid \exists x.P \\
 & \mid P * Q \mid P \multimap Q \mid e.f \mapsto e' \mid v : C \mid e = e
 \end{aligned}$$

Expressions, e , are extended from Figure 3.1 to allow ghost variables as well.

For intuitionistic logic we require all assertions to satisfy the monotonicity, or persistency, constraint.

Definition 3.3.1 (Monotonicity constraint). P satisfies the monotonicity constraint if whenever $VS, H, I \models P$ and $H \subseteq H'$, then $VS, H', I \models P$

The monotonic subsets of the powerset of states form a Heyting algebra. We provide the formal semantics of formulae in Figure 3.3. The intuitionistic connectives ($\vee, \wedge, \Rightarrow$) and quantifiers (\forall, \exists) are interpreted in the usual way [48]. In addition to the intuitionistic connectives we have the new spatial connectives $*$ and \multimap , along with the predicates $e = e'$, $e : C$ and $e.f \mapsto e'$. Taking these in reverse order: the predicate $e.f \mapsto e'$ consists of all the tuples (VS, H, I) where the heap, H , consists of at least the single mapping from the location given by the meaning of e and the field f to the value given by the meaning of e' . $e : C$ is true if the evaluation of e gives

$$\begin{aligned}
\text{mods}(x = \dots) &\stackrel{\text{def}}{=} \{x\} \\
\text{mods}(x.f = e;) &\stackrel{\text{def}}{=} \text{mods}(;) \stackrel{\text{def}}{=} \{\} \\
\text{mods}(C \ x; \bar{s}) &\stackrel{\text{def}}{=} \text{mods}(\bar{s}) \setminus x \\
\text{mods}(s\bar{s}) &\stackrel{\text{def}}{=} \text{mods}(s) \cup \text{mods}(\bar{s}) \\
\text{mods}(\{\bar{s}\}) &\stackrel{\text{def}}{=} \text{mods}(\bar{s}) \\
\text{mods}(\text{if}(e_1 == e_2)\{\bar{s}_1\} \text{else}\{\bar{s}_2\}) &\stackrel{\text{def}}{=} \text{mods}(\bar{s}_1) \cup \text{mods}(\bar{s}_2)
\end{aligned}$$

Figure 3.4: Modifies property for Inner MJ

a value of dynamic type³ C . $e = e'$ is true in any state where the expressions evaluate to the same value. The spatial conjunction $P * Q$ means the heap can be split into two disjoint parts in which P and Q hold respectively. Heaps of more than one element are specified by using the $*$ connective to join smaller heaps. Finally, the adjoint to $*$, written $P \multimap Q$, means that if the heap can be extended with any disjoint heap satisfying P then the resultant heap will satisfy Q .

The essence of “local reasoning” is that, to understand how a piece of code works, it should only be necessary to reason about the memory the code actually accesses (its so-called “footprint”). Ordinarily, aliasing precludes such a principle but the separation enforced by the $*$ connective allows this intuition to be captured formally by the following rule.

L-FRAME

$$\frac{\vdash \{P\}\bar{s}\{Q\}}{\vdash \{P * R\}\bar{s}\{Q * R\}}$$

provided \bar{s} does not modify the free variables of R : $\text{mods}(\bar{s}) \cap FV(R) = \emptyset$

The side-condition is required because $*$ only describes the separation of heap locations and not variables, without this restriction \bar{s} could modify stack variables in R altering R 's validity; see [13] for more details. We give the definition of $\text{mods}(_)$ in Figure 3.4.

By using this rule, a local specification concerning only the variables and parts of the heap that are used by \bar{s} can be arbitrarily extended as long as the extension's free variables are not modified by \bar{s} . Thus, from a local specification we can infer a global specification that is appropriate to the larger footprint of an enclosing program. Consider a function that operates on a list:

$$\{list(i)\}m(i)\{list(i)\}$$

This specification says that if provided a list, i , then if m returns it will leave i still satisfying $list$. Implicitly, anything not mentioned is unchanged by the call. Hence we can derive a specification

$$\{list(i) * P\}m(i)\{list(i) * P\}$$

where P represents the other state we might know about. We have gone from a local specification describing what we use, to a global specification that additionally describes the context: this means we do not need to provide a different specification for each call site.

3.4 Java rules

Now we provide the logic for reasoning about Java. A judgement in our assertion language is written as follows:

$$\Gamma \vdash \{P\}\bar{s}\{Q\}$$

³We use dynamic type to mean the class the object was created with, and static type for the type of an expression denoting that object. For example $C \ x = \text{new } D ();$ The static type of x is C but the dynamic type of x is D .

This is read: the statement, \bar{s} , satisfies the specification $\{P\}\bar{s}\{Q\}$, given the method and constructor hypotheses, Γ . These hypotheses are given by the following grammar:

$$\begin{aligned}\Gamma &:= \epsilon \mid \{P\}\eta(\bar{x})\{Q\}, \Gamma \\ \eta &:= C.m \mid C\end{aligned}$$

However, when it simplifies the presentation, we will treat Γ as a partial function from method and defining class name, $C.m$, and constructor names, C , to specifications. For the hypotheses, Γ , to be well-formed each method, $C.m$, and constructor, C , should appear at most once; and each specification's free program variables should be contained in the method's arguments and *ret*. We will only consider well-formed Γ .

First we present the rules for field manipulation. We give small axioms in the style of O'Hearn, Reynolds, and Yang [65]. Field write requires the heap to contain at least the single field being written to, and the post-condition specifies it has the updated value. Implicitly the rule specifies no other fields are modified, and hence it can be extended by the frame rule to talk about additional state. The field access rule simply requires the state to be known and sets the variable equal to the contents. The ghost variables X and Y are used to allow x and y to be syntactically the same variable without needing a case split or substitution in the rule definition.

$$\begin{array}{l} \text{L-FWRITE} \qquad \qquad \qquad \Gamma \vdash \{x.f \mapsto _ \} x.f = y; \{x.f \mapsto y\} \\ \text{L-FREAD} \qquad \qquad \qquad \Gamma \vdash \{X = x \wedge X.f \mapsto Y\} y = x.f; \{X.f \mapsto Y \wedge y = Y\} \end{array}$$

Note: These rules highlight our choice to have fields assertions as the primitive assertion about the heap. The rules are given in terms of just the fields they use, and the frame rule dictates what happens to the other values. If the primitive assertions described entire objects, we would need to give the rules in terms of the whole object and say how the other fields were affected by this assignment.

Next we present the rules associated with types. The first ensures that a cast will complete successfully: the object identifier being cast must be a subtype of the target type or null. The L-STYPE rule allows static typing information to be exploited in the proof: if we can prove that it meets the specification for all subtypes and `null`, then it meets the specification without any typing information.

$$\begin{array}{l} \text{L-DOWNCAST} \qquad \qquad \qquad \Gamma \vdash \{P[x/y] \wedge (x : C \vee x = \text{null})\} y = (D)x\{P\} \\ \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \text{provided } C \prec D. \\ \text{L-STYPE} \qquad \qquad \qquad \frac{\Gamma \vdash \{P \wedge x = \text{null}\}\bar{s}\{Q\} \quad \forall C \prec D. \Gamma \vdash \{P \wedge x : C\}\bar{s}\{Q\}}{\Gamma \vdash \{P\}\bar{s}\{Q\}} \\ \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \text{provided } x \text{ has static type } D. \end{array}$$

Note: The L-STYPE proof rule is not local: it requires knowledge of all the classes that exist. Hence it will not lead to open proofs, i.e. proofs open to the extension of the class hierarchy. Later we will use this rule with a restriction on the class hierarchy to derive the rule for dynamic dispatch.

Next we give the rules for reasoning about methods. These are essentially the same as the standard function rules [75]. The method call rule states that if the pre-condition of $C.m$ holds and the receiver is of type C , then the post-condition will hold after the call. The L-RETURN rule allows the method's return value to be matched with the call's return variable using the special

variable *ret*. This command is logically equivalent to $ret = v;$: we give this rule to match the Java syntax.

L-CALL	$\{P\}C.m(\bar{x})\{Q\}, \Gamma \vdash \{P \wedge \text{this} : C\}ret = \text{this}.m(\bar{x}); \{Q\}$
L-RETURN	$\Gamma \vdash \{P[v/ret]\}return v; \{P\}$

The arguments, receiver and return variable of a method, and constructor, can all be specialized to arbitrary variables using the variable substitution rule, L-VARSUBST. The rule is similar to that given by Reynolds [75]. The first side-condition is the same as given by Reynolds, and prevents the substitution creating assignment to an expression. The second is because we make a distinction between ghost variables and program variables, and hence ghost variables cannot appear in the program text. The third restriction is because `return` implicitly uses a variable *ret*, but *ret* is not syntactically free in `return x;`,

L-VARSUBST	$\frac{\Gamma \vdash \{P\}\bar{s}\{Q\}}{\Gamma \vdash \{\theta(P)\}\theta(\bar{s})\{\theta(Q)\}}$
------------	---

provided $\theta = [e_1, \dots, e_n/x_1, \dots, x_n];$

- if x_i is modified by \bar{s} , then e_i is just a variable not free in any other e_j ;
- if x_i is in the free variables of \bar{s} , then e_i is restricted to program expressions, expressions not containing ghost variables; and
- if \bar{s} contains `return` then $\theta(ret) = ret$.

Note: We allow θ to alter `this`, as it simplifies the L-CALL rule.

Next, we give the rules for reasoning about constructors. Before we give the rules we define the shorthand $new(C, x)$ as meaning $x.f_1 \mapsto \text{null} * \dots * x.f_n \mapsto \text{null}$ where $dom(\delta_F(C)) = f_1, \dots, f_n$. The L-NEW rule is similar to the L-CALL rule with the additional post-condition giving the object's type. The first statement in a constructor definition is always a `super` call. We validate a `super` call in the same way as `new`, but the post-condition about the object's type is not added, and the pre-condition requires the fields to exist (and be initialized to `null`).

L-NEW	$\Gamma, \{P\}C(\bar{x})\{Q\} \vdash \{P\}\text{this} = \text{new } C(\bar{x}); \{Q \wedge \text{this} : C\}$
L-SUPER	$\Gamma, \{P\}D(\bar{x})\{Q\} \vdash \{P * \text{new}(D, \text{this})\}\text{super}(\bar{x}); \{Q\}$ provided the <code>super</code> call is in class C constructor and $C \prec_1 D$.

Now we give the rules for introducing assumptions about definitions of methods and constructors, given in Figure 3.5. We define a new judgement $\Gamma_1 \Vdash \Gamma_2$ to allow the introduction of mutually recursive methods. $\Gamma_1 \Vdash \Gamma_2$ means the bodies of the methods in Γ_2 can be verified assuming all the method calls satisfy the specifications in Γ_1 . The method introduction rule, L-DMETHOD, checks that the body, \bar{s} , meets the specification assuming it is invoked on the correct class. The constructor introduction rule, L-DCONSTR, verifies the body with both the pre-condition and the new fields for the object. It is not given the pre-condition $\text{this} : C$ as the constructor is also called by the `super` call. The remaining three rules are used to introduce and manipulate these definitions. These rules are similar in style to those used by von Oheimb [83] to reason about mutually recursive procedures.

Finally, the standard rules are presented in Figure 3.6. As we allow allocation of new variables at any point in a sequence, we require two sequencing rules. The first is for the case where the first command is not a variable allocation. The second is where a new variable is being created. This requires that the new variable does not get passed out of this sequence.

L-DMETHOD	$\frac{\Gamma \vdash \{P \wedge \text{this} : C\} \bar{s} \{Q\}}{\Gamma \Vdash \{P\} C.m(\bar{x}) \{Q\}}$	provided $mbody(C, m) = (\bar{x}, \bar{s})$
L-DCONSTR	$\frac{\Gamma \vdash \{new(C, \text{this}) * P\} \text{super}(\bar{e}); \bar{s} \{Q\}}{\Gamma \Vdash \{P\} C(\bar{x}) \{Q\}}$	provided $cnbody(C) = (\bar{x}, \text{super}(\bar{e}); \bar{s})$.
L-DSPLIT	$\frac{\Gamma \Vdash \Gamma_1 \quad \Gamma \Vdash \Gamma_2}{\Gamma \Vdash \Gamma_1, \Gamma_2}$	
L-DWEAKEN	$\frac{\Gamma \Vdash \Gamma_1}{\Gamma, \Gamma_2 \Vdash \Gamma_1}$	
L-DINTRO	$\frac{\Gamma, \Gamma' \Vdash \Gamma' \quad \Gamma, \Gamma' \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{P\} \bar{s} \{Q\}}$	

Figure 3.5: Definition rules

L-IF	$\frac{\Gamma \vdash \{P \wedge x = y\} \bar{s}_1 \{Q\} \quad \Gamma \vdash \{P \wedge \neg(x = y)\} \bar{s}_2 \{Q\}}{\Gamma \vdash \{P\} \text{if } (x == y) \{ \bar{s}_1 \} \text{ else } \{ \bar{s}_2 \} \{Q\}}$	
L-SKIP	$\Gamma \vdash \{P\}; \{P\}$	
L-BLOCK	$\frac{\Gamma \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{P\} \bar{s} \{Q\}}$	
L-SEQ1	$\frac{\Gamma \vdash \{P\}_{s_1} \{R\} \quad \Gamma \vdash \{R\}_{s_2 \dots s_n} \{Q\} \quad s_1 \neq Cx;}{\Gamma \vdash \{P\}_{s_1 s_2 \dots s_n} \{Q\}}$	
L-SEQ2	$\frac{\Gamma \vdash \{P\}_{s_1 \dots s_n} \{Q\} \quad x \notin FV(P) \cup FV(Q)}{\Gamma \vdash \{P\} C x; s_1 \dots s_n \{Q\}}$	
L-CONSEQUENCE	$\frac{P \Rightarrow P' \quad \Gamma \vdash \{P'\} C \{Q'\} \quad Q' \Rightarrow Q}{\Gamma \vdash \{P\} C \{Q\}}$	
L-VARELIM	$\frac{\Gamma \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{\exists x.P\} \bar{s} \{\exists x.Q\}}$	provided x is not free in \bar{s} .

Figure 3.6: Standard Rules

The whole program is verified by

$$\{true\}\text{Object}()\{true\} \vdash \{P\}\bar{s}\{Q\}$$

Note: There is an implicit context of the program π .

3.5 Dynamic dispatch

The method call rule, L-CALL, defined earlier does not account for dynamic dispatch. As explained in §1.1.1, the usual approach to dealing with dynamic dispatch is to use behavioural subtyping. We present a generalised notion of behavioural subtyping: *specification compatibility*.

Definition 3.5.1 (Specification compatibility). We define specification compatibility, $\vdash \{P_D\}\bar{s}\{Q_D\} \Rightarrow \{P_C\}\bar{s}\{Q_C\}$, iff for all \bar{s} if $\text{modifies}(\bar{s}) \subseteq \{ret\}$ then $\Gamma \vdash \{P_C\}\bar{s}\{Q_C\}$ is derivable from $\Gamma \vdash \{P_D\}\bar{s}\{Q_D\}$, that is a proof of $\Gamma \vdash \{P_C\}\bar{s}\{Q_C\}$ exists whose only assumption is $\Gamma \vdash \{P_D\}\bar{s}\{Q_D\}$.

The quantification over the statement sequence \bar{s} in the definition restricts the derivation to the structural rules: L-CONSEQUENCE, L-FRAME, L-VARELIM and L-VARSUBST. The frame rule, L-FRAME, can only introduce formula that do not mention *ret*. If the derivation only uses the rule L-CONSEQUENCE, specification compatibility degenerates to standard behavioural subtyping.

From now on in this thesis we will restrict method and constructor environments to *well-behaved environments*, that is where subtypes must have compatible specifications with their supertypes' specification.

Definition 3.5.2 (Well-behaved environments). For each $\{P_C\}C.m(\bar{x})\{Q_C\}$ in Γ , if $D \prec C$, then $\{P_D\}D.m(\bar{y})\{Q_D\}$ is in Γ and $\vdash \{P_D[\bar{x}/\bar{y}]\}\{Q_D[\bar{x}/\bar{y}]\} \Rightarrow \{P_C\}\{Q_C\}$

The substitution, $[\bar{x}/\bar{y}]$, in the definition is required to allow methods to use different argument names.

Specification compatibility generalises behaviour subtyping to allow variable manipulation. Consider the following specifications

$$\{\text{this}.f \mapsto _ \}\{ \text{this}.f \mapsto _ \} \tag{3.1}$$

$$\{\text{this}.f \mapsto X\}\{ \text{this}.f \mapsto X \} \tag{3.2}$$

If code expects the behaviour of (3.1) then clearly it would be fine with (3.2): it is a refinement. However, the standard behavioural subtyping implications would not allow this as $\text{this}.f \mapsto _ \Rightarrow \text{this}.f \mapsto X$ does not hold. Hence we must add rules to allow variable manipulation in behavioural subtyping.

Not only does specification compatibility generalize behavioural subtyping in its manipulation of variables, but it also allows specification intersections. Consider the following pair of specifications

Method	Pre-condition	Post-condition
C.m	P_C	Q_C
D.m	$(P_C \wedge X = 1) \vee (P_D \wedge X = 2)$	$(Q_C \wedge X = 1) \vee (Q_D \wedge X = 2)$

Note: We use the ghost variable X to allow us to encode the intersection of the specification $\{P_C\}\{Q_C\}$ with $\{P_D\}\{Q_D\}$. We assume P_C and Q_C do not contain X .


```

class Tree extends Object {
    Tree l; Tree r;

    Tree(Tree l, Tree r) {
        super();
        this.l = l;
        this.r = r;
    }

    Tree clone() {
        Tree l, r, l2, r2;
        l = this.l;
        if(l != null) l2 = l.clone();
        else l2 = null;

        Tree r = this.r;
        if(r != null) r2 = r.clone();
        else r2 = null;

        Tree nt = new Tree(l2, r2);
        return nt;
    }
}

```

Figure 3.7: Tree example source code.

□

In Section 3.8, we will show that specification compatibility is too restrictive for many programs, and in Chapter 6 we will address this.

3.6 Example: Tree copy

Now we illustrate the use of this logic with a simple example of a tree class and a clone method. The source code is given in Figure 3.7.

Before we can specify the clone method, we extend our logic slightly to reason about trees, by adding the following predicate to the logic. In the next chapter we will formally define recursive predicates.

$$tree(i) \stackrel{\text{def}}{=} i = \text{null} \vee \exists j, k. i.l \mapsto j * i.r \mapsto k * tree(j) * tree(k)$$

Using this predicate, we can then specify the constructor and clone method as

Method	Pre-condition	Post-condition
Tree(l, r)	$tree(l) * tree(r)$	$tree(\text{this})$
Tree.clone()	$tree(\text{this})$	$tree(\text{this}) * tree(\text{ret})$

Note: Although the specification does not prevent the clone method being called on a null objects, the L-CALL rule does.

The constructor is verified with its precondition and the fields defined by the class:

```

{tree(l) * tree(r) * this.l ↦ null * this.r ↦ null}
  {true}
  super()
  {true}
{tree(l) * tree(r) * this.l ↦ null * this.r ↦ null}

```

We verify the super call using the frame rule. The superclass is Object, and Object's constructor's specification is $\{true\} _ \{true\}$.

```

{tree(l) * tree(r) * this.l ↦ null * this.r ↦ null}
  {this.l ↦ null}
  this.l = l;
  {this.l ↦ l}
{tree(l) * tree(r) * this.l ↦ l * this.r ↦ null}

```

We verify the assignment to the `l` field using the frame rule and the assignment rule.

```
{tree(l) * tree(r) * this.l ↦ l * this.r ↦ null}
  {this.r ↦ null}
  this.r = r;
  {this.r ↦ r}
{tree(l) * tree(r) * this.l ↦ l * this.r ↦ r}
```

and the `r` field in the same way.

```
{tree(l) * tree(r) * this.l ↦ l * this.r ↦ r}
{tree(this)}
```

Finally, we use the rule of consequence to put the post-condition in as it appears in the specification.

We can verify the `clone` method as follows. First we verify the code for copying the left hand subtree.

```
{this.l ↦ j * tree(j)}
  {this.l ↦ j}
  l = this.l;
  {this.l ↦ l ∧ l = j}
{this.l ↦ l ∧ l = j * tree(j)}
{this.l ↦ l * tree(l)}
  {tree(l)}
  if(l != null)
    {tree(l) ∧ l ≠ null}
    l2 = l.clone();
    {tree(l) * tree(l2)}
  else
    {tree(l) ∧ l = null}
    {tree(l) ∧ l = null ∧ null = null}
    l2 = null;
    {tree(l) ∧ l = null ∧ l2 = null}
    {tree(l) * tree(l2)}
  {tree(l) * tree(l2)}
{this.l ↦ l * tree(l) * tree(l2)}
```

We can provide a similar proof for the right hand subtree. Using these proofs we can complete the proof as

```
{tree(this) ∧ this : Tree}
{this.l ↦ j * this.r ↦ k * tree(j) * tree(k)}
  Tree l2, r2, l, r;
{this.l ↦ j * this.r ↦ k * tree(j) * tree(k)}
  {this.l ↦ j * tree(j)}
  l = this.l;
  if(l != null)
    l2 = l.clone();
  else
    l2 = null;
  {this.l ↦ l * tree(l) * tree(l2)}
{this.l ↦ l * this.r ↦ k * tree(l) * tree(k) * tree(l2)}
  {this.r ↦ k * tree(k)}
  r = this.r;
  if(r != null)
    r2 = r.clone();
```

```

else
  r2 = null;
  {this.r ↦ r * tree(r) * tree(r2)}
{this.l ↦ l * this.r ↦ r * tree(l) * tree(r) * tree(l2) * tree(r2)}
{tree(this) * tree(l2) * tree(r2)}
  {tree(l2) * tree(r2)}
  Tree nt = new Tree(l2, r2);
  {tree(nt)}
{tree(this) * tree(nt)}
  return nt;
{tree(this) * tree(ret)}

```

Again the frame rule is used to allow the additional state to be preserved by the call, without any additional proof burden.

3.7 Semantics

Now we will relate the semantics of the previous chapter with the logic given in this chapter, and hence prove its soundness. The typing of configurations from the previous chapter would consider heaps containing objects with missing fields as invalid. However, the logic works explicitly with partially defined objects so we relax the conditions on a heap being well-formed, originally given on page 34. We no longer require that every field of an object exists, but still require it to point to the correct type if it does. This relaxed rule is as follows:

$$\text{TH-OBJECTOK} \quad \frac{\forall f \in \text{dom}(H(o, -)). \quad \delta, H \vdash H(o, f) : \delta_{\mathbb{F}}(H(o))(f)}{\delta, H \vdash o \text{ ok}}$$

This weakening of the type system leads to two additional reduction rules

$$\begin{array}{l} \text{E-FIELDACCESSFAIL} \quad (H, VS, o.f, FS) \rightarrow (H, VS, \mathbf{NSFE}, FS) \quad \text{provided } (o, f) \notin \text{dom}(H) \\ \text{E-FIELDWRITEFAIL} \quad (H, VS, o.f = v;, FS) \rightarrow (H, VS, \mathbf{NSFE}, FS) \quad \text{provided } (o, f) \notin \text{dom}(H) \end{array}$$

These exceptional states correspond to the `NoSuchFieldException` in Java.

To simplify the proofs we want to consider $(H, VS, \bar{s}, \square)$ as well-typed when \bar{s} is a method body, i.e. ends in `return`. This is forbidden by the type rules in the previous chapter. We modify the case for the empty framestack to accept any variable scoping.

$$\text{TF-STACKEMPTY} \quad \delta, H, VS \vdash \square : \tau \rightarrow \tau$$

and relax the definition of a completed configuration:

$$\text{compl}(\text{config}) \stackrel{\text{def}}{=} ; \exists H, VS, \zeta, v. (H, VS, \zeta, \square) = \text{config} \vee (H, VS, v, \square) = \text{config}$$

where ζ ranges over `return v`; and $;$.

The logic will be used to remove the runtime exceptions: **NSFE** and the null pointer and class cast exception from the previous chapter. We define safety, in the style of Yang [88], to mean a configuration cannot get to an exceptional state.

Definition 3.7.1 (Safety). A configuration is safe iff it cannot reduce to a stuck state that is not completed configuration, i.e.

$$\text{config safe} \stackrel{\text{def}}{=} \forall \text{config}'. (\text{config} \rightarrow^* \text{config}' \wedge \text{stuck}(\text{config}') \Rightarrow \text{compl}(\text{config}'))$$

Note: Safety is a stronger property than type safety enforced in the previous chapter: a safe program will not dereference null, or perform an invalid cast.

We then prove the standard properties, following Yang [88], of the operational semantics for the soundness of the frame rule. For compactness we use $(H, VS, F, FS) * H'$ as a shorthand for $(H * H', VS, F, FS)$.

Lemma 3.7.2 (Safety monotonicity). *If $config$ safe then for any disjoint extension H' the extended configuration is safe: $config * H'$ safe.*

Proof. Assume $config * H'$ reduces to an error state. Then $config$ must also reduce, possibly earlier, to an error state. But we know $config$ safe, so the original assumption must be wrong. \square

Lemma 3.7.3 (Heap locality). *If $(H_0, VS, \bar{s}, \square)$ safe and $(H_0 * H_1, VS, \bar{s}, \square) \rightarrow^* (H', VS', \zeta, \square)$ then $\exists H'_0. H'_0 * H_1 = H'$ and $(H_0, VS, \bar{s}, \square) \rightarrow^* (H'_0, VS', \zeta, \square)$.*

Proof. We prove

$$\begin{aligned} config \text{ safe} \wedge config * H_0 \rightarrow (H, VS, \bar{s}, FS) \\ \Rightarrow \exists H_1. config \rightarrow (H_1, VS, \bar{s}, FS) \wedge H_1 * H_0 = H \end{aligned}$$

by case analysis of the reduction relation. \square

In what follows we find it useful to use a subset of the reduction relation that satisfies certain properties.

Definition 3.7.4 (Predicated reduction). Assuming we have a reduction relation $config \rightarrow config$ and a predicate, ϕ , on a configuration, we define a restricted reduction relation where every configuration must satisfy a predicate as

$$\frac{\phi(config)}{config \xrightarrow{\phi}^* config} \quad \frac{\phi(config'') \quad config \xrightarrow{\phi}^* config' \quad config' \rightarrow config''}{config \xrightarrow{\phi}^* config''}$$

We find it particularly useful to induct over the recursive depth of a computation.

Definition 3.7.5 (Recursive depth). We define recursive depth, $rdepth$, as the number of return statements in a given framestack:

$$rdepth(FS) = \begin{cases} 0 & FS = \square \vee FS = \text{return } e; \circ \square \\ 1 + rdepth(FS') & FS = \text{return } e; \circ FS' \wedge FS' \neq \square \\ rdepth(FS') & FS = F \circ FS' \wedge F \neq \text{return } e; \end{cases}$$

We consider $\text{return } e; \circ \square$ to have a zero recursive depth as the `return` will remain on the stack as there is no where to return to. We extend this to configurations in the obvious way, and define a predicate $rdepth_n$ that holds if the depth is less than or equal to n .

Lemma 3.7.6 (Bounded depth reductions contains all finite reductions).

$config \rightarrow^* config'$ iff there exists an n such that $config \xrightarrow{rdepth_n}^* config'$.

Proof. No finite reduction sequence can generate an infinite depth frame stack: each reduction step can introduce at most one `return`. \square

Rather than proving properties about the reduction relation in all possible contexts, we will show that the context does not affect the reduction. First we define a relation for when a computation does not affect its context.

Definition 3.7.7 (Stack contains predicate). We define a predicate, $ends$, for checking a frametack ends in a particular frametack.

$$ends_{FS, VS} \stackrel{\text{def}}{=} \{(H, VS' \circ VS, CF, FS') \mid FS' = FS'' \circ FS \wedge VS' = MS \circ VS''\}$$

Lemma 3.7.8 (Context change).

$$(H, MS \circ [], \bar{s}, []) \rightarrow^* (H', MS' \circ [], \zeta, []) \iff (H, MS \circ VS, \bar{s}, FS) \xrightarrow{ends_{FS, VS}}^* (H', MS' \circ VS, \zeta, FS)$$

Proof. We prove the stronger fact:

$$(H, VS \circ VS_1, CF, FS \circ FS_1) \xrightarrow{ends_{FS_1, VS_1}}^* (H', VS' \circ VS_1, CF', FS'' \circ FS_1) \Rightarrow (H, VS \circ VS_2, CF, FS \circ FS_2) \xrightarrow{ends_{FS_2, VS_2}}^* (H', VS' \circ VS_2, CF', FS'' \circ FS_2)$$

By induction on the length of the reduction. The base case holds trivially. The inductive case requires

$$(H, VS \circ VS_1, CF, FS \circ FS_1) \rightarrow (H', VS' \circ VS_1, CF', FS' \circ FS_1) \Rightarrow (H, VS \circ VS_2, CF, FS \circ FS_2) \rightarrow (H', VS' \circ VS_2, CF', FS' \circ FS_2)$$

This can be shown by case analysis of the rules. \square

This lemma shows that if we consider the empty context, then that is equivalent to considering all computations that ignore their context.

Lemma 3.7.9 (Method reduction ignores context).

$$(H, MS \circ VS, \bar{s}, FS) \rightarrow^* (H'', VS'', CF'', FS'') \wedge \neg ends_{FS, VS}(H'', VS'', CF'', FS'') \Rightarrow \exists H', MS', \zeta. (H, MS \circ VS, \bar{s}, FS) \xrightarrow{ends_{FS, VS}}^* (H', MS' \circ VS, \zeta, FS)$$

Proof. We can ignore exceptional reduction as it is ruled out by the assumption.⁴ Consider the last state in the reduction that satisfies the $ends$ predicate. The only reduction rules that can leave the $ends$ relation are: E-RETURN, E-SKIP and E-SUB, but these all operate on states of the form $(H', MS' \circ VS, \zeta, FS)$. \square

By showing method reduction ignores its context we can see it is valid to only consider method reduction in the empty context.

Lemma 3.7.10 (Expression evaluation). $\llbracket e \rrbracket_{VS, I}$ respects the operational semantics,

$$(H, VS, e, []) \rightarrow^* (H, VS, v, []) \iff v = \llbracket e \rrbracket_{VS, I}$$

Proof. Trivial as e can only be `null` or a variable. \square

The rules given earlier used substitution to model assignment; next we link update and substitution semantically.

⁴If we consider exceptions to be caught, then this would not hold. We would need to extend ζ to exceptions.

Definition 3.7.11 (Substitution on a stack). Let θ be a substitution of the form: $[e_1, \dots, e_n/x_1, \dots, x_n]$. We define a relation between stacks

$$VS, I \prec_{\theta} VS', I' \stackrel{\text{def}}{=} \forall e. \llbracket e \rrbracket_{VS, I} = v \Leftrightarrow \llbracket \theta(e) \rrbracket_{VS', I'} = v$$

It is important to note $VS[x \mapsto \llbracket e \rrbracket_{VS, I}] \prec_{[e/x]} VS$ holds, that is this relation contains the update on stacks under the syntactic substitution used in the logic.

Lemma 3.7.12 (Substitution preserves truth). *If VS and I are related by θ to VS' and I' then $\theta(P)$ is satisfied by VS, H, I iff P is satisfied by VS', H, I' .*

Proof. We prove

$$VS, I \prec_{\theta} VS', I' \Rightarrow (VS, H, I \models P \Leftrightarrow VS', H, I' \models \theta(P))$$

by induction on the structure of P . Holds at leaves by definition of the relation. \square

The following lemma is used to prove L-VARSUBST is sound.

Lemma 3.7.13 (Substitution preserves reduction).

$$VS_1 \prec_{\theta} VS_2 \Rightarrow ((H, VS_1, \bar{s}, \llbracket \cdot \rrbracket) \text{ safe} \Leftrightarrow (H, VS_2, \theta(\bar{s}), \llbracket \cdot \rrbracket) \text{ safe}) \quad (3.5)$$

$$\begin{aligned} (H, VS_1, \bar{s}, \llbracket \cdot \rrbracket) \rightarrow^* (H', VS'_1, \zeta, \llbracket \cdot \rrbracket) \wedge VS_1 \prec_{\theta} VS_2 \\ \Rightarrow ((H, VS_2, \theta(\bar{s}), \llbracket \cdot \rrbracket) \rightarrow^* (H', VS'_2, \zeta, \llbracket \cdot \rrbracket) \wedge VS'_1 \prec_{\theta} VS'_2) \end{aligned} \quad (3.6)$$

$$\begin{aligned} (H, VS_2, \theta(\bar{s}), \llbracket \cdot \rrbracket) \rightarrow^* (H', VS'_2, \zeta, \llbracket \cdot \rrbracket) \wedge VS_1 \prec_{\theta} VS_2 \\ \Rightarrow ((H, VS_1, \bar{s}, \llbracket \cdot \rrbracket) \rightarrow^* (H', VS'_1, \zeta, \llbracket \cdot \rrbracket) \wedge VS'_1 \prec_{\theta} VS'_2) \end{aligned} \quad (3.7)$$

provided $\theta = [e_1, \dots, e_n/x_1, \dots, x_n]$;

- if x_i is modified by \bar{s} , then e_i is just a variable not free in any other e_j ;
- if x_i is in the free variables of \bar{s} , then e_i is restricted to program expressions, expressions not containing ghost variables; and
- if \bar{s} contains `return` then $\theta(\text{ret}) = \text{ret}$.

Proof. We can see that this holds for expressions from the definition of the substitution relation. Hence we can consider statements with only values and no expressions. The method call and constructor cases follow directly from Lemmas 3.7.8 and 3.7.9. The only case of interest is the update variable. We must show $VS_1[x \mapsto v] \prec_{\theta} VS_2[\theta(x) \mapsto v]$, which holds trivially. \square

Now we define the semantics of the proof system.

Definition 3.7.14. We write $\Gamma \models \{P\}\bar{s}\{Q\}$ to mean if the specifications in Γ are true of the methods and constructors in the program, then so is $\{P\}\bar{s}\{Q\}$, i.e.

$$\Gamma \models \{P\}\bar{s}\{Q\} \stackrel{\text{def}}{=} \models \Gamma \Rightarrow \models \{P\}\bar{s}\{Q\}$$

where

$$\models \Gamma \stackrel{\text{def}}{=} \forall \{P\}\eta(\bar{x})\{Q\} \in \Gamma. \models \{P\}\eta(\bar{x})\{Q\}$$

$$\models \{P\}C.m(\bar{x})\{Q\} \stackrel{\text{def}}{=} \models \{P \wedge \text{this} : C\}\bar{s}\{Q\}$$

where $mbody(C, m) = (\bar{x}, \bar{s})$

$$\models \{P\}C(\bar{x})\{Q\} \stackrel{\text{def}}{=} \models \{P * \text{new}(\text{this}, C)\}\bar{s}\{Q\}$$

where $cnbody(C) = (\bar{x}, \bar{s})$

$$\models \{P\}\bar{s}\{Q\} \stackrel{\text{def}}{=} \forall VS, H, I.$$

$$VS, H, I \models P \wedge \exists \tau (\delta \vdash (H, VS, \bar{s}, \llbracket \cdot \rrbracket) : \tau)$$

$$\Rightarrow \left(\begin{aligned} & (H, VS, \bar{s}, \llbracket \cdot \rrbracket) \text{ safe} \\ & \wedge ((H, VS, \bar{s}, \llbracket \cdot \rrbracket) \rightarrow^* (H, VS', \zeta, \llbracket \cdot \rrbracket) \Rightarrow (VS', H', I \uplus \llbracket \zeta \rrbracket \models Q)) \end{aligned} \right)$$

where $\llbracket \text{return } v; \rrbracket = \{\text{ret} \mapsto v\}$ and $\llbracket ; \rrbracket = \emptyset$.

Note: The semantics are defined in an implicit context of a program, π , and a class table, δ .

Now we come to the key contribution of this chapter: the soundness of the logic.

Theorem 3.7.15 (Soundness). *If $\Gamma \vdash \{P\}\bar{s}\{Q\}$, then $\Gamma \models \{P\}\bar{s}\{Q\}$*

Proof. As we have recursive method calls we must prove this initially by induction on the recursive depth [71, 83].⁵ We define

$$\begin{aligned} \models_n \{P\}\bar{s}\{Q\} &\stackrel{\text{def}}{=} \forall VS, H, I. \\ &VS, H, I \models P \wedge \exists \tau (\delta \vdash (H, VS, \bar{s}, []): \tau) \\ &\Rightarrow \left(\begin{array}{l} (H, VS, \bar{s}, []) \text{ safe} \\ \wedge ((H, VS, \bar{s}, []) \xrightarrow{rdepth_n} (H, VS', \zeta, []) \Rightarrow (VS', H', I \uplus [\zeta]) \models_{\Delta_f} Q) \end{array} \right) \end{aligned}$$

By Lemma 3.7.6, $\models \{P\}\bar{s}\{Q\}$ is equivalent to $\exists n. \models_n \{P\}\bar{s}\{Q\}$. Hence, we prove

$$\begin{aligned} \Gamma \vdash \{P\}\bar{s}\{Q\} &\Rightarrow \left(\begin{array}{l} (\forall n. \models_n \Gamma \Rightarrow \models_{n+1} \{P\}\bar{s}\{Q\}) \\ \wedge \models_0 \{P\}\bar{s}\{Q\} \end{array} \right) \\ \Gamma \Vdash \Gamma' &\Rightarrow (\forall n \models_n \Gamma \Rightarrow \models_{n+1} \Gamma') \wedge \models_0 \Gamma' \end{aligned}$$

by rule induction. This induction is valid by Lemma 3.7.6. The definition introduction rules, Figure 3.5, all follow directly from their definition. The frame rule follows from Lemmas 3.7.2 and 3.7.3 and the free variable restriction. The variable substitution rule follows from Lemma 3.7.13. The L-STYPE rule follows directly from the definition of a well-formed configuration. Sequencing follows from the type preservation proof and the usual argument. The L-CALL, L-NEW, L-SUPER rules all follow from induction on the recursive depth, and that they do not modify their parameters. \square

Finally in this section, we show that the transformation preserves the semantics, but first we must show that adding fresh variables does not affect the reduction.

Lemma 3.7.16 (Fresh variables or new blocks do not affect reduction). *Adding a fresh variable, or a fresh block of variables, does not affect expression or statement reduction, i.e. if*

$$(H, MS \circ [], CF, []) \rightarrow^* (H', MS' \circ [], \zeta, [])$$

where $MS = MS_a \circ BS \circ MS_b$, $MS' = MS'_a \circ BS \circ MS'_b$, and CF is a statement sequence or an expression.

then

$$(H, MS_1 \circ [], CF, []) \rightarrow^* (H', (MS'_1 \circ [], \zeta, []))$$

where $MS_1 = MS_a \circ BS[z \mapsto v] \circ MS_b$ and $MS'_1 = MS'_a \circ BS[z \mapsto v] \circ MS'_b$ and z not mentioned in CF .

and

$$(H, (MS_2 \circ \{\} \circ BS \circ MS_2) \circ VS, CF, []) \rightarrow^* (H', (MS'_1 \circ \{\} \circ BS' \circ MS'_2) \circ VS, \zeta, [])$$

where $MS_2 = MS_a \circ \{\} \circ BS \circ MS_b$ and $MS'_1 = MS'_a \circ \{\} \circ BS \circ MS'_b$ and if CF is of the form $s_1 \dots Cx; \dots s_2$ then MS_a is non-empty.

Note: Above we abuse the use of \circ to additionally mean stack append, in addition to cons.

⁵The $(H, VS, \bar{s}, [])$ **safe** does not depend on the recursive depth, so even programs that do not terminate are safe.

```

class Cell extends Object {
  Object cnts;
  ...
  void set(Object o) {
    this.cnts = o;
  }

  Object get() {
    Object temp;
    temp = this.cnts;
    return temp;
  }
}

class Recell extends Cell {
  Object bak;
  ...
  void set(Object o) {
    Object temp;
    temp = this.cnts;
    this.bak = temp;
    this.cnts = o;
  }
}

```

Figure 3.8: Source of Cell and Recell (Constructors omitted)

Proof. By induction on the structure of CF . □

Lemma 3.7.17 (Transformation preserves semantics). *The transformation defined in Figure 3.2 preserves the semantics: i.e.*

$$\pi \vdash (H, VS, \bar{s}, []) \rightarrow^* (H', VS', ;, []) \Rightarrow (\pi) \vdash (H, VS, (\bar{s}), []) \rightarrow^* (H', VS'_+, ;, [])$$

where VS_+ is either VS or VS extended with a single variable in the top frame.

Proof. We prove the following

$$\Phi(n) \wedge \Psi(n) \Rightarrow \Phi(n+1) \wedge \Psi(n+1)$$

where

$$\begin{aligned} \Phi(n) &\stackrel{\text{def}}{=} \forall H, VS, \pi, \bar{s}, H', VS'. \\ &\pi \vdash (H, VS, \bar{s}, []) \xrightarrow{rdepth_n}^* (H', VS', \zeta, []) \Rightarrow \\ &(\pi) \vdash (H, VS, (\bar{s}), []) \xrightarrow{rdepth_n}^* (H', VS'_+, \zeta, []) \end{aligned}$$

and

$$\begin{aligned} \Psi(n) &\stackrel{\text{def}}{=} \forall \pi, H, VS, e, H', VS', v, z. \\ &\pi \vdash (H, VS, e, []) \xrightarrow{rdepth_n}^* (H', VS', v, []) \Rightarrow \\ &(\pi) \vdash (H, VS + [z \mapsto -, C], (e)_z, []) \xrightarrow{rdepth_n}^* (H', VS' + [z \mapsto v, C], ;, []) \end{aligned}$$

where $((BS \circ MS) \circ VS) + [x \mapsto (v, C)]$ is a shorthand for $(BS[x \mapsto (v, C)] \circ MS) \circ VS$.

We can prove this by induction on the translation, and use Lemma 3.7.16 to deal with the variables and block scopes the translation introduces. □

3.8 Encapsulation and inheritance

Although this logic allows reasoning about programs that manipulate pointer datastructures, it does not provide support for abstraction (encapsulation) or inheritance. Consider the definition of the `Cell` class given in Figure 3.8. We can provide the obvious specifications to the `Cell` class's methods:

Method	Pre-condition	Post-condition
<code>Cell.set(o)</code>	$\text{this.cnts} \mapsto X$	$\text{this.cnts} \mapsto o$
<code>Cell.get()</code>	$\text{this.cnts} \mapsto X$	$\text{this.cnts} \mapsto X \wedge \text{ret} = X$

These specifications can easily be verified using the rules described earlier. However, the specifications are sensitive to changes in the implementation. Renaming the `cnts` field would alter all the proofs using the `Cell`, even if they only used the `get` and `set` methods. We are lacking the common abstractions afforded to us by object-oriented programming.

Worse still, we cannot provide support for inheritance. Consider the standard subtype of `Cell`, `Recell` [1], also given in Figure 3.8. Consider giving its `set` method the obvious specification.

Method	Pre-condition	Post-condition
<code>Recell.set(o)</code>	$\text{this.cnts} \mapsto X$ $* \text{this.bak} \mapsto _$	$\text{this.cnts} \mapsto o$ $* \text{this.bak} \mapsto X$

This code is clearly a well-behaved subtype of `Cell` yet it is not specification compatible with its parent, because `Recell.set` has a footprint of two fields, while `Cell.set` has a footprint of one field. This is an example of the extended state problem mentioned in Chapter 1. The only way to gain specification compatibility is to modify the `Cell`'s specification to account for the `Recell`'s specification, i.e.

Method	Pre-condition	Post-condition
<code>Cell.set(o)</code>	$\text{this.cnts} \mapsto X$ $* (\text{this} : \text{Recell} \Rightarrow \text{this.bak} \mapsto _)$	$\text{this.cnts} \mapsto o$ $* (\text{this} : \text{Recell} \Rightarrow \text{this.bak} \mapsto X)$
<code>Recell.set(o)</code>	$\text{this.cnts} \mapsto X$ $* (\text{this} : \text{Recell} \Rightarrow \text{this.bak} \mapsto _)$	$\text{this.cnts} \mapsto o$ $* (\text{this} : \text{Recell} \Rightarrow \text{this.bak} \mapsto X)$

Adapting specifications in this way is perhaps feasible in small programs, but does not yield a scalable solution. In particular it would not be possible for open programs, i.e. programs where only part of the source is available.

The remainder of this thesis presents solutions to the complex problems of encapsulation and inheritance. In the next chapter we add support for encapsulation to the logic using a novel concept of an abstract predicate. In Chapter 6 we extend this concept to additionally deal with inheritance. The result is, we claim, the first program logic that is expressive enough to reason about real-world code written in realistic object-oriented languages, such as Java.

4

Abstract predicates

In the previous chapter we presented a logic for Java. However, that logic was not powerful enough to express the abstraction that encapsulation provides to object-oriented programs. In this chapter we extend the logic from the previous chapter with reasoning that supports encapsulation.

Various researchers have proposed the enrichment of a program logic to view the data abstractly (as in data groups [54]), or the methods/procedures abstractly (as in method groups [52, 80]). In contrast, we propose to add the abstraction to the logical framework itself, by introducing the notion of an *abstract predicate*. Intuitively, abstract predicates are used like abstract data types. Abstract data types have a name, a scope and a concrete representation. Operations defined within this scope can freely exchange the datatype's name and representation, but operations defined outside the scope can only use the datatype's name. Similarly abstract predicates have a name and a formula. The formula is *scoped*: code verified inside the scope can use both the predicate's name and its body, while code verified outside the scope must treat the predicate atomically. We do not reason directly about Java access modifiers `private` and `protected`, but provide a more abstract notion of scoping.

In separation logic it is common to use inductively defined predicates to represent data types, for example §3.6. In essence we allow predicates to additionally encapsulate state and not just represent it. This gives us two key advantages: (1) the impact of changing a predicate is easy to define; and (2) by encapsulating state we are able to reason about *ownership transfer* [66].

The rest of this chapter is structured as follows. We begin by giving the proof rules associated with abstract predicates. In §4.2, we provide examples to illustrate the use and usefulness of abstract predicates. Then in §4.3 we give the semantics and prove the new rules are sound. We conclude with a discussion of related work.

Note: A different version of this work, which uses a simple language of functions rather than Java, was presented at POPL [67].

4.1 Proof rules

We extend the assertion language given in Chapter 3 with predicates. We write α to range over predicate names and use a function $arity()$ from predicate names to a natural number representing the predicate's arity.

We extend judgements to contain predicate definitions as follows:

$$\Lambda; \Gamma \vdash \{Q_1\} \bar{s} \{Q_2\}$$

This is read “the statements, \bar{s} , satisfy the specification $\{Q_1\} \bar{s} \{Q_2\}$, given the method and constructor hypotheses, Γ , and predicate definitions, Λ .” Predicate definitions are given by the following grammar:

$$\Lambda ::= \epsilon \mid (\alpha(\bar{x}) \stackrel{\text{def}}{=} P), \Lambda$$

For the predicate definitions, Λ , to be well-formed we require that each predicate, α , is contained at most once; the free variables of the body, P , are contained in the arguments, \bar{x} ; and P is a positive formula.¹ We will only consider well-formed Λ . When it simplifies the presentation, we will treat Λ as a partial function from predicate names to formulae. We define $\Lambda(\alpha)[\bar{e}]$ as $P[\bar{e}/\bar{x}]$ where Λ contains $\alpha(\bar{x}) \stackrel{\text{def}}{=} P$. Additionally, we extend the definition rules in Figure 3.5 (page 51) to include the predicate definitions.

Note: In this thesis we only use abstract predicates to represent inductively defined datastructures (such as lists and trees), so the definitions are all continuous. However, the rules and semantics do not require the definitions to be continuous, only that the definitions have a fixed point.

The exchange between definitions and names occurs in the following two axioms² concerning abstract predicates:

$$\begin{array}{l} \text{OPEN} \\ \text{CLOSE} \end{array} \quad \frac{}{\begin{array}{l} (\alpha(\bar{x}) \stackrel{\text{def}}{=} P), \Lambda \models \alpha(\bar{e}) \Rightarrow P[\bar{e}/\bar{x}] \\ (\alpha(\bar{x}) \stackrel{\text{def}}{=} P), \Lambda \models P[\bar{e}/\bar{x}] \Rightarrow \alpha(\bar{e}) \end{array}}$$

These axioms embody our intuition that if (and only if) an abstract predicate is in scope (that is, it is contained in Λ) then we can freely move between its name and its definition. These axioms are used in the rule of consequence, L-CONSEQUENCE.

$$\text{L-CONSEQUENCE} \quad \frac{\Lambda \models P \Rightarrow P' \quad \Lambda; \Gamma \vdash \{P'\} C \{Q'\} \quad \Lambda \models Q' \Rightarrow Q}{\Lambda; \Gamma \vdash \{P\} C \{Q\}}$$

Taking the analogy with abstract datatypes: the rule of consequence can be seen as a subtyping rule, and the axioms open and close are the coercions, or subtyping relation, allowed when the datatype is in scope.

Next we present the rule to enforce the scoping of predicate definitions:

$$\text{L-DINTROABS} \quad \frac{\Lambda', \Lambda; \Gamma, \Gamma' \Vdash \Gamma' \quad \Lambda; \Gamma, \Gamma' \vdash \{P\} \bar{s} \{Q\}}{\Lambda; \Gamma \vdash \{P\} \bar{s} \{Q\}}$$

provided P, Q, Γ and Λ' do not contain the predicate names in $\text{dom}(\Lambda')$; and $\text{dom}(\Lambda)$ and $\text{dom}(\Lambda')$ are disjoint.

This rule allows a class, or package, author to use the definition of some abstract predicates, yet the client can only use the abstract predicate names. The methods and constructors in Γ' are within the scope of the predicates defined in Λ hence verifying the the bodies can use the predicate definitions. The client code, \bar{s} , is *not* in the scope of the predicates, so it can only use

¹A positive formula is one where predicate names appear only under an even number of negations. This ensures that a fixed point can be found; this is explained in further detail in Definition 4.3.3 and Lemma 4.3.4.

²We present them as semantic implications as a complete proof theory for separation logic is not known.

the predicates atomically through the specifications in Γ' . The predicate names cannot occur in the conclusion's specification, $\{P\}_{-}\{Q\}$.

The side-conditions for this rule prevent both the predicates escaping the scope of the module, and repeated definitions of a predicate.

In fact, the previous definition introduction rule, L-DINTROABS, is a derived rule in our system. It is derived from the standard definition introduction rule and two new rules for manipulating abstractions. The first, L-ABSWEAK, allows the introduction of new definitions.

$$\text{L-ABSWEAK} \frac{\Lambda; \Gamma \vdash \{P\}C\{Q\}}{\Lambda, \Lambda'; \Gamma \vdash \{P\}C\{Q\}} \text{ provided } \text{dom}(\Lambda') \text{ and } \text{dom}(\Lambda) \text{ are disjoint}$$

The second, L-ABSELIM, allows any unused predicate to be removed.

$$\text{L-ABSELIM} \frac{\Lambda, \Lambda'; \Gamma \vdash \{P\}C\{Q\}}{\Lambda; \Gamma \vdash \{P\}C\{Q\}} \text{ provided the predicate names in } P, Q, \Gamma \text{ and } \Lambda \text{ are not in } \text{dom}(\Lambda').$$

We can derive the L-DINTROABS rule in the following way:

$$\text{L-DINTRO} \frac{\Lambda, \Lambda'; \Gamma, \Gamma' \Vdash \Gamma' \quad \text{L-ABSWEAK} \frac{\Lambda; \Gamma, \Gamma' \vdash \{P\}\bar{s}\{Q\}}{\Lambda, \Lambda'; \Gamma, \Gamma' \vdash \{P\}\bar{s}\{Q\}}}{\text{L-ABSELIM} \frac{\Lambda, \Lambda'; \Gamma \vdash \{P\}\bar{s}\{Q\}}{\Lambda; \Gamma \vdash \{P\}\bar{s}\{Q\}}}$$

4.2 Examples

To demonstrate the utility of abstract predicates, we shall present a couple of examples.

4.2.1 Connection pool

Our first example is a database connection pool. Constructing a database connection is generally an expensive operation, so this cost is reduced by pooling connections using the object pool design pattern [42, Chapter 5]. Programs regularly access several different databases, hence we require multiple connection pools and dynamic instantiation. The connection pool must prevent the connections being used after they are returned: ownership must be transferred between the client and the pool.

We assume a library constructor, `Conn(db)`, to construct a database connection. This routine takes a single parameter that specifies the database, and returns a handle to a connection. The specification of this constructor uses a predicate `conn` to represent the state of the connection.

```
{true} Conn(db) {conn(this,db)}
```

Note: In a more realistic implementation, such as JDBC [35], several arguments would be used to specify how to access a database. Additionally an instance of the Factory pattern would be used so that the correct database driver could provide the connection.

We use two classes, given in Figure 4.1, to represent the state of a connection pool: `ConnPool` represents the connection pool, and `ConnList` holds the connections. We reflect this in the logic by defining two abstract predicates: `cpool` and `clist`. The `cpool` predicate is used to represent a connection pool; and the `clist` predicate is used inside the `cpool` to represent a list of connection


```

{cpool(this, db) * conn(y, db) ∧ this : ConnPool}
{∃i. this.l ↦ i * this.db ↦ db * clist(i, db) * conn(y, db)}
  t = this.l;
{this.l ↦ t * this.db ↦ db * clist(t, db) * conn(y, db)}
  n = new CList(y, t);
{this.l ↦ t * this.db ↦ db * n.hd ↦ y * n.tl ↦ t * clist(t, db) * conn(y, db)}
  this.l = n
{this.l ↦ n * this.db ↦ db * n.hd ↦ y * n.tl ↦ t * clist(t, db) * conn(y, db)}
{this.l ↦ n * this.db ↦ db * clist(n, db)}
{cpool(this, db)}

```

In this proof the definitions from Λ of both *cpool* and *clist* are used with OPEN and CLOSE to give the following three implications

$$\begin{aligned}
cpool(this, db) &\Rightarrow \exists i. this.l \mapsto i * this.db \mapsto db * clist(i, db) \\
n.hd \mapsto y * n.tl \mapsto t * clist(t, db) * conn(y, db) &\Rightarrow clist(n, db) \\
this.l \mapsto n * this.db \mapsto db * clist(n, db) &\Rightarrow cpool(this, db)
\end{aligned}$$

These are used with the rule of L-CONSEQUENCE to complete the proof. It is essential that Λ is known in (B) for these implications to hold.

Next we present, and attempt to verify, a fragment of client code using the connection pool (branch (C)). It demonstrates both correct and incorrect usage, which causes the verification to fail. The example calls a function, *useConn*, that uses a connection.

```

{cpool(x, db)}
  y = x.getConn();
{cpool(x, db) * conn(y, db)}
  {conn(y, db)}
  useConn(y);
  {conn(y, db)}
{cpool(x, db) * conn(y, db)}
  x.freeConn(y);
{cpool(x, db)}
  useConn(y)
{??}

```

The client gets a connection from the pool, uses it and then returns it. However, after returning it, the client tries to use the connection. This command cannot be validated as the precondition does not contain the *conn* predicate. Even though this predicate is contained in *cpool*, the client is unable to expand the definition because it is out of scope. We do not have Λ in this branch of the proof. This illustrates how abstract predicates capture “ownership transfer”, a notion first coined by O’Hearn *et al.* [66]. The connection passes from the client into the connection pool stopping the client from directly accessing it, even though the client has a pointer to the connection.

It is essential to be able to create new instances of a datatype. We conclude this section with a second example client that demonstrates the instantiation of abstract predicates. A connection pool library wants many instances of a connection pool; generally one per database. This can be easily handled by constructing the required number of `ConnPool` instances. Assume we have two different databases, `db1` and `db2`.

```

{true}
  y = new ConnPool(db1);
{cpool(y, db1)}
  z = new ConnPool(db2);

```

```
{cpool(y, db1) * cpool(z, db2)}
```

This code creates two connection pools. The parameter prevents us returning the connection to the incorrect pool.

```
{cpool(y, db1) * cpool(z, db2)}
  x = z.getConnection();
{cpool(y, db1) * cpool(z, db2) * conn(x, db2)}
  y.freeConn(x)
{??}
```

The `freeConn` call can only be validated if $db1 = db2$.³

This example has illustrated that abstract predicates capture the notion of “ownership transfer”. We will compare our support for ownership transfer to theirs in §4.4.2.

4.2.2 Ticket machine

Our next example is a simple ticket machine. Consider the following specifications:

Function	Pre-condition	Post-condition
<code>getTicket()</code>	<i>true</i>	<i>Ticket(ret)</i>
<code>useTicket(x)</code>	<i>Ticket(x)</i>	<i>true</i>

To call `useTicket` you must have called `getTicket`; each usage consumes a ticket. Trying to use a ticket twice fails:

```
{true}
  x = getTicket();
{Ticket(x)}
  useTicket(x);
{true}
  useTicket(x);
{??}
```

The second call to `useTicket` fails, because the first call removed the *Ticket*.

Any client that is validated against this specification must use the ticket discipline correctly. In fact the module is free to define the ticket in any way, e.g. $Ticket(x) \stackrel{\text{def}}{=} true$. Although duplication of this ticket would be logically valid, $true * true \Leftrightarrow true$, the client does not know this, and hence cannot.

4.2.3 Aside: Permissions interpretation

To better understand the previous example we turn to the permissions interpretation of separation logic. O’Hearn [63] has recently given separation logic an ownership, or permissions, interpretation: \mapsto is the permission to read and write a memory location. (This has been extended by Bornat, Calcagno, O’Hearn and Parkinson [14] to separate read and total permissions; we will detail this in Chapter 5.) In the previous example, the *Ticket* predicate is the permission to call `useTicket` once. The call consumes this permission. Using the permissions interpretation of separation logic, abstract predicates allow modules to define their own permissions. The concept of ownership transfer can be seen as transferring permission to and from the client.

³Given the specification it is always valid to return a connection to a pool if it is to the correct database. A tighter specification could be given to restrict returning to the allocating pool. See the `malloc` example in §4.4.1.

4.3 Semantics

In the previous sections we have introduced informally the notion of abstract predicates and detailed a couple of examples to demonstrate their usefulness. In this section we formalize them precisely and show that the two abstract predicate rules are sound: L-ABSWEAK and L-ABSELIM.

First we define the semantics of an abstract predicate. We define semantic predicate environments, Δ , as follows:

$$\Delta : \prod \alpha : \mathcal{A}. (\text{Values}^{\text{arity}(\alpha)} \rightarrow \mathcal{P}(\mathcal{H}))_{\perp}$$

where \mathcal{A} is the set of predicate names. We use an indexed, or dependent, product to represent a function from predicate name to definition of the correct arity. Later we require the combination of disjoint Δ s, so the always *false* definition is different from an undefined predicate, hence we lift the lattice. The reader might have expected the use of $\mathcal{P}(\mathcal{H} \times \mathcal{S} \times \mathcal{I})$ where \mathcal{S} is the set of all stacks and \mathcal{I} is the set of all ghost stacks. However, this breaks substitution as the predicate can depend on variables that are not syntactically free.

The semantics of a predicate is as follows:

$$VS, H, I \models_{\Delta} \alpha(\bar{e}) \stackrel{\text{def}}{=} \alpha \in \text{dom}(\Delta) \wedge H \in (\Delta\alpha)[[\bar{e}]]_{VS, I}$$

We lookup the definition in Δ , evaluate the arguments with respect to the stack, VS , supply the arguments to the definition, and check the heap is contained in the result. The remaining semantics are defined the same as in §3.3 with the predicate environment added in the obvious way.

Lemma 4.3.1. *Semantic predicate environments form a complete lattice.*

Proof. The powerset gives us a complete lattice. Any function from a set to a complete lattice is also a complete lattice. Lifting a complete lattice is a complete lattice, and a product of complete lattices is a complete lattice. \square

The complete lattice has the following ordering

$$\Delta \sqsubseteq \Delta' \stackrel{\text{def}}{=} \forall \alpha. \forall \bar{v} : \text{Values}^{\text{arity}(\alpha)}. \Delta(\alpha) \neq \perp \Rightarrow \Delta(\alpha)(\bar{v}) \subseteq \Delta'(\alpha)(\bar{v})$$

and the least upper bound of this order is written \sqcup and defined as:

$$\bigsqcup_{i \in I} (\Delta_i) \stackrel{\text{def}}{=} \lambda \alpha. \lambda \bar{v} : \text{Values}^{\text{arity}(\alpha)}. \bigcup_{i \in I}^{\perp} (\Delta_i(\alpha)(\bar{v}))$$

$$\text{where } \bigcup_{i \in I}^{\perp} X_i \stackrel{\text{def}}{=} \begin{cases} \perp & \forall i \in I. X_i = \perp \\ \bigcup_{i \in I \wedge X_i \neq \perp} X_i & \text{otherwise} \end{cases}$$

Lemma 4.3.2. *Formulae only depend on the predicate names they mention, i.e. if Δ defines all the predicate names in P , and Δ and Δ' are disjoint, then*

$$\forall VS, H, I. \quad VS, H, I \models_{\Delta} P \Leftrightarrow VS, H, I \models_{\Delta \sqcup \Delta'} P$$

Proof. By induction on the structure of P . \square

Definition 4.3.3 (Positive and negative formulae).

$$\begin{array}{ll}
\text{pos}(P \wedge Q) = \text{pos}(P) \wedge \text{pos}(Q) & \text{neg}(P \wedge Q) = \text{neg}(P) \wedge \text{neg}(Q) \\
\text{pos}(P \vee Q) = \text{pos}(P) \wedge \text{pos}(Q) & \text{neg}(P \vee Q) = \text{neg}(P) \wedge \text{neg}(Q) \\
\text{pos}(P * Q) = \text{pos}(P) \wedge \text{pos}(Q) & \text{neg}(P * Q) = \text{neg}(P) \wedge \text{neg}(Q) \\
\text{pos}(e = e) = \text{true} & \text{neg}(e = e) = \text{true} \\
\text{pos}(e.f \mapsto e) = \text{true} & \text{neg}(e.f \mapsto e) = \text{true} \\
\text{pos}(P * Q) = \text{neg}(P) \wedge \text{pos}(Q) & \text{neg}(P * Q) = \text{pos}(P) \wedge \text{neg}(Q) \\
\text{pos}(P \Rightarrow Q) = \text{neg}(P) \wedge \text{pos}(Q) & \text{neg}(P \Rightarrow Q) = \text{pos}(P) \wedge \text{neg}(Q) \\
\text{pos}(\alpha(\bar{v})) = \text{true} & \text{neg}(\alpha(\bar{v})) = \text{false}
\end{array}$$

Lemma 4.3.4. *Positive formulae are monotonic with respect to semantic predicate environments, i.e.*

$$\text{pos}(P) \wedge \Delta \sqsubseteq \Delta' \wedge VS, H, I \models_{\Delta} P \Rightarrow VS, H, I \models_{\Delta'} P$$

Proof. We prove the following stronger result

$$\begin{aligned}
\Delta \sqsubseteq \Delta' \quad \Rightarrow \quad & (\text{pos}(P) \wedge VS, H, I \models_{\Delta} P \Rightarrow VS, H, I \models_{\Delta'} P) \\
& \wedge (\text{neg}(P) \wedge VS, H, I \models_{\Delta'} P \Rightarrow VS, H, I \models_{\Delta} P)
\end{aligned}$$

by induction on the structure of P . □

Now let us consider the construction of a semantic predicate environment from an abstract one, Λ . The abstract predicate environment does not necessarily define every predicate, so constructing a solution requires additional semantic definitions, Δ , to fill the holes. We use the following function to generate a fixed point:

$$\text{step}_{(\Delta, \Lambda)}(\Delta') \stackrel{\text{def}}{=} \lambda \alpha \in \text{dom}(\Lambda). \lambda \bar{n} \in \mathbb{N}^{\text{arity}(\alpha)}. \{H \mid VS, H, I \models_{\Delta' \sqcup \Delta} \Lambda(\alpha)[\bar{n}]\}$$

where Λ contains the definitions we want to solve; and Δ are the predicates not defined in Λ . step is monotonic on predicate environments, because of Lemma 4.3.4 and that all the definitions are positive. Hence by Tarski's theorem and Lemma 4.3.1 we know a least fixed point always exists. We write $\llbracket \Lambda \rrbracket_{\Delta}$ for the least fixed point⁴ of $\text{step}_{\Delta, \Lambda}$.

Consider the set of all solutions of Λ of the form $\Delta \sqcup \llbracket \Lambda \rrbracket_{\Delta}$:

$$\text{close}(\Lambda) \stackrel{\text{def}}{=} \{\Delta \sqcup \llbracket \Lambda \rrbracket_{\Delta} \mid \text{dom}(\Delta) = \mathcal{A} \setminus \text{dom}(\Lambda)\}$$

Next we consider a property of fixed points, and a couple of properties of this function.

Lemma 4.3.5. *Consider two monotonic functions: $f_1 : A \times B \rightarrow A$ and $f_2 : A \times B \rightarrow B$ where A and B are complete lattices. Consider the simultaneous fixed point (a, b) :*

$$\text{fix}(x, y).(f_1(x, y), f_2(x, y)) = (a, b) \tag{4.1}$$

and the fixed point with one argument provided:

$$\text{fix}(x).(f_2(a', x)) = b'(a') \tag{4.2}$$

The fixed points are equal: $b'(a) = b$

Note: This Lemma is actually a corollary of Bekič's theorem [85, Page 163]. We present the following simplified proof for completeness.

⁴We could generalise the semantics to consider all possible fixed points. We use least fixed points to simplify the semantics.

Proof. We know b is a solution to fixed point equation in (4.2) from (4.1) as $f_1(a, b) = a$. Therefore $b'(a) \subseteq b$. Now prove $(a, b'(a))$ is a pre-fix point of equation in (4.1). This requires

$$(f_1(a, b'(a)), f_2(a, b'(a))) \subseteq (a, b'(a)) \quad (4.3)$$

The second element of the pair follows from (4.2). We know $b'(a) \subseteq b$ and f_1 monotone, so $f_1(a, b'(a)) \subseteq f_1(a, b)$. From (4.1) we know $f_1(a, b) = a$. Hence, proving the inequality, (4.3) for the first element of the pair. Therefore $(a, b'(a))$ is a prefix point of 4.1, but (a, b) is the least prefix point, so $b \subseteq b'(a)$. Hence they are equal. \square

Lemma 4.3.6. *Adding new predicate definitions refines the set of possible semantic predicate environments, i.e.*

$$close(\Lambda_1) \supseteq close(\Lambda_1, \Lambda_2)$$

Proof. Consider an arbitrary element of $close(\Lambda_1, \Lambda_2)$: $\llbracket \Lambda_1, \Lambda_2 \rrbracket_{\Delta} \sqcup \Delta = \Delta \sqcup \Delta_1 \sqcup \Delta_2$, $dom(\Delta_1) = dom(\Lambda_1)$ and $dom(\Delta_2) = dom(\Lambda_2)$. First show $\llbracket \Lambda_1 \rrbracket_{\Delta \sqcup \Delta_2} = \Delta_1$. This follows directly from Lemma 4.3.5 by defining $f_1(a, b) \stackrel{\text{def}}{=} step_{\Lambda_1; \Delta \sqcup a}(b)$ and $f_2(a, b) \stackrel{\text{def}}{=} step_{\Lambda_2; \Delta \sqcup b}(a)$. From these definitions we get $fix(x).f_1(x, \Delta_2) = \llbracket \Lambda_1 \rrbracket_{\Delta \sqcup \Delta_2}$ and $fix(x, y).(f_1(x, y), f_2(x, y)) = (\Delta_1, \Delta_2)$. So by the Lemma 4.3.5 we know $\llbracket \Lambda_1 \rrbracket_{\Delta \sqcup \Delta_2} = \Delta_1$, therefore $\llbracket \Lambda_1, \Lambda_2 \rrbracket_{\Delta} \sqcup \Delta = \llbracket \Lambda_1 \rrbracket_{\Delta \sqcup \Delta_2} \sqcup \Delta \sqcup \Delta_2 \in close(\Lambda_1)$. Therefore every element is contained in $close(\Lambda_1)$. \square

Lemma 4.3.7. *The removal of predicate definitions does not affect predicates that do not use them. Given Λ which is disjoint from Λ' and does not mention predicate names in its domain; we have*

$$\forall \Delta \in close(\Lambda). \exists \Delta' \in close(\Lambda, \Lambda'). \Delta \upharpoonright dom(\Lambda') = \Delta' \upharpoonright dom(\Lambda')$$

where $f \upharpoonright S$ is $\{a \mapsto b \mid a \mapsto b \in f \wedge a \notin dom(S)\}$

Proof. We only need consider $\alpha \in dom(\Lambda)$. Consider $\Delta_1 \sqcup \Delta_2 = \llbracket \Lambda, \Lambda' \rrbracket_{\Delta}$ where $dom(\Delta_1) = dom(\Lambda)$ and $dom(\Delta_2) = dom(\Lambda')$. Using Lemma 4.3.2 we get $\forall \Delta'. \llbracket \Lambda \rrbracket_{\Delta \sqcup \Delta'} = \llbracket \Lambda \rrbracket_{\Delta} = \llbracket \Lambda \rrbracket_{\Delta \sqcup \Delta_2}$. By the proof of the previous Lemma we know

$$\forall \alpha \in dom(\Lambda). \llbracket \Lambda \rrbracket_{\Delta \sqcup \Delta'}(\alpha) = \llbracket \Lambda \rrbracket_{\Delta \sqcup \Delta_2}(\alpha) = \llbracket \Lambda, \Lambda' \rrbracket_{\Delta}(\alpha)$$

Therefore, for every element in $close(\Lambda)$, there is an element in $close(\Lambda, \Lambda')$ that agrees on all the definitions not in $dom(\Lambda')$. \square

We define validity with respect to an abstract predicate environment, written $\Lambda \models P$, as follows:

$$\forall VS, H, I, \Delta \in close(\Lambda). VS, H, I \models_{\Delta} P$$

Theorem 4.3.8. *OPEN and CLOSE are valid, i.e.*

$$\begin{aligned} (\alpha(\bar{x}) \stackrel{\text{def}}{=} P), \Lambda \models \alpha(\bar{e}) &\Rightarrow P[\bar{e}/\bar{x}] \\ (\alpha(\bar{x}) \stackrel{\text{def}}{=} P), \Lambda \models P[\bar{e}/\bar{x}] &\Rightarrow \alpha(\bar{e}) \end{aligned}$$

Proof. Directly from definitions. \square

We are now in a position to define a semantics for our reasoning system. We write $\Lambda; \Gamma \models \{P\} \bar{s} \{Q\}$ to mean that, if every specification in Γ holds for the methods and constructors, and every abstract predicate definition in Λ is true of a predicate environment, then so is $\{P\} \bar{s} \{Q\}$:

$$\forall \Delta \in close(\Lambda). \Delta \models \Gamma \Rightarrow \Delta \models \{P\} \bar{s} \{Q\}$$

The rest of the definitions are modified in the obvious way to pass the abstract predicate environment around.

Given this definition we can show that the two new rules for abstract predicates are sound.

Command	Description	Rule
$x = [e]$	read heap location e	$\frac{\left\{ \begin{array}{l} e = Z \\ \wedge Z \mapsto Y \end{array} \right\} x = [e] \left\{ \begin{array}{l} Z \mapsto Y \\ \wedge x = Y \end{array} \right\}}{\text{Rule}}$
$[e] = e'$	write e' to heap location e	$\{e \mapsto _ \} [e] = e' \{e \mapsto e'\}$
$\text{dispose } e$	dispose of heap location e	$\{e \mapsto _ \} \text{dispose } e \{ \text{empty} \}$
$\text{while}(B) \bar{s}$	while B holds execute \bar{s}	$\frac{\{R \wedge B\} \bar{s} \{R\}}{\{R\} \text{while}(B) \bar{s} \{R \wedge \neg B\}}$
$\text{newvar } x \text{ in } \bar{s}$	new local variable x	$\frac{\{P\} \bar{s} \{Q\}}{\{P\} \text{newvar } x \text{ in } \bar{s} \{Q\}}$ where $x \notin FV(P) \cup FV(Q)$

Figure 4.2: Details required for Malloc and free

Theorem 4.3.9. *Abstract weakening, L-ABSWEAK, is sound.*

Proof. Direct consequence of definition of judgements and Lemma 4.3.6. □

Theorem 4.3.10. *Abstract elimination, L-ABSELIM, is sound.*

Proof. Follows from Lemmas 4.3.2 and 4.3.7. □

4.4 Discussion

We conclude this chapter by considering, in detail, an example of a memory manager and three related works: the hypothetical frame rule, higher-order separation logic, and tpestates.

4.4.1 Malloc and free

In this subsection we present an example of a memory manager that allocates and deallocates variable sized blocks of memory. We present the example for two reasons: it demonstrates that abstract predicates are applicable to non-object-oriented programming; and it allows us to present a better comparison with the hypothetical frame rule. (This example was suggested to us by O'Hearn.)

In this example we use a different model of separation logic, as we require disposal and address arithmetic. Recalling the discussion from §3.3 this requires a move from intuitionistic separation logic to classical separation logic to exclude weakening on the spatial connective, $*$: we do not want memory leaks. This alters the semantics of \mapsto to mean that the heap contains a single cell, rather than at least a single cell. We have an additional assertion *empty* for the empty heap. The heap is accessed by integers, rather than *oids* and fields, and we use e to range over integer expressions. We present a summary of the programming language in Figure 4.2. This language is similar to those used in most recent separation logic papers [65, 75, 67].

We also use some additional features for handling arrays, described by Reynolds [75]: the iterated separating conjunctions, $\otimes_{x=e_1}^{e_2}.P$; and a system routine `allocate` that allocates variable sized blocks. Intuitively, the iterated separating conjunction, $\otimes_{x=e_1}^{e_2}.P$, is the expansion $P[e_1/x] * \dots * P[e_2/x]$ where x ranges from e_1 to e_2 . If e_2 is less than e_1 , it is equivalent to *empty*. More formally its semantics are:

$$\begin{aligned}
VS, H, I \models_{\Delta} \textcircled{*}_{x=e_1}^{e_2}. P \stackrel{\text{def}}{=} & ([e_1]_{VS, I} = n_1 \wedge [e_2]_{VS, I} = n_2) \Rightarrow \\
& ((n_1 \leq n_2 \Rightarrow VS, H, I \models_{\Delta} P[n_1/x] * \textcircled{*}_{x=n_1+1}^{n_2}. P) \\
& \wedge (n_1 > n_2 \Rightarrow VS, H, I \models_{\Delta} \textit{empty}))
\end{aligned}$$

Returning to the example, consider the following naïve specifications, which demonstrate the difficulties in reasoning about a memory manager:

Function	Pre-condition	Post-condition
<code>malloc(n)</code>	<i>empty</i>	$\textcircled{*}_{i=0}^{n-1}. (ret + i \mapsto _)$
<code>free(x)</code>	$\textcircled{*}_{i=0}^{n-1}. x + i \mapsto _$	<i>empty</i>

`malloc` returns an array n cells long starting at `ret`, and `free` consumes an array n cells long starting at `x`. The problem is with `free`'s specification: it does not specify how much memory is returned as n is a free variable.

The standard specification [50] of `free` only requires it to deallocate blocks provided by `malloc`. It is undefined on all other arguments. Using abstract predicates we *are* able to provide an adequate specification.

Function	Pre-condition	Post-condition
<code>malloc(n)</code>	<i>empty</i>	$\textcircled{*}_{i=0}^{n-1}. (ret + i \mapsto _) * \textit{Block}(ret, n)$
<code>free(x)</code>	$\textcircled{*}_{i=0}^{n-1}. x + i \mapsto _ * \textit{Block}(x, n)$	<i>empty</i>

The *Block* predicate is used as a modular certificate, or permission, that `malloc` actually produced the block. The client cannot construct a *Block* predicate as its definition is not in scope.

Typical implementations of `malloc` and `free` store the block's size in the cell before the allocated block [50]. This can be specified directly by defining the *Block* predicate as follows.

$$\textit{Block}(x, n) \stackrel{\text{def}}{=} (x - 1) \mapsto n$$

This allows `free` to determine the quantity of memory returned. More complicated specifications can be used which account for padding and other bookkeeping.

We can give a simple implementations of both `malloc` and `free`, which call system routines to construct (`allocate`) and dispose (`dispose`) the blocks.⁵

```

malloc n = (newvar x in x=allocate(n+1);
           [x]=n; return x+1)
free x = (newvar n in n=[x-1];
         while(n≥0) (n=n-1; dispose(x+n))

```

Both of their implementations can be verified; here we present the proof of `malloc`:

```

{empty}
  x=allocate(n+1);
  { $\textcircled{*}_{i=0}^n. (x + i) \mapsto \_$ }
  { $x \mapsto \_ * \textcircled{*}_{i=1}^n. (x + i) \mapsto \_$ }
  [x]=n;
  { $x \mapsto n * \textcircled{*}_{i=1}^n. (x + i) \mapsto \_$ }
  return x+1
  {(ret - 1) \mapsto n *  $\textcircled{*}_{i=1}^n. (ret - 1 + i) \mapsto \_$ }
  {(ret - 1) \mapsto n *  $\textcircled{*}_{i=0}^{n-1}. (ret + i) \mapsto \_$ }
  { $\textcircled{*}_{i=0}^{n-1}. (ret + i) \mapsto \_ * \textit{Block}(ret, n)$ }

```

⁵One could extend the specification to have an additional memory manager predicate that contains a free list, as in the connection pool example.

The final step in this proof abstracts the cell containing the block's length, hence the client cannot directly access it. The following code fragment attempts (but fails) to break this abstraction:

```
{empty}
  x=malloc(30);
  {⊗i=029. (x+i) ↦ _ * Block(x,30)}
  [x-1]=15;
  {???}
  free(x);
```

The client attempts to modify the information about the block's size. This would be a clear failure in modularity as the client is dependent on the implementation of *Block*. Fortunately, we are unable to validate the assignment as the pre-condition does not contain $x - 1 \mapsto _$. Although the *Block* contains the cell, the client does not have the definition in scope and hence cannot use it.

This example has demonstrated that abstract predicates are applicable even in languages without direct support for abstraction. The C memory manager uses encapsulation, but this abstraction is not protected by the run-time or the language design. However, it can be protected at the logical level using abstract predicates. Abstract predicates provide encapsulation to languages that do not have encapsulation.

4.4.2 Hypothetical frame rule

The first attack on modularity in separation logic was made by O'Hearn, Reynolds and Yang [66]. They added *static* modularity to separation logic. They hide the internal resources of a module from its clients using the *hypothetical frame rule*.

In this section we present a detailed comparison between abstract predicates and the hypothetical frame rule. In particular we compare the different forms of modularity they allow.

O'Hearn, Reynolds and Yang extend separation logic with the following rule,⁶ known as hypothetical frame rule, to allow the hiding of information.

$$\frac{\Gamma, \{P_1\}k_1\{Q_1\}[X_1], \dots, \{P_n\}k_n\{Q\}[X_n] \vdash \{P\}\bar{s}\{Q\}}{\Gamma, \{P_1 * R\}k_1\{Q_1 * R\}[X_1, Y], \dots, \{P_n * R\}k_n\{Q * R\}[X_n, Y] \vdash \{P * R\}\bar{s}\{Q * R\}}$$

This rule can be used to derive a modular procedure definition rule, which they use to explain the modularity they enable.

$$\frac{\Gamma \vdash \{P_1 * R\}\bar{s}_1\{Q_1 * R\} \quad \vdots \quad \Gamma \vdash \{P_n * R\}\bar{s}_n\{Q_n * R\} \quad \Gamma, \{P_1\}k_1\{Q_1\}[X_1], \dots, \{P_n\}k_n\{Q_n\}[X_n] \vdash \{P\}\bar{s}\{Q\}}{\Gamma \vdash \{P * R\} \text{let } k_1 = \bar{s}_1, \dots, k_n = \bar{s}_n \text{ in } \bar{s} \text{ end}\{Q * R\}}$$

Note: The rule additionally has several side-conditions on free variables, and also on the form of R , which we discuss later.

In this rule, R is the hidden resource that is only available to the procedures, or functions, k_1, \dots, k_n . The parameterless procedures can modify the resource R , but the procedures must always restore R before they return. Returning to the motivating example of the memory manager, R could be a list of unallocated fixed size⁷ blocks of memory. `malloc` and `free` can

⁶Side-conditions omitted.

⁷The memory manager verified with the hypothetical frame rule can only allocate fixed size blocks of memory. It cannot be used with variable sized blocks as described earlier.

temporarily break this invariant but it must always hold when they return. The client code, \bar{s} , is verified without R , and hence is unable to *directly* modify the free list. Returning to the memory manager, a client should only modify the free list using `malloc` and `free`, it should not have direct access to this list. This rule describes the partitioning between the client's and the module's state.

This partitioning of resources allows “ownership transfer”: state can safely be transferred between the module and the client without fear of dereferencing dangling pointers. This allows reasoning about examples such as a simple memory manager, which allocates fixed size blocks of memory, and a queue.

Though this is a significant advance, their work is severely limited as it only models static modularity. Their modules are based on Parnas' work on information hiding [68], which deals with single instances of the hidden data structure: there is only one R and it is statically scoped. It cannot be used for many common forms of abstraction, including ADTs and classes, where we require multiple dynamically scoped instances of the hidden resource.

Abstract predicates can represent dynamically scoped abstractions. Their use captures “ownership transfer.” Our connection pool example captures all the key properties of O'Hearn *et al.*'s idealization of a memory manager. In addition in §4.4.1 we showed that abstract predicates could be used to represent modular permissions, or certificates, and hence verify a memory manager for variable sized blocks. O'Hearn, Reynolds and Yang's [66] idealization of a memory manager does not support variable sized blocks. Their specifications cannot be extended to cover this without exposing the representation of the block. Additionally, it is impossible for them to enforce that `malloc` must provide the blocks that `free` deallocates without extending the logic.

There are several side-conditions to the modular procedure definition rule:

1. \bar{s} does not modify variables in R , *except through using* k_1, \dots, k_n ;
2. Y is disjoint from P, Q, \bar{s} , and the context $\Gamma, \{P_1\}k_1\{Q_1\}[X_1], \dots, \{P_n\}k_n\{Q_n\}[X_n]$;
3. \bar{s}_i only modifies variables in X_i and Y ; and
4. R is precise.

The first three clauses restrict the use of variables. The variables X_1, \dots, X_n can appear in both the procedure specifications and the invariant, but cannot be modified by the client. These variables are required to specify the protocol, or behaviour, of the procedures. In addition global “hidden” variables Y are also used to represent the invariant's state. These variables all require careful restriction on how they can be modified.

Abstract predicates do not suffer these complications because they express behaviour using arguments to the predicates: these arguments can be seen by anyone, but can only be modified if the underlying definition is known. There is no need for a special modifies clause.

Clause 4 restricts the forms of invariants the modular procedure definition rule can use. Without these restrictions it is unsound. Consider the following counterexample, in classical separation logic, due to Reynolds. 1 represents a one element heap, 0 represents a zero element

heap.

$$\begin{array}{c}
\frac{\frac{\frac{\frac{\Gamma \vdash \{0 \vee 1\}k\{0\}}{\Gamma \vdash \{0\}k\{0\}}}{\Gamma \vdash \{0 \vee 1\}k\{0\}}}{\Gamma \vdash \{0 * 1\}k\{0 * 1\}}}{\Gamma \vdash \{1\}k\{1\}}}{\Gamma \vdash \{1 \wedge 1\}k\{1 \wedge 0\}}}{\Gamma \vdash \{1\}k\{\text{false}\}} \\
\frac{\frac{\frac{\frac{\Gamma \vdash \{0 \vee 1\}k\{0\}}{\Gamma \vdash \{0 * 1\}k\{0 * 1\}}}{\Gamma \vdash \{1\}k\{0\}}}{\Gamma \vdash \{1 * \text{true}\} \text{let } k = \text{skip in } k \text{ end}\{\text{true} * \text{false}\}}}{\Gamma \vdash \{1 * \text{true}\} \text{let } k = \text{skip in } k \text{ end}\{\text{false}\}}
\end{array}$$

where $\Gamma = \{0 \vee 1\}k\{0\}$.

In classical separation logic precise sizes of heaps can be expressed so $1 \wedge 0$ is false. Clearly the code above will terminate but the proof above shows it will not terminate. The logic is unsound!

The problem arises in the rule of conjunction:

$$\frac{\Gamma \vdash \{P\}\bar{s}\{Q\} \quad \Gamma \vdash \{P_2\}\bar{s}\{Q_2\}}{\Gamma \vdash \{P \wedge Q\}\bar{s}\{P_2 \wedge Q_2\}}$$

If we are using the information hiding outlined above, then we really have a hidden $* R$ attached to each pre- and post-condition. However, we cannot directly rewrite the proof rule with these. We must insert a use of the rule of consequence.

$$\frac{\frac{\frac{\Gamma \vdash \{P * R\}k\{Q * R\} \quad \Gamma \vdash \{P_2 * R\}k\{Q_2 * R\}}{\Gamma \vdash \{P * R \wedge P_2 * R\}k\{Q * R \wedge Q_2 * R\}}}{\Gamma \vdash \{(P \wedge P_2) * R\}k\{(Q \wedge Q_2) * R\}}}{\Gamma \vdash \{(P \wedge P_2) * R\}k\{(Q \wedge Q_2) * R\}} \tag{4.4}$$

This leaves us with two implications to prove:

$$\begin{aligned}
(P \wedge P_2) * R &\Rightarrow P * R \wedge P_2 * R \\
Q * R \wedge Q_2 * R &\Rightarrow (Q \wedge Q_2) * R
\end{aligned}$$

However, in classical separation logic, the second implication only holds for *precise* assertions [66].

A predicate P is *precise* if and only if for all states (VS, H, I) there is at most one subheap H_p of H for which $VS, H_p, I \models P$ holds.

For soundness, O'Hearn, Reynolds and Yang restrict their rule to use only these precise assertions.⁸

The soundness of abstract predicates is less subtle and has no need to restrict the forms of predicates; except to ensure recursive definitions exist.

To summarize, the hypothetical frame rule can only deal with static modularity: modularity with a single hidden resource. It cannot deal with dynamic modularity: modularity where the hidden resource can be dynamically allocated. Hence it cannot be used to reason about abstract datatypes or objects. When reasoning about static modularity the proofs are more compact and elegant with the hypothetical frame rule. However, all these proofs can be encoded into abstract predicates. We do not believe the converse is true: the static nature of the hypothetical frame rule prevents simple reasoning about abstract datatypes.

⁸In intuitionistic separation logic there is a different restriction called *supported* [66].

```

class Footballer {
  {true}
  Footballer() {...}
  {Healthy(this)}

  {Healthy(this)}
  play() {...}
  {Tired(this)}

  {Tired(this)}
  pub() {...}
}
{Aching(this)}
{Tired(this)}
warmDown() {...}
{Healthy(this)}
{Aching(this)}
longRest() {...}
{Healthy(this)}

```

Figure 4.3: Simple protocol example

4.4.3 Higher-order separation logic

In this chapter we have built a logic for reasoning about encapsulation and abstract types. It is well-known that abstract types are related via the Curry-Howard correspondence to existential types [58]. Abstract predicates appear to be analogous, but at the level of the propositions themselves. In recent unpublished work, Birkedal and Torp-Smith [10] show how to use existential quantifiers to model abstract predicates. They use the higher-order separation logic of Biering, Birkedal and Torp-Smith [8] and present the connection pool example given earlier. In addition to modelling abstract predicates with existentially quantified predicates, they also show how to use universal quantifiers to model parametric datatypes.

Similarly Reddy [73] has extended specification logic [78] to allow existential quantification over predicates. This quantification allows him to represent the encapsulation of state, and hence objects and abstract datatypes.

4.4.4 Typestates

Finally we discuss typestates as they allow similar styles of proofs to those presented in this chapter.

Fähndrich and DeLine [36] have developed a simple method for protocol checking in class-based languages. This work is based on Typestates of Strom and Yemini [81]. Typestates were initially developed to help compiler writers in formalizing checking for programming errors such as accessing uninitialized variables. Each variable has a Typestate associated to it, e.g. $x : (int \times \text{uninitialized}), y : (int \times \text{initialized})$. The Typestate can be altered by operations, e.g. assigning an initialized variable to an uninitialized variable alters the Typestate to initialized.

Fähndrich and DeLine extend this to allow the programmer to express protocols for a class. We illustrate the protocol ideas with a simple example of a footballer and how certain actions affect the footballer's state. We give the specifications in Figure 4.3 and present a state diagram in Figure 4.4. Note the dotted transition in the state diagram is not represented in the specification. Unfortunately this specification might leave the average footballer unhappy as he would not be able to go to the pub without playing football and ending up *Aching*. We could change the specification to the `pub()` method to be:⁹

```

{ Tired(this)  $\vee$  Healthy(this) }
pub() {...}
{ Aching(this) }

```

⁹Not shown in Figure 4.4.




Figure 4.4: State diagram for protocol

This is still unsatisfactory as we really want the footballer to be able to go to the pub if healthy and still be healthy afterwards. Fähndrich and DeLine’s logic cannot express this, because they cannot express dependencies between input and output states. Our logic is able to express this through the use of ghost variables

$$\{(X = 1 \wedge Tired(\text{this})) \vee (X = 2 \wedge Healthy(\text{this}))\}$$

pub() {...}

$$\{(X = 1 \wedge Aching(\text{this})) \vee (X = 2 \wedge Healthy(\text{this}))\}$$

The logic developed in this thesis is more expressive than Fähndrich and DeLine, but their aims are different: their logic is designed to automatically check protocols, rather than to prove correctness of programs. Their logic also deals with inheritance in a similar way to that developed in Chapter 6, but they give more structure to the Typestates than we give to abstract predicate families, which allows them to more easily inherit method bodies, but restricts the potential changes a subclass can make. This is discussed in more detail in § 6.5.

5

Read sharing

In the previous chapter we presented abstract predicates and showed their utility in reasoning about encapsulation. However, abstract predicates necessarily prevent sharing between data-structures.¹ If a predicate represents some state, then no other disjoint predicate can represent the same piece of state. This is required to prevent updates to one data-structure affecting another’s representation (the *raison d’être* of separation logic!).

However, it is perfectly reasonable to have several data-structures with *read* access to the same state, for example list iterators. A list iterator requires access to the internal structure of a list: several iterators can validly traverse a single list as long as they do not modify it. To reason about this we need shared read access.

In this chapter we will extend separation logic to allow shared read access, and explain this extension using the permissions interpretation of separation logic, given in §4.2.3. In the previous chapter we saw how abstract predicates could be used to generate modular permissions: the permission to use a method or function. The permissions interpretation liberates us from seeing the field predicate, $e.f \mapsto e'$, as simply describing a part of the heap, but instead as the permission to read and write to that part of the heap. With this interpretation in mind it is reasonable to consider other predicates that give us weaker permissions, in particular, read-only permissions.

In this chapter, we introduce a “read-only points-to” predicate, \multimap . We must ensure that a write permission and a separate read permission cannot exist to the same cell, otherwise our model would be unsound. The classic concurrent example of multiple readers and a single writer [25] keeps a count of the number of readers to the resource. We will mirror this in the logic.

One problem Bornat *et al.* [14] identified with counting read-only permissions is that $*$ can no longer be used to express disjointness. In particular, the obvious definition of a read-only tree admits directed acyclic graphs. We address this problem by considering a new model of separation logic which uses sets of named permissions.

The rest of the chapter is organised as follows. In §5.1 we present a sound logic for read-only assertions, and in §5.2 demonstrate its use. In §5.3 we show that $*$ no longer expresses disjointness, and present an extension to the logic that solves this problem. In §5.4, we present an abstract model that embodies both permissions idioms, and then in §5.5 we present the instantiations of these models and show the named model properly contains the counting model. We

¹This is not completely true because one can use standard classical conjunction, but this removes much of the benefit of using separation logic.

conclude by demonstrating some other idioms of permission accounting that can be expressed in the named permissions model.

Collaboration This chapter is based on joint work with Richard Bornat, Cristiano Calcagno, and Peter O’Hearn. The author discovered the counting model independently of Calcagno, O’Hearn and Bornat. This led to a collaborative publication about read-only permissions in a non-object-oriented setting [14]. This chapter extends that published work by demonstrating how read-only permissions can be used with datatypes, and solves the open problem of combining disjointness and shared access. The presentation of the abstract semantics, in §5.4, is due to Calcagno.

5.1 Counting

In this section we present a simple extension that allows the counting of read-only permissions. The counting logic mirrors the programming paradigm of reference counting by keeping a count of the number of read permissions to a field.

This model uses two types of permissions: source permissions, $e.f \mapsto^n e'$, and read-only permissions, $e.f \rightarrow e'$. We define a source permission $e.f \mapsto^n e'$ as a permission that has issued n read permissions to $e.f$. This definition gives rise to the following axiom:

$$e.f \mapsto^n e' \Leftrightarrow e.f \rightarrow e' * e.f \mapsto^{n+1} e'$$

If we have a source permission that has issued n read permissions, it is equivalent to a source permission that has issued $n + 1$ read permissions and a read permission. This axiom allows the source permission to dispense and collect read permissions to the cell. We have a distinguished source permission written $e.f \mapsto^0 e'$, which represents a total, or write, permission. We use the shorthand $e.f \mapsto e'$ for such a total permission.

There is an additional axiom that is required for some proofs:

$$e.f \mapsto^0 e' * e.f \rightarrow e'' \Leftrightarrow \text{false}$$

This prevents us having a total permission and a read permission. It also prevents us having two source permissions to the same field.

We generalise the field read axiom to only require a read permission to the field. The other axioms are unchanged as they require total permissions.²

$$\text{L-FREAD} \quad \Gamma \vdash \{x = X \wedge X.f \rightarrow Y\} z = x.f; \{X.f \rightarrow Y \wedge z = Y\}$$

Note: As we are working in an intuitionistic setting we can have permissions leaks: state can no longer be modified because part of the permission has been lost. This may at first seem unsettling but it allows us to define an immutable location. $\exists n. e.f \mapsto^n e'$ is a location that can never be modified as we do not know how many permissions have been removed.

5.2 Example: List iterators

In this section we present our motivating example of a list class with iterators. An iterator has a constructor and two methods, `hasNext` and `next`. The `hasNext` method returns `true` if there

²One might say the other axioms have all the occurrences of $e.f \mapsto e'$ replaced with $e.f \mapsto^0 e'$. Our shorthand makes this syntactic substitution.


```

class Iter {
  List n;
  Iter(List x) {
    this.n = x;
  }
  boolean hasNext() {
    return this.n!=null;
  }
  Object next() {
    List hd = this.n.hd;
    List temp = this.n.tl;
    this.n = temp;
    return hd;
  }
}

class List {
  Object Hd;
  List tl;
  ...
  Iter iterator() {
    return new Iter(this);
  }
  ...
}

```

Figure 5.1: List iterator source

Method	Pre-condition	Post-condition
<code>x.iterator()</code>	$list(x, n)$	$list(x, n + 1) * iter(ret, x)$
<code>x.hasNext()</code>	$iter(x, y)$	$(ret = true \wedge iter_b(x, y))$ $\vee (ret = false \wedge iter(x, y))$
<code>x.next()</code>	$iter_b(x, y)$	$iter(x, y)$
<code>Iter(x)</code>	$list_R(x)$	$iter(this, x)$

Table 5.1: Specifications for iterator methods

are more elements to be examined, and `next` returns the next element. We give the source code in Figure 5.1 and its specification in Table 5.1. The implementation uses `booleans` to aid clarity; this can easily be encoded into Inner MJ. The specifications use two predicates to represent an iterator: $iter$ is just the datastructure for an iterator, and $iter_b$ is the same with a constraint that it has a current element: the current node is not `null`. The predicates are defined as:

$$\begin{aligned}
 iter(x, y) &\stackrel{\text{def}}{=} \exists z. x.n \mapsto z * list_R(z) * (list_R(z) \multimap list_R(y)) \\
 iter_b(x, y) &\stackrel{\text{def}}{=} \exists z. x.n \mapsto z * list_R(z) * (list_R(z) \multimap list_R(y)) \wedge z \neq \text{null}
 \end{aligned}$$

The second predicate is used to prevent clients trying to access the next element when it does not exist.

Note: The use of two predicates represents the calling protocol of the class, and mirrors the work on `Typestates` [36] mentioned in §4.4.4.

Now we define the predicates for representing a list:

$$\begin{aligned}
 list(i, n) &\stackrel{\text{def}}{=} (\exists k, l. i.hd \xrightarrow{n} k * i.tl \xrightarrow{n} l * list(l, n)) \vee i = \text{null} \\
 list_R(i) &\stackrel{\text{def}}{=} (\exists k, l. i.hd \rightarrow k * i.tl \rightarrow l * list_R(l)) \vee i = \text{null}
 \end{aligned}$$

The n parameter is used to count the number of passive readers of the list. $list_R$ is the read-only list predicates that allows multiple iterators over the same list.³

We have a lemma about these list definitions: it allows read-only list segments to be formed from a list source.

³In what follows we will adopt the convention that a subscript R on predicates represents read-only state.

Lemma 5.2.1. $list(i, n) \Leftrightarrow list_R(i) * list(i, n + 1)$

Proof. By induction on the length of the list segment. □

Next we present proofs that the constructor, `Iter()`, and the methods, `iterator()` and `hasNext()` meet their specifications. First we consider `Iter()`:

```
{ListR(x) * this.n ↦ null}
  this.n = x;
{ListR(x) * this.n ↦ x}
{ListR(x) * (ListR(x) -* ListR(x)) * this.n ↦ x}
{iter(this, x)}
```

This proof follows directly from the definitions of `-*` and `iter`. We can prove the `iterator` method with Lemma 5.2.1 as follows.

```
{list(this, n)}
{list(this, n + 1) * listR(this)}
  return new Iter(this);
{list(this, n + 1) * iter(ret, this)}
```

Finally we give the proof for `hasNext`:

```
{iter(this, y)}
{∃z. this.n ↦ z * listR(z) * (listR(z) -* list(y))}
  Object t = this.n;
{this.n ↦ t * listR(t) * (listR(t) -* listR(y))}
  return t != null;
{ret = (t ≠ null) ∧ this.n ↦ t * listR(t) * (listR(t) -* listR(y))}
{((ret = true ∧ t ≠ null) ∨ (ret = false ∧ t = null))}
{  * this.n ↦ t * listR(t) * (listR(t) -* listR(y))}
{(ret = true ∧ iterb(this, y)) ∨ (ret = false ∧ iter(this, y))}
```

Now we present an example of client code incorrectly using an iterator:

```
{list(1, 0)}
  Iter x;
{list(1, 0)}
  x = l.iterator();
{list(1, 1) * iter(x, 1)}
  {iter(x, 1)}
  if(x.hasNext())
  {iterb(x, 1)}
    o = x.next();
  {iter(x, 1)}
  else
  {iter(x, 1)}
    x.next();
  {??}
  {??}
{list(1, 1) * ??}
```

The client constructs an iterator, checks if it has a next element, but then attempts to move to the next element in both branches of the `if`. However, this cannot be validated as the `iterb` predicate is not in the pre-condition. The verification has prevented a null pointer exception without any additional runtime checks.

Now let us consider how we can modify a list. With Java's linked list implementation, changing a list and then using an iterator created before the change throws an exception. Here is an excerpt from the API:⁴

The iterators returned by [the list's] `iterator` and `listIterator` methods are fail-fast: if the list is structurally modified at any time after the iterator is created, in any way except through the Iterator's own `remove` or `add` methods, the iterator will throw a `ConcurrentModificationException`. Thus, in the face of concurrent modification, the iterator fails quickly and cleanly, rather than risking arbitrary, non-deterministic behavior at an undetermined time in the future.

We want the logic to prevent this exception statically. We could disallow updates if the list has iterators, e.g.

Method	Pre-condition	Post-condition
<code>List.reverse()</code>	$list(this, 0)$	$list(ret, 0)$

The specification stops us calling `reverse` if the list has any iterators. This is too restrictive: using an iterator after a modification is the problem, not their existence. Consider the following specification:

Method	Pre-condition	Post-condition
<code>List.reverse()</code>	$list(this, n) * $ $\otimes_{i=1}^n (\exists z. (iter(z, this) \vee iter_b(z, this)))$	$list(ret, 0)$

See §4.4.1 for an explanation of the iterated separating conjunction, $\otimes_{i=1}^n$.

Although the pre-condition implies $list(this, 0)$ the client must treat the predicates abstractly, so cannot know this.

This allows us to have iterators to the list, but at the point we call `reverse` all n iterators will be absorbed back into the list and hence cannot be used. Consider the following usage example

```

{list(x, 0)}
y = x.iterator();
{list(x, 1) * iter(y, x)}
z = y.hasNext();
{list(x, 1) * (z = true ∧ iter_b(y, x)) ∨ (iter(y, x) ∧ z = false)}
if( z )
  {list(x, 1) * iter_b(y, x)}
  {list(x, 1) *  $\otimes_{i=1}^1 (iter_b(y, x))$ }
  x = x.reverse();
  {list(x, 0)}
  w = y.next(); (1)
  {??}
else
  ;

```

When we reach (1) we will not be able to validate the function call as the `reverse` will have swallowed the iterator predicate. Pleasingly the static verification fails at the same point the actual libraries would throw an exception.

Note: We would also prevent a call to `hasNext` if the iterator has been invalidated, which does not throw an exception in the actual Java 1.4.1 implementation (but arguably should).

⁴<http://java.sun.com/j2se/1.4.2/docs/api/java/util/LinkedList.html>

5.3 Disjointness

So far we have seen a logic for read-only permissions and used it to reason about list iterators. Bornat *et al.* [14] have noted that this logic, and the other logics in [14], are not capable of expressing disjointness: a formula cannot *simply*⁵ enforce that a read-only tree is not a direct acyclic graph (DAG). Consider the following formula:

$$tree_r(i) \stackrel{\text{def}}{=} i = \text{null} \vee (i.l \rightarrow j * tree_r(j) * i.r \rightarrow k * tree_r(k))$$

From a single node we can generate several read only references to that node:

$$k.l \mapsto \text{null} * k.r \mapsto \text{null} \Rightarrow k.l \xrightarrow{2} \text{null} * k.r \xrightarrow{2} \text{null} * tree_r(k) * tree_r(k)$$

Clearly, if we have another node where both branches point to k , then we can construct a “tree” in the following way:

$$i.l \mapsto k * i.r \mapsto k * tree_r(k) * tree_r(k) \Rightarrow tree_r(i)$$

This is not a *tree*!

This problem arises as $*$ no longer represents disjoint heaps, or even disjoint permissions: we can have the same permission twice, i.e. two read permissions to a single cell. Any two read permissions on a field look identical and we cannot distinguish them.

To distinguish permissions we need to give them names. Consider a countably infinite set of permissions \mathbb{P} . The key idea is that we can use subsets of \mathbb{P} . Hence we can count by considering the size of sets, and disjointness by considering the names of the elements. Later we will see how to allow the division of sets. We annotate the points-to relation with non-empty subsets of \mathbb{P} : $e.f \xrightarrow{p} e'$. Any non-empty subset of \mathbb{P} is a read permission and the whole set \mathbb{P} is the total permission. Again we define a shorthand $e.f \mapsto e'$ for $e.f \xrightarrow{\mathbb{P}} e'$. Additionally, to use this logic we require some operations on sets: disjoint union, set difference, set cardinality and set equality; and also an existential over subsets of \mathbb{P} . The subsets, p , are semantically treated like ghost variables, and can be eliminated using L-VARELIM. We could redefine the list predicates given earlier, on Page 81, as:

$$\begin{aligned} list(i, n) &\stackrel{\text{def}}{=} (\exists k, l, p. i.\text{hd} \xrightarrow{\mathbb{P} \setminus p} k * i.\text{tl} \xrightarrow{\mathbb{P} \setminus p} l * list(l, n) * |p| = n) \vee i = \text{null} \\ list_R(i) &\stackrel{\text{def}}{=} (\exists k, l, p. i.\text{hd} \xrightarrow{p} k * i.\text{tl} \xrightarrow{p} l * list_R(l) * |p| = 1) \vee i = \text{null} \end{aligned}$$

We present a general encoding in §5.5.1 of the counting model into this model.

There are two axioms for permissions in the logic:

DISJOINT	$e.f \xrightarrow{p} e' * e.f \xrightarrow{p} e' \Leftrightarrow \text{false}$
SPLIT	$(p_1 \uplus p_2 = p) \Leftrightarrow (e.f \xrightarrow{p_1} e' * e.f \xrightarrow{p_2} e' \Leftrightarrow e.f \xrightarrow{p} e')$
	where p, p_1 and p_2 are all non-empty

The first, DISJOINT, means we cannot have the same permission set twice and the second, SPLIT, allows us to split and combine disjoint permission sets.

⁵Elaborate mechanisms where trees carry sets of locations could be used, but this solution could hardly be called *simple*.

Note: One might have expected DISJOINT to be defined as $e.f \xrightarrow{p} e' * e.f \xrightarrow{p'} e' \wedge p \cap p' \neq \emptyset \Leftrightarrow \text{false}$. This is derivable as

$$\begin{aligned} p \cap p' \neq \emptyset \wedge \xrightarrow{p \cap p'} * \xrightarrow{p \cap p'} &\Leftrightarrow \text{false} \\ \Rightarrow p \cap p' \neq \emptyset \wedge \xrightarrow{p \setminus p'} * \xrightarrow{p' \setminus p} * \xrightarrow{p \cap p'} * \xrightarrow{p \cap p'} &\Leftrightarrow \text{false} * \xrightarrow{p \setminus p'} * \xrightarrow{p' \setminus p} \\ \Rightarrow p \cap p' \neq \emptyset \wedge \xrightarrow{p} * \xrightarrow{p'} &\Leftrightarrow \text{false} \end{aligned}$$

Again the only programming axiom that needs changing is the field read axiom:

$$\text{L-FREAD} \quad \Gamma \vdash \{x = X \wedge X.f \xrightarrow{p} Y \wedge |p| \geq 1\} y = x.f \{X.f \xrightarrow{p} Y \wedge Y = y\}$$

Earlier we showed that the counting model cannot *simply* express read-only trees without admitting DAGs. However, in this model we can simply state that each node in the tree must have the *same* permission set. Consider the formula:

$$\text{tree}(i, p) \stackrel{\text{def}}{=} i = \text{null} \vee (\exists j, k. i.l \xrightarrow{p} j * i.r \xrightarrow{p} k * \text{tree}(j, p) * \text{tree}(k, p))$$

This definition cannot admit sharing. We can prove this by induction on the definition of *tree*. By induction we can see that $\text{tree}(x, p)$ can only contain locations with permission set p . Hence $\text{tree}(j, p) * \text{tree}(k, p)$ implies the locations representing $\text{tree}(j, p)$ and $\text{tree}(k, p)$ are disjoint. Hence, we cannot have sharing: only a tree.

5.4 Abstract semantics

In this section, we provide a semantic model of read-only permissions in separation logic. We use the abstract model of permissions due to Calcagno [14]. We begin by giving this model in terms of a partial commutative semi-group of permissions: $(M, *)$. We write m for an element of M and redefine heaps, from page 45, as

$$H : \mathcal{H} \stackrel{\text{def}}{=} (\text{Locations} \times \text{Fields} \rightarrow M \times \text{Values}) \times (\text{Locations} \rightarrow \text{Classes})$$

That is each field is represented by a value and an element from the permissions semi-group. We define composition on field values as

$$(m_1, v_1) * (m_2, v_2) \stackrel{\text{def}}{=} \begin{cases} (m_1 * m_2, v_1) & v_1 = v_2, \text{ and } m_1 * m_2 \text{ is defined.} \\ \text{undefined} & \text{otherwise} \end{cases}$$

and then extend this to value heaps as

- $H_{v_1} * H_{v_2}$ is defined iff $H_{v_1}(l, f) * H_{v_2}(l, f)$ is defined for all $(l, f) \in (\text{dom}(H_{v_1}) \cap \text{dom}(H_{v_2}))$
- $(H_{v_1} * H_{v_2})(l, f) \stackrel{\text{def}}{=} \begin{cases} H_{v_1}(l, f) & H_{v_2}(l, f) \text{ is undefined} \\ H_{v_2}(l, f) & H_{v_1}(l, f) \text{ is undefined} \\ H_{v_1}(l, f) * H_{v_2}(l, f) & \text{otherwise} \end{cases}$

We write $H_1 * H_2$ for composition of heaps, which is defined as:

$$(H_{v_1}, H_{t_1}) * (H_{v_2}, H_{t_2}) \stackrel{\text{def}}{=} \begin{cases} ((H_{v_1} * H_{v_2}), H_{t_1}) & H_{v_1} * H_{v_2} \wedge H_{t_1} = H_{t_2} \\ \text{undefined} & \text{otherwise} \end{cases}$$

E-FIELDACCESS	$(H, VS, o.f, FS) \rightarrow (H, VS, v, FS)$ where $(o, f) \in \text{dom}(H)$ and $H(o, f) = m, v$ and $m \in \mathbf{M}$.
E-FIELDWRITE	$(H, VS, o.f = v; , FS) \rightarrow (H', VS, ; , FS)$ where $(o, f) \in \text{dom}(H)$ and $H' = H[(o, f) \mapsto m_w, v]$ and $H(o, f) = m_w, v'$
E-NEW	$(H, VS, \text{new } C(\bar{v}), FS) \rightarrow (H', (BS \circ []) \circ VS, \bar{s}, (\text{return } o;) \circ FS)$ where $\text{cnbody}(C) = (\bar{x}, \bar{s}), \Delta_c(C) = \bar{C}, H' = H[o \mapsto C][(o, \bar{f}) \mapsto (m_w, \text{null})]$ $o \notin \text{dom}(H)$, and $BS = \{\text{this} \mapsto (o, C), \bar{x} \mapsto (\bar{v}, \bar{C})\}$

Figure 5.2: Alterations to operational semantics for read-only permissions

We define the semantics of the points-to relation as

$$VS, H, I \models e.f \xrightarrow{m} e' \stackrel{\text{def}}{=} H = \{(\llbracket e \rrbracket_{VS, I}, f) \mapsto (m, \llbracket e' \rrbracket_{VS, I})\} \wedge m \in \mathbf{M}$$

and define the logical $*$ using the new heap composition operator, $*$.

We define a distinguished element of \mathbf{M} , m_w that corresponds to a write permission such that

$$\forall m \in \mathbf{M}. m * m_w \text{ is undefined} \quad (5.1)$$

The semantics of the language must be altered to accomodate the permissions model. Only the rules for manipulating the heap require any changes: the new definitions of E-FIELDACCESS, E-FIELDWRITE and E-NEW are given in Figure 5.2.

The alterations to the semantics correspond to run-time checks that the relevant permissions are available. These new semantics with the constraint (5.1) allow us to show the heap locality, Lemma 3.7.3, and the safety monotonicity, Lemma 3.7.2, properties hold. Hence the frame rule is still sound.

Theorem 5.4.1. *Permissions model is sound.*

Proof. The field read rule is sound with these semantics, and so is the frame rule. The rest of the soundness proof remains unchanged. \square

5.5 Concrete models

Now we consider instantiating the abstract model for the two logics given earlier.

Counting model We can define the permissions semi-group for the counting model as: $(\mathbb{Z}, *_1)$ where

$$z_1 *_1 z_2 = \begin{cases} \text{undefined} & z_1 \geq 0 \wedge z_2 \geq 0 \\ \text{undefined} & (z_1 \geq 0 \vee z_2 \geq 0) \wedge z_1 + z_2 < 0 \\ z_1 + z_2 & \text{otherwise} \end{cases}$$

Note: We equate $\xrightarrow{-1}$ with the read permission \rightarrow .

Theorem 5.5.1. *The counting axioms are sound.*

$$\begin{aligned} e.f \xrightarrow{n} e' &\iff e.f \xrightarrow{-1} e' * e.f \xrightarrow{n+1} e' \\ e.f \xrightarrow{0} e' * e.f \xrightarrow{-1} e' &\iff \text{false} \end{aligned}$$

Proof. Directly from the definition of the semi-group. \square

Named model We can define the permissions semi-group for the named permissions model as: $(\mathcal{P}_{\text{non-empty}}(\mathbb{P}), \uplus)$, where $\mathcal{P}_{\text{non-empty}}(\mathbb{P})$ is the set of non-empty subsets of \mathbb{P} .

Additionally, we must add a few simple set operations to the syntax of formula, e.g. cardinality, set minus and quantification over sets, and provide them with the obvious semantics. We will not go into the rather tedious detail here.

Theorem 5.5.2. *DISJOINT and SPLIT are sound.*

Proof. Direct from the definition of the semi-group. □

5.5.1 Encoding the counting model in the named model.

Next, we show the semantics of the named model properly encodes the semantics of the counting model. To aid clarity we write, H^c , H^n , P^c , and P^n for heaps (H) in the counting and named models, and formulae (P) in the counting and named models respectively. We define the translation from terms in the counting logic to terms in the named logic as follows:

$$\begin{aligned} \llbracket e.f \stackrel{n}{\rightarrow} e' \rrbracket &\stackrel{\text{def}}{=} \exists p. e.f \stackrel{\mathbb{P} \setminus p}{\rightarrow} e' \wedge |p| = n \\ \llbracket e \rightarrow e' \rrbracket &\stackrel{\text{def}}{=} \exists p. e.f \stackrel{p}{\rightarrow} e' \wedge |p| = 1 \end{aligned}$$

The source permission with n missing read-only permissions is the set of permissions with n elements missing. The read-only permission is represented by a singleton set of permissions. We omit the obvious definition for the other predicates and connectives.

We must also define a translation from the named model's heaps to counting model's heaps. First, we define a predicate on the permissions in the named model. The predicate identifies the permissions that correspond to permission in the counting model:

$$\text{fin}(p) \stackrel{\text{def}}{=} \exists n. |p| = n \vee |\mathbb{P} \setminus p| = n$$

We restrict permissions to finite and cofinite sets as their cardinality can be represented by an integer. We then define a translation from the elements of the named permissions semi-group to the counting permissions semi-group.

$$\llbracket p \rrbracket \stackrel{\text{def}}{=} \begin{cases} -n & |p| = n \\ n & |\mathbb{P} \setminus p| = n \\ \text{undefined} & \text{otherwise} \end{cases}$$

Lemma 5.5.3. *Translation preserves the semi-group action for finite and cofinite permissions, that is*

$$n_1 *_1 n_2 = \llbracket p_3 \rrbracket \Rightarrow \exists p_1, p_2. \llbracket p_1 \rrbracket = n_1 \wedge \llbracket p_2 \rrbracket = n_2 \wedge p_1 \uplus p_2 = p_3 \quad (5.2)$$

$$\llbracket p_1 \rrbracket *_1 n_2 = n_3 \Rightarrow \exists p_2, p_3. \llbracket p_2 \rrbracket = n_2 \wedge \llbracket p_3 \rrbracket = n_3 \wedge p_1 \uplus p_2 = p_3 \quad (5.3)$$

$$\forall p_1, p_2 \in \text{fin}. p_1 \uplus p_2 \text{ is defined} \Rightarrow \llbracket p_1 \rrbracket *_1 \llbracket p_2 \rrbracket = \llbracket p_1 \uplus p_2 \rrbracket \quad (5.4)$$

Proof. The three properties can be shown from the definitions of the semi-group and using case analysis: (5.2) and (5.4) by case analysis on positive and negative n s; and for (5.4) by case analysis on finite and cofinite sets. □

We extend the restriction on permissions, $\text{fin}(\cdot)$, to heaps in the obvious way, and extend the translation on permissions to heaps in the obvious way. We can extend the previous lemma to heaps.

Theorem 5.5.4. *Encoding counting model into named model is sound and complete. For all formulae, P^c , in the counting model, the following holds for all heaps in the named model:*

$$\llbracket H^n \rrbracket = H^c \quad \Rightarrow \quad (VS, H^n, I \models \llbracket P^c \rrbracket) \quad \iff \quad VS, H^c, I \models P^c$$

Proof. By induction on the structure of P^c . The $*$ and $\neg*$ cases require the previous lemma, and \mapsto follows from encoding. Other cases follow trivially. \square

5.5.2 Other models

To conclude this chapter, we will discuss how the named permissions model can be used to represent other permissions idioms. Bornat *et al.* [14] identify two key idioms for the dissemination of read-only permissions: counting, given earlier, and splitting. In this subsection we first show how to represent splitting and then we how to properly combine splitting and counting.

Infinite splitting

Boylard [16] first proposed the idea of permissions with weights. The key idea is to split a permission into two pieces such that they sum to the same weight as the original. This idea of splitting is very useful in algorithms which divide work recursively, such as parallel tree substitution [14]. The key to this idiom is that splitting can occur to an unbounded depth: a permission can always be split into two permissions. This style of reasoning exists in the general model due to the following fact: any countably infinite set can be split into two countably infinite sets, i.e.

$$|p| = |\mathbb{N}| \quad \iff \quad \exists p_1, p_2. |p_1| = |p_2| = |\mathbb{N}| \wedge p = p_1 \uplus p_2$$

When combined with SPLIT this allows any infinite read permissions to be split into two infinite read permissions.

Note: This does not precisely model the fractional permissions of Bornat *et al.* [14], as they use rational numbers to name these permissions. This additional structure can be encoded into this model, if required, by considering a bijection between the permissions set, \mathbb{P} , and an interval in \mathbb{Q} . However, this does not present any interesting problems.

Combining splitting and counting

We conclude by showing how splitting and counting can be combined. We want to allow the following:

1. split permissions that have been counted;
2. split permissions that can count, i.e. split source permissions; and
3. count permissions that can count, i.e. count source permissions.

Notation: We use $X \leftrightarrow Y$ to mean a bijection exists from X to Y .

Counting removes finite sets from an infinite set: $\mathbb{P} \leftrightarrow \mathbb{P} + 1$; and splitting takes an infinite set and produces two infinite sets: $\mathbb{P} \leftrightarrow \mathbb{P} + \mathbb{P}$. Everything needs an infinite set, but counting only provides finite sets: we need to count infinite sets!

We can make the following two observations:

$$\begin{aligned} p \leftrightarrow p' &\Rightarrow p \times \mathbb{P} \leftrightarrow p' \times \mathbb{P} \\ \mathbb{P} &\leftrightarrow \mathbb{P} \times \mathbb{P} \end{aligned}$$

Hence, we can count infinite sets with

$$(\mathbb{P} \times \mathbb{P}) \leftrightarrow (\mathbb{P} \times \mathbb{P}) + (1 \times \mathbb{P})$$

Clearly, $1 \times \mathbb{P}$ is infinite so we can count permissions, that themselves can count or be split.

Now let us provide an example of how one might use these ideas for multilevel counting. We will write $prod_{\mathbf{p}}$ for a bijection from $\mathbb{P} \times \mathbb{P}$ to \mathbf{p} , a countably infinite set. We will write $f(p)$ as a shorthand for $\{v | f(v') = v \wedge v' \in p\}$. We have three important facts about this bijection:

$$\begin{aligned} \mathbf{p} &= prod_{\mathbf{p}}(\mathbb{P} \times \mathbb{P}) \\ prod_{\mathbf{p}}((\mathbb{P} \setminus (p' \uplus p'')) \times \mathbb{P}) \uplus prod_{\mathbf{p}}(p'' \times \mathbb{P}) &= prod_{\mathbf{p}}((\mathbb{P} \setminus p') \times \mathbb{P}) \\ |p'| \geq 1 \Rightarrow |prod_{\mathbf{p}}(p' \times \mathbb{P})| &= |\mathbb{P}| \end{aligned}$$

Let us consider how we could use this to do multilevel counting. We define a predicate to represent counting permissions for a resource:

$$counter(x, \mathbf{p}, n) \stackrel{\text{def}}{=} \exists p''. x \overset{\mathbf{p} \setminus prod_{\mathbf{p}}(p'' \times \mathbb{P})}{\mapsto} _ \wedge |p''| = n \wedge |\mathbf{p}| = |\mathbb{P}|$$

This definition leads to the following equivalence, which can be used to construct and destruct counters.

$$|\mathbf{p}| = |\mathbb{P}| \wedge x \overset{\mathbf{p}}{\mapsto} _ \Leftrightarrow counter(x, \mathbf{p}, 0)$$

Now we want implications to allow us to count out permissions

$$counter(x, \mathbf{p}, n) \Leftrightarrow counter(x, \mathbf{p}, n+1) * (x \overset{p'}{\mapsto} _ \wedge |p'| = |\mathbb{P}|) * counted(\mathbf{p}, p')$$

We need an additional predicate to represent the information that p' is a permission counted from \mathbf{p} . This can be defined as

$$counted(\mathbf{p}, p') \stackrel{\text{def}}{=} \exists p''. p' = prod_{\mathbf{p}}(p'' \times \mathbb{P}) \wedge |p''| = 1$$

These predicates show how this model supports multilevel counting.

5.6 Concluding remarks

The named permissions model is highly expressive and for many examples using a simpler model, such as counting, may be useful. One issue is a sensible syntax for representing the different idioms. The predicates for counting given above abstract the details of the required bijections, and similar predicates can be provided for splitting. However, it may be that a better syntax can be found for expressing the permissions. Ross and Bornat have independently suggested using strings to represent permissions. This may prove a more concise notion than the sets and bijections used above, but this remains further work.

6

Inheritance

So far we have given a logic that is capable of reasoning about programs which use encapsulation. There is a second form of abstraction used in object-oriented programming: inheritance. In this chapter we present the extensions required to allow reasoning about inheritance.

We present an extension to abstract predicates, called abstract predicate families, that allow each class to specify its own definition for that class. This allows subtypes to have a different internal representation. We demonstrate the usefulness of abstract predicate families by considering the `Cell/Recell` classes, and an extended example of the visitor design pattern [40].

The rest of the chapter is structured as follows. We begin, in §6.1 by discussing the difficulties in reasoning about inheritance. In §6.2, we present our solution: abstract predicate families. In §6.3, we illustrate their use with two examples. In §6.4, we present the formal semantics of this extension. We conclude, in §6.5, by discussing an open issue of reverifying inherited methods and present some possible solutions.

6.1 The problem

First let us revisit the problem shown in §3.8 to highlight why standard separation logic (even extended with abstract predicates) cannot deal with inheritance. Consider the `Cell` class and a subclass that has backup, `Recell`, presented in Figure 6.1. We give the obvious specifications

```
class Cell extends Object {
  Object cnts;

  void set(Object o) {
    this.cnts = o;
  }

  Object get() {
    Object temp;
    temp = this.cnts;
    return temp;
  }
}

class Recell extends Cell {
  Object bak;

  void set(Object o) {
    Object temp;
    temp = this.cnts;
    this.bak = temp;
    this.cnts = o;
  }
}
```

Figure 6.1: Source code for `Cell` and `Recell` classes

Method	Pre-condition	Post-condition
<code>Cell.set(n)</code>	$\text{this.cnts} \mapsto _$	$\text{this.cnts} \mapsto n$
<code>Recell.set(n)</code>	$\text{this.cnts} \mapsto X$ $* \text{this.bak} \mapsto _$	$\text{this.cnts} \mapsto n$ $* \text{this.bak} \mapsto X$

Table 6.1: Specifications of `set` methods.

of the `set` methods in Table 6.1. Unfortunately these specifications do not satisfy specification compatibility, or behavioural subtyping.

We require the following two implications to hold

$$\begin{aligned} \text{pre}(\text{Cell}, \text{set}) &\Rightarrow \text{pre}(\text{Recell}, \text{set}) \\ \text{post}(\text{Recell}, \text{set}) &\Rightarrow \text{post}(\text{Cell}, \text{set}) \end{aligned}$$

where $\text{pre}(C, m)$ denotes the pre-condition for the method m in class C , and post denotes the post-condition.

Given the specifications in Table 6.1, these implications can never hold as they require a one element heap to be the same size as a two element heap. What about abstract predicates? The specification for `Recell` must use a different abstract predicate to `Cell` as it has a different body, that is we define the abstract predicates as

$$\text{Val}_{\text{Cell}}(y, x) \stackrel{\text{def}}{=} y.\text{cnts} \mapsto x \tag{6.1}$$

$$\text{Val}_{\text{Recell}}(x, y, z) \stackrel{\text{def}}{=} x.\text{cnts} \mapsto y * x.\text{bak} \mapsto z \tag{6.2}$$

and the specifications as

Method	Pre-condition	Post-condition
<code>Recell.set(n)</code>	$\text{Val}_{\text{Recell}}(\text{this}, X, _)$	$\text{Val}_{\text{Recell}}(\text{this}, n, X)$
<code>Cell.set(n)</code>	$\text{Val}_{\text{Cell}}(\text{this}, _)$	$\text{Val}_{\text{Cell}}(\text{this}, n)$

Unfortunately the predicates are treated parametrically: no implications hold between them.

As it stands, abstract predicates do not, by themselves, help with behavioural subtyping. They provide support for encapsulation but not inheritance. In an object-oriented setting we require predicates to have multiple definitions, hence we introduce *abstract predicate families* where the families are sets of definitions indexed by class. An instance of an abstract predicate family is written $\alpha(x; \vec{v})$ to indicate that the object x satisfies *one* definition from the abstract predicate family α with arguments \vec{v} . The particular definition satisfied depends on the dynamic type of x . In object-oriented programming an object could be from one of many classes; abstract predicate families provide a similar choice of predicate definitions when considering their behaviour.

6.2 Abstract predicate families

We define abstract predicate family definitions, Λ^f , with the following syntax.

$$\Lambda^f := \epsilon \mid (\alpha_C(x; \vec{x}) \stackrel{\text{def}}{=} P), \Lambda^f$$

Λ^f is well-formed if it has at most one entry for each predicate and class name pair, and the free variables of the body, P , are in its argument list, $x; \vec{x}$. The first argument is distinguished as it

is used to index by class. We use a semi-colon, ;, to separate this distinguished arguments, and commas to separate the remaining arguments. We treat Λ^f as a function from predicate and class name pairs to predicate definitions. Each entry corresponds to the definition of an abstract predicate family for a particular class.

One can think of an abstract predicate family as a predicate that existentially hides the type:

$$\alpha(e; \bar{e}) \approx \exists C. \Lambda^f(\alpha, C)[e; \bar{e}] \wedge e : C$$

This is *not* the complete truth as it does not account for partial definitions, where the whole class hierarchy is unavailable. This intuition mirrors the static typing of an object: `void m(C c) { ... }` is a method that expects something that behaves like a `C`: it might not be a `C` but it will be a subtype of `C`. Abstract predicate families mirror this in the logic by allowing assertions that require a state satisfying a property but where the property's precise definition is unknown.

This semantic intuition, leads to the obvious adaptation of the OPEN and CLOSE axioms from Chapter 4.

OPEN	$\Lambda^f \models (x : C \wedge \alpha(x; \bar{x})) \Rightarrow \Lambda^f(\alpha, C)[x; \bar{x}]$
CLOSE	$\Lambda^f \models (x : C \wedge \Lambda^f(\alpha, C)[x; \bar{x}]) \Rightarrow \alpha(x; \bar{x})$
where $\alpha, C \in \text{dom}(\Lambda^f)$.	

To OPEN or CLOSE a predicate we must know which class contains the definition, and must have that definition in scope.

Our example shows the need to alter the arity of the predicate to reflect casting; the `Recell`'s `Val` predicate has three arguments while the `Cell`'s only has two. Hence we provide the following pair of implications:

WIDEN	$\Lambda^f \models \alpha(x; \bar{x}) \Rightarrow \exists \bar{y}. \alpha(x; \bar{x}, \bar{y})$
NARROW	$\Lambda^f \models \exists \bar{y}. \alpha(x; \bar{x}, \bar{y}) \Rightarrow \alpha(x; \bar{x})$

As we allow change in arity we must provide an operation to provide too few or too many arguments to the definition. We use square brackets to denote this unusual argument application operation.

$$(\alpha_C(x; \bar{x}) \stackrel{\text{def}}{=} P)[e; \bar{e}] \stackrel{\text{def}}{=} \begin{cases} P[e/x, \bar{e}_1/\bar{x}] & |\bar{e}_1| = |\bar{x}| \text{ and } \bar{e} = \bar{e}_1, \bar{e}_2 \\ \exists \bar{y}. P[e/x, (\bar{e}, \bar{y})/\bar{x}] & |\bar{e}, \bar{y}| = |\bar{x}| \end{cases}$$

If we give a predicate more variables than its definition requires, it ignores them, and if too few, it treats the missing arguments as existentially quantified. This definition of substitution is used to give the families' version of OPEN and CLOSE.

Note: We can OPEN predicates at incorrect arities as the substitution will correctly manipulate the arguments. An alternative approach would be to restrict opening to the correct arity, and use WIDEN and NARROW to get the correct arity. However, this alternative approach complicates the semantics.

Again we have two rules for introducing and eliminating abstract predicate families.

ABSTRACT WEAKENING	$\Lambda^f; \Gamma \vdash \{P\}C\{Q\}$
	$\Lambda^f, \Lambda^{f'}; \Gamma \vdash \{P\}C\{Q\}$
	provided $dom(\Lambda^{f'})$ and $dom(\Lambda^f)$ are disjoint.
ABSTRACT ELIMINATION	$\Lambda^f, \Lambda^{f'}; \Gamma \vdash \{P\}C\{Q\}$
	$\Lambda^f; \Gamma \vdash \{P\}C\{Q\}$
	provided the predicate names in P, Q, Γ and Λ^f are not in $dom_a(\Lambda^{f'})$ and $dom_a(\Lambda^f)$ is defined as $\{\alpha \mid (\alpha, C) \in dom(\Lambda^f)\}$.

Abstract predicate families are less symmetric than abstract predicates: weakening allows the introduction for a particular class and predicate, while elimination requires the entire family of definitions to be removed, i.e. it must remove all the classes' definitions for a predicate. This is because it is not possible to give a simple syntactic check for which parts, i.e. classes, of a family will be used.

6.3 Examples

Next we will present a couple of examples to illustrate the use of abstract predicate families.

6.3.1 Cell/Recell

Let us return to our original motivating example. We define an abstract predicate family, Val , with the definitions for `Cell` and `Recell` given earlier in equations (6.1) and (6.2).

We have to validate four methods: `Cell.set`, `Cell.get`, `Recell.set` and `Recell.get`. Even though the bodies of `Cell.get` and `Recell.get` are the same, we must validate both, because they have different predicate definitions. We will return to this point in §6.5.

The proofs all follow in a straightforward manner: we present the proof of the `Recell`'s `set` method here.

```

{ Val(this; X, _) ∧ this : Recell }
{ this.cnts ↦ X * this.bak ↦ _ ∧ this : Recell }
  temp = this.cnts;
{ this.cnts ↦ X * this.bak ↦ _ ∧ this : Recell ∧ X = temp }
  this.bak = temp;
{ this.cnts ↦ X * this.bak ↦ temp ∧ this : Recell ∧ X = temp }
  this.cnts = 0;
{ this.cnts ↦ 0 * this.bak ↦ temp ∧ this : Recell ∧ X = temp }
{ this.cnts ↦ 0 * this.bak ↦ X ∧ this : Recell }
{ Val(this; 0, X) }

```

Additionally, we must prove the method specifications are compatible, defined in the sense of definition 3.5.2 (Page 52). The compatibility of the `set` method follows from the rule of L-CONSEQUENCE and L-AUXVARELIM.

$$\frac{\frac{\vdash \{Val(this; X, -)\} - \{Val(this; n, X)\}}{\vdash \{Val(this; -, -)\} - \{Val(this; n, -)\}}}{\vdash \{Val(this; -)\} - \{Val(this; n)\}}$$

The two `get` methods have the same specification, so are obviously compatible.

A client that uses these methods does not need to worry about dynamic dispatch. Consider the following method:

Method	Pre-condition	Post-condition
<code>NonNullCell.set(x)</code>	$Val(\text{this}; _) \wedge x \neq \text{null}$	$Val(\text{this}; x)$
<code>NonNullCell.get()</code>	$Val(\text{this}; X)$	$Val(\text{this}; X) \wedge X \neq \text{null}$ $\wedge X = \text{ret}$

$$Val_{\text{NonNullCell}}(x; y) \stackrel{\text{def}}{=} x.\text{cnts} \mapsto y \wedge y \neq \text{null}$$

Table 6.2: Specification for `NonNullCell`.

```
m(Cell c) {
  c.set(c);
}
```

This code simply sets the `Cell` to point to itself. The code is specified as

Method	Pre-condition	Post-condition
<code>m(c)</code>	$Val(c; _)$	$Val(c; c)$

Now consider calling `m` with a `Recell` object.

```
{ true }
Recell r = new Recell(x);
{ Val(r; x, -) * r : Recell }
{ Val(r; x, -) }
{ Val(r; -) }
m(r);
{ Val(r; r) }
{ Val(r; r, -) }
{ Val(r; r, -) * r : Recell }
```

We use L-CONSEQUENCE to cast the `Val` predicate to have the correct arity. We need not consider dynamic dispatch at all because of specification compatibility.

The specification of method `m` is weaker than we might like. Based on the implementation we might expect the example’s post-condition $\{Val(r; r, x)\}$. However, there are several bodies that satisfy `m`’s specification: for example `c.set(x); c.set(c);`. We can set the `Cell` to have any value, as long as the last value we set is the `Cell` itself. This body acts identically on a `Cell` to the previous body, however on a `Recell` it acts differently. Hence only using the specification we cannot infer the tighter post-condition. This could be deduced if `m` was specified for a `Recell` as well.

Now let us briefly consider another `Cell` class: a `NonNullCell`. It specifies the arguments to `set` must not be `null`. We present its specification in Figure 6.2. A `NonNullCell` is not a subtype of `Cell` as we cannot replace `Cell` with a `NonNullCell`, i.e. `x.set(null);` can only be called on a `Cell`. However, it is possible to make `Cell` a subtype of `NonNullCell`. We must alter the specification of `Cell` so that it has two predicate families: one for the non-null state, and one for the null state. We present the specifications and predicate definitions in Figure 6.2. Now we can show the specification compatibility of these two methods as follows. First we

Method	Pre-condition	Post-condition
set (y)	$Val(\text{this}; _)$ $\vee NVal(\text{this})$	$(y = \text{null} \wedge NVal(\text{this})) \vee$ $(y \neq \text{null} \wedge Val(\text{this}; y))$
get ()	$Val(\text{this}; x) \wedge X = 1$ $\vee NVal(\text{this}) \wedge X = 2$	$(X = 1 \wedge Val(\text{this}; x) \wedge ret = x \wedge x \neq \text{null})$ $\vee (X = 2 \wedge NVal(\text{this}) \wedge ret = \text{null})$

$$Val_{\text{Cell}}(x; y) \stackrel{\text{def}}{=} x.\text{cnts} \mapsto y \wedge y \neq \text{null}$$

$$NVal_{\text{Cell}}(x) \stackrel{\text{def}}{=} x.\text{cnts} \mapsto \text{null}$$

Figure 6.2: New specification for Cell class.

present the set method.

$$\frac{\left\{ \begin{array}{l} Val(\text{this}; x) \\ \vee NVal(\text{this}) \end{array} \right\} - \left\{ \begin{array}{l} (y = \text{null} \wedge NVal(\text{this})) \\ \vee (Val(\text{this}; y) \wedge y \neq \text{null}) \end{array} \right\}}{\left\{ y \neq \text{null} * \left(\begin{array}{l} Val(\text{this}; x) \\ \vee NVal(\text{this}) \end{array} \right) \right\} - \left\{ \left(\begin{array}{l} (y = \text{null} \wedge NVal(\text{this})) \\ \vee (Val(\text{this}; y) \wedge y \neq \text{null}) \end{array} \right) * y \neq \text{null} \right\}} \\ \{y \neq \text{null} \wedge Val(\text{this}; x)\} - \{Val(\text{this}; y)\}$$

Although it is possible for the Cell's set method to end in a state that is not expected by the NonNullCell's post-condition, this proof shows it is not possible to end in an unexpected state if the NonNullCell's pre-condition is satisfied. We use the frame rule to preserve the pre-condition that $y \neq \text{null}$ across the call.

Note: We have used the ghost variable to encode the specification intersection as described in §3.5.

Next we present a similar proof for the get () method.

$$\frac{\left\{ \begin{array}{l} Val(\text{this}; x) \wedge X = 1 \\ \vee NVal(\text{this}) \wedge X = 2 \end{array} \right\} - \left\{ \begin{array}{l} (X = 1 \wedge Val(\text{this}; x) \wedge ret = x \wedge x \neq \text{null}) \\ \vee (X = 2 \wedge NVal(\text{this}) \wedge ret = \text{null}) \end{array} \right\}}{\left\{ \left(\begin{array}{l} Val(\text{this}; x) \wedge X = 1 \\ \vee NVal(\text{this}) \wedge X = 2 \end{array} \right) * X = 1 \right\} - \left\{ \left(\begin{array}{l} (X = 1 \wedge Val(\text{this}; x) \wedge ret = x \wedge x \neq \text{null}) \\ \vee (X = 2 \wedge NVal(\text{this}) \wedge ret = \text{null}) \end{array} \right) * X = 1 \right\}} \\ \frac{\{Val(\text{this}; x) \wedge X = 1\} - \{Val(\text{this}; x) \wedge x = ret \wedge x \neq \text{null}\}}{\frac{\{\exists X. Val(\text{this}; x) \wedge X = 1\} - \{Val(\text{this}; x) \wedge x = ret \wedge x \neq \text{null}\}}{\{Val(\text{this}; x)\} - \{Val(\text{this}; x) \wedge x = ret \wedge x \neq \text{null}\}}}$$

This proof is slightly more complicated as we must eliminate the ghost variable X and preserve its value across the call using the frame rule. However, once again we see that although the method has a more general post-condition, it is a valid replacement for the non-null version.

6.3.2 Visitor pattern

Now let us consider an extended example using the visitor design pattern [40]. The visitor pattern allows code in object-oriented programs to be function-oriented rather than data-oriented. That is, all the methods for a particular operation are grouped in a single class, rather than


```

class Visitor {
    void visitC(Const x) {;}
    void visitP(Plus x) {;}
}

class Ast extends Object {
    void accept(Visitor x) {;}
}

class Const extends Ast {
    int v;
}

void accept(Visitor x) {
    x.visitC(this);
}

class Plus extends Ast {
    Ast l; Ast r;
    void accept(Visitor x) {
        x.visitP(this);
    }
}

```

Figure 6.3: Source code for formulae and visitor classes.

grouping all the operations on a particular type of data in a single class. In functional languages, such as ML, a particular operation (function) is defined using pattern matching. The visitor pattern allows object-oriented programmers to express this common idiom by allowing matching on an object's type without using the `instanceof` keyword found in Java. We consider an imperative, external visitor [18] as we do not wish to focus on any of the complexities in typing visitors [17].

We consider a visitor over a very simple syntax of formulae with just addition of two terms and constants. We present the source code in Figure 6.3. In this example we will use a primitive type of integer, `int`, but for compactness we will not formally extend the logic or language's semantics. The `Visitor` class has two methods: the first, `visitC` is invoked when the visitor visits a `Const` node; and the second, `visitP` is for a `Plus` node. The `Visitor` class is a template that should be overridden to produce more interesting visitors.¹ We define three classes to represent an abstract syntax tree of these formulae: `Ast` which is used as a common parent for the term constructors; `Const` that represents constant terms; and `Plus` that represents the addition of two terms. We define a single method, `accept`, in `Ast` that is overridden in each of the subclasses. The `Const` class represents an integer constant, and calls `visitC` when it accepts a visitor. The `Plus` class represents an addition of two formulae and calls `visitP` when it accepts a visitor. The double invocation of calling `accept` followed by `visit` is called double dispatch, and allows the case switch to occur without any downcasts or `instanceof` checks. We will return to this point in the proof later.

Before we can formally specify the visitor, we must extend the logic to have values of the formulae's syntax:

$$\begin{aligned}
 \tau & ::= n \mid \underline{\tau} \pm \tau \\
 \mu & ::= \bullet \mid \tau \underline{\pm} \mu \mid \mu \underline{\pm} \tau
 \end{aligned}$$

We underline the plus constructor, $\underline{\pm}$, to distinguish it from the arithmetic operation, and use $\mu[\mu']$ to mean replace \bullet by μ' in μ . We give the specifications for the `visit` and `accept` methods in Table 6.3. There are many choices one can make for the specification of a visitor. The μ parameter is used to allow the evaluation to depend on the context. In the example that follows we use the context to allow us to accumulate the value of the expression as we traverse it, rather than having to calculate it in a bottom up fashion. One can remove the μ parameter from the predicates if we do not want context sensitive visitors. That is, the state of the visitor does not depend on its context. One might have expected the methods all to have post-conditions of

¹One would normally make the `Visitor` class abstract, but this is not in our syntax. Instead, we can make its constructor's pre-condition *false* (not presented). Thus it can never be instantiated.

Method	Pre-condition	Post-condition
<code>visitC(x)</code>	$x : \text{Const} \wedge \text{Ast}(x, \tau) * \text{Visitor}(\text{this}, \mu)$	$\text{Visited}(\text{this}; x, \tau, \mu)$
<code>visitP(x)</code>	$x : \text{Plus} \wedge \text{Ast}(x, \tau) * \text{Visitor}(\text{this}, \mu)$	$\text{Visited}(\text{this}; x, \tau, \mu)$
<code>accept(x)</code>	$\text{Ast}(\text{this}, \tau) * \text{Visitor}(x, \mu)$	$\text{Visited}(x; \text{this}, \tau, \mu)$

Table 6.3: Specifications for accept and visit methods.

```

class Calc
  extends Visitor {
  int amount;
  void visitC(Const x) {
    this.amount += x.n;
  }
  void visitP(Plus x) {
    x.l.accept(this);
    x.r.accept(this);
  }
}

```

Figure 6.4: Source code for Calc visitor

the form $\text{Visited}(\dots) * \text{Ast}(x, \tau)$. However, this kind of specification prevents us altering the structure of the expression.

We define the Ast predicate family for the three classes.

$$\begin{aligned}
 \text{Ast}_{\text{Ast}}(x, \tau) &\stackrel{\text{def}}{=} \text{false} \\
 \text{Ast}_{\text{Const}}(x, \tau) &\stackrel{\text{def}}{=} x.v \mapsto n \wedge \tau = n \\
 \text{Ast}_{\text{Plus}}(x, \tau) &\stackrel{\text{def}}{=} x.l \mapsto i * x.r \mapsto j * \text{Ast}(i, \tau_l) * \text{Ast}(j, \tau_r) \wedge \tau = \tau_l \pm \tau_r
 \end{aligned}$$

Note: Defining the predicate for the `Ast` class as *false* prevents any invocation of its methods.

We are now in a position to verify the `accept` methods for the `Plus` and `Const` classes. We only present the `Plus` case as the `Const` case is very similar.

```

{Ast(this : τ) * Visitor(x; μ) ∧ this : Plus}
  x.visitP(this);
{Visitor(x; this, τ, μ)}

```

The verification of the method call introduces the current class's type: this is exactly what is required to meet `visitP`'s specification. This method call simply provides information about which type of node this is to the `Visitor`: it is providing a case select using dynamic dispatch. This proof does not need changing no matter how many subclasses of `Visitor` are added.

Let us consider an actual visitor implementation. The implementation given in Figure 6.4 calculates the value of the formulae. We can define the Visitor predicate for this class as

$$\begin{aligned}
 \text{Visitor}_{\text{Calc}}(x; \mu) &\stackrel{\text{def}}{=} x.\text{amount} \mapsto \text{lcalc}(\mu) \\
 \text{Visited}_{\text{Calc}}(x; y, \tau, \mu) &\stackrel{\text{def}}{=} x.\text{amount} \mapsto (\text{lcalc}(\mu) + \text{calc}(\tau)) * \text{Ast}(y; \tau)
 \end{aligned}$$

where we define the function *calc* and *lcalc* as:

$$\begin{aligned} \text{calc}(\tau) &\stackrel{\text{def}}{=} \begin{cases} n & \tau = n \\ \text{calc}(\tau_1) + \text{calc}(\tau_2) & \tau = \tau_1 \pm \tau_2 \end{cases} \\ \text{lcalc}(\mu) &\stackrel{\text{def}}{=} \begin{cases} 0 & \mu = \bullet \\ \text{lcalc}(\mu_1) & \mu = \mu_1 \pm \tau \\ \text{calc}(\tau) + \text{lcalc}(\mu_1) & \mu = \tau \pm \mu_1 \end{cases} \end{aligned}$$

The *lcalc* function is used to calculate the accumulated total from the context: the sum of everything to the left of the hole, \bullet , i.e. the nodes we have already visited.

We verify the `visitP` method as²

```
{x : Plus  $\wedge$  Ast(x;  $\tau$ ) * Visitor(this;  $\mu$ )  $\wedge$  this : Calc}
{P * x.l  $\mapsto$  i * x.r  $\mapsto$  j * Ast(i,  $\tau_1$ ) * Ast(j,  $\tau_2$ ) * Visitor(this;  $\mu$ )}
  {x.l  $\mapsto$  i * Ast(i,  $\tau_1$ ) * Visitor(this;  $\mu$ ,  $\bullet$ )}
    x.l.accept(this);
    {x.l  $\mapsto$  i * Visited(this; i,  $\tau_1$ ,  $\mu$ )}
  {P * x.l  $\mapsto$  i * x.r  $\mapsto$  j * Ast(j,  $\tau_2$ ) * Visited(this; i,  $\tau_1$ ,  $\mu$ )}
  {P * x.l  $\mapsto$  i * x.r  $\mapsto$  j * Ast(i,  $\tau_1$ ) * Ast(j,  $\tau_2$ ) * Visitor(this;  $\mu[\tau_1 \pm \bullet]$ )}
    {x.r  $\mapsto$  j * Ast(j,  $\tau_2$ ) * Visitor(this;  $\mu[\tau_1 \pm \bullet]$ )}
      x.r.accept(this);
      {x.r  $\mapsto$  j * Visited(this; j,  $\tau_2$ ,  $\mu[\tau_1 \pm \bullet]$ )}
  {P * x.l  $\mapsto$  i * x.r  $\mapsto$  j * Ast(i,  $\tau_1$ ) * Visited(this; j,  $\tau$ ,  $\mu$ )}
  {Visited(this; x,  $\tau$ ,  $\mu$ )}
```

where $P = x : \text{Plus} \wedge \text{this} : \text{Calc} \wedge \tau = \tau_1 \pm \tau_2$

This proof follows from the following two formulae.

$$\begin{aligned} x : \text{Calc} &\Rightarrow \text{Visited}(x; y, \tau, \mu) \Leftrightarrow \text{Visitor}(x; \mu[\tau \pm \bullet]) * \text{Ast}(y, \tau) \\ &\left(\begin{array}{l} \text{Ast}(i, \tau_1) * x.l \mapsto i * x.r \mapsto j \wedge y : \text{Calc} \\ * \text{Visited}(y; j, \tau_2, \mu[\tau_1 \pm \bullet]) \wedge x : \text{Plus} \end{array} \right) \Leftrightarrow \text{Visited}(\text{this}; x, \tau_1 \pm \tau_2, \mu) \end{aligned}$$

Both of which follow directly from $\text{lcalc}(\mu) + \text{calc}(\tau) = \text{lcalc}(\mu[\tau \pm \bullet])$ which can be shown by induction on μ .

We can produce different subclasses of `Visitor` to perform many different functions, such as checking whether an AST contains a particular constant, or to clone an AST. As the proof is modular we only need to verify the new classes we write: we know all the other classes interact correctly.

Next we consider a subclass of `Visitor` that mutates the structure of an expression by removing any addition of zero. First, we present a function that operates on the abstract syntax.

$$\begin{aligned} rz(\tau_1 \pm \tau_2) &\stackrel{\text{def}}{=} \begin{cases} rz(\tau_1) & rz(\tau_2) = 0 \\ rz(\tau_2) & rz(\tau_1) = 0 \\ rz(\tau_1) \pm rz(\tau_2) & \text{otherwise} \end{cases} \\ rz(n) &\stackrel{\text{def}}{=} n \end{aligned}$$

This function would simplify $0 \pm (3 \pm 4)$ to 3 ± 4 , and 0 ± 0 to 0. It cannot remove all occurrences of zero, but it removes all additions of zero. We present a program that does this in Figure 6.5. We have used MJ syntax extended with booleans for compactness and legibility. The `Visitor` does an in-place removal: it does not allocate any storage; instead it just modifies the original

```

class RemoveZeros
    extends Visitor {
    boolean isZero;
    boolean isChanged;
    Ast newl;

    void visitC(Const c) {
        if(c.n==0) {
            this.isZero = true;
        }
    }

    void visitP(Plus p) {
        p.l.accept(this);
        if(this.isZero) {
            this.isChanged = false;
            this.isZero = false;
            p.r.accept(this);
            if(!this.isChanged) {
                this.newl = p.r;
            }
        } else {
            if(this.isChanged) {
                p.l = this.newl;
                this.isChanged = false;
            }
            p.r.accept(this);
            if(this.isZero) {
                this.isChanged = true;
                this.newl = p.l;
                this.isZero = false;
            }
            else if(isChanged) {
                p.r = this.newl;
                this.isChanged = false;
            }
        }
        this.isChanged = true;
    }
}

```

Figure 6.5: Source code of RemoveZeros visitor

structure. Rather than explain the rather difficult corners cases we will present the proof. This highlights why each item is required. We use RZ as a shorthand for `RemoveZeros`. We specify the *Visitor* and *Visited* predicates as follows

$$\begin{aligned}
 Visitor_{RZ}(x; \mu) &\stackrel{\text{def}}{=} x.newl \mapsto _ * x.isZero \mapsto false * x.isChanged \mapsto false \\
 Visited_{RZ}(x; y, \tau, \mu) &\stackrel{\text{def}}{=} x.newl \mapsto a * x.isZero \mapsto b * x.isChanged \mapsto c \\
 &\quad * ((c \wedge Ast(a; rz(\tau))) \vee (\neg c \wedge Ast(y; rz(\tau)))) \wedge (b \Leftrightarrow rz(\tau) = 0)
 \end{aligned}$$

The `isZero` variable specifies that the expression just visited is equivalent to zero: it might have been a single constant 0 or any composite expression with only zero constants, e.g. 0 ± 0 . The `isChanged` variable indicates if the removal of zeros altered the top node, and hence we must now link to a sub-term and the `newl` variable denotes this sub-term. We present an overview of the proof in Figure 6.6. We do not present all the tedious details, but instead consider the highlighted part of the proof in Figure 6.7. The most interesting step is at the end where we use weakening to remove the `p.l` and `p.r` fields.

This proof demonstrates the support that abstract predicate families provides for inheritance. The imperative nature of the visitor is well suited to separation logic. The reader should note that this algorithm is non-trivial to implement: indeed, the author's initial implementation had several small mistakes. Although testing would have caught some of these errors, formal verification guarantees none still exist.

6.4 Semantics

In this section we consider the extensions to the semantics of §4.3 sufficient to model abstract predicate families. The semantic predicate environment as defined in §4.3 has to be extended to handle the arity changes that predicate families require. We define a semantic predicate family

²Rather than present the code broken into simple Inner MJ statements, we will present a proof skeleton using MJ statements.

```

{p : Plus  $\wedge$  Ast(p;  $\tau$ ) * Visitor(this;  $\mu$ )  $\wedge$  this : RZ}
{p : Plus  $\wedge$  p.l  $\mapsto$  i * Ast(i;  $\tau_1$ ) * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visitor(this;  $\mu$ )  $\wedge$  this : RZ}
  p.l.accept(this);
{p : Plus  $\wedge$  p.l  $\mapsto$  i * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visited(this; i,  $\tau_1$ ,  $\mu$ )  $\wedge$  this : RZ}
  if(this.isZero) {
    {p : Plus  $\wedge$  p.l  $\mapsto$  i * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visited(this; i,  $\tau_1$ ,  $\mu$ )  $\wedge$  this.isZero  $\mapsto$  true  $\wedge$  this : RZ}
    this.isZero = false;
    this.isChanged = false;
    {p : Plus  $\wedge$  p.l  $\mapsto$  i * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visitor(this;  $\mu$ )  $\wedge$  this : RZ  $\wedge$  rz( $\tau_1$ ) = 0}
    p.r.accept(this);
    {p : Plus  $\wedge$  p.l  $\mapsto$  i * p.r  $\mapsto$  j *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visited(this; j,  $\tau_2$ ,  $\mu$ )  $\wedge$  this : RZ  $\wedge$  rz( $\tau_1$ ) = 0}
    if(!this.isChanged) {
      this.newl = p.r;
      this.isChanged = true;
    }
    {Visited(this; p,  $\tau$ ,  $\mu$ )}
  } else {
    {p : Plus  $\wedge$  p.l  $\mapsto$  i * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visited(this; i,  $\tau_1$ ,  $\mu$ )  $\wedge$  this.isZero  $\mapsto$  false  $\wedge$  this : RZ}
    if(this.isChanged) {
      p.l = this.newl;
      this.isChanged = false;
    }
    { $\exists k$ . p : Plus  $\wedge$  p.l  $\mapsto$  k * p.r  $\mapsto$  j * Ast(j;  $\tau_2$ ) * Ast(k, rz( $\tau_1$ )) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visitor(this;  $\mu$ )  $\wedge$  this : RZ  $\wedge$  rz( $\tau_1$ )  $\neq$  0}
    p.r.accept(this);
    { $\exists k$ . p : Plus  $\wedge$  p.l  $\mapsto$  k * p.r  $\mapsto$  j * Ast(k, rz( $\tau_1$ )) *  $\tau_1 \pm \tau_2 = \tau$ 
 * Visited(this; j,  $\tau_2$ ,  $\mu$ )  $\wedge$  this : RZ  $\wedge$  rz( $\tau_1$ )  $\neq$  0}
    if(this.isZero) {
      this.isChanged = true;
      this.newl = p.l;
      this.isZero = false;
    }
    else if(isChanged) {
      p.r = this.newl;
      this.isChanged = false;
    }
    {Visited(this; p,  $\tau$ ,  $\mu$ )}
  }
{Visited(this; p,  $\tau$ ,  $\mu$ )}

```

Figure 6.7

Figure 6.6: RemoveZeros' proof overview.

$$\left\{ \begin{array}{l} p : \text{Plus} \wedge p.l \mapsto i * p.r \mapsto j * \tau_1 \pm \tau_2 = \tau \wedge \text{this} : \text{RZ} \wedge rz(\tau_1) = 0 \\ \quad * \text{Visited}(\text{this}; j, \tau_2, \mu) \end{array} \right\}$$

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} \\ * \text{Visited}(\text{this}; j, \tau_2, \mu) \wedge \text{this.isChanged} \mapsto _ \end{array} \right\}$$

b = this.isChanged;

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} \\ * \text{Visited}(\text{this}; j, \tau_2, \mu) \wedge \text{this.isChanged} \mapsto b \end{array} \right\}$$

if (!b) {

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} \\ * \text{Visited}(\text{this}; j, \tau_2, \mu) \wedge \text{this.isChanged} \mapsto \text{false} \end{array} \right\}$$

t = p.r;

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto t * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} \\ * \text{Visited}(\text{this}; t, \tau_2, \mu) \wedge \text{this.isChanged} \mapsto \text{false} \end{array} \right\}$$

By expanding the definition of *Visited*

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto t * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} * \text{Ast}(t; rz(\tau_2)) \\ \quad * \text{this.newl} \mapsto _ * \text{this.isChanged} \mapsto \text{false} \\ \quad * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau_2) = 0) \end{array} \right\}$$

By replacing $rz(\tau_2)$ with $rz(\tau)$

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto t \wedge \text{this} : \text{RZ} * \text{Ast}(t; rz(\tau)) * \text{this.newl} \mapsto _ \\ * \text{this.isChanged} \mapsto \text{false} * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau) = 0) \end{array} \right\}$$

this.newl = t;

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto t \wedge \text{this} : \text{RZ} * \text{Ast}(t; rz(\tau)) * \text{this.newl} \mapsto t \\ * \text{this.isChanged} \mapsto \text{false} * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau) = 0) \end{array} \right\}$$

this.isChanged = true;

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto t \wedge \text{this} : \text{RZ} * \text{Ast}(t; rz(\tau)) * \text{this.newl} \mapsto t \\ * \text{this.isChanged} \mapsto \text{true} * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau) = 0) \end{array} \right\}$$

By the definition of *Visited*

$$\left\{ p.l \mapsto i * p.r \mapsto t * \text{Visited}(\text{this}; p, \tau, \mu) \right\}$$

} else {

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} \\ * \text{Visited}(\text{this}; j, \tau_2, \mu) \wedge \text{this.isChanged} \mapsto \text{true} \end{array} \right\}$$

By expanding the definition of *Visited*

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j * rz(\tau) = rz(\tau_2) \wedge \text{this} : \text{RZ} * \text{this.newl} \mapsto k \\ \quad * \text{Ast}(k; rz(\tau_2)) * \text{this.isChanged} \mapsto \text{true} \\ \quad * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau_2) = 0) \end{array} \right\}$$

By replacing $rz(\tau)$ with $rz(\tau_2)$

$$\left\{ \begin{array}{l} p.l \mapsto i * p.r \mapsto j \wedge \text{this} : \text{RZ} * \text{this.newl} \mapsto k * \text{Ast}(k; rz(\tau)) \\ * \text{this.isChanged} \mapsto \text{true} * \text{this.isZero} \mapsto c * (c \Leftrightarrow rz(\tau) = 0) \end{array} \right\}$$

By the definition of *Visited*

$$\left\{ p.l \mapsto i * p.r \mapsto j * \text{Visited}(\text{this}; p, \tau, \mu) \right\}$$

;

}

$$\left\{ \text{Visited}(\text{this}; p, \tau, \mu) * p.l \mapsto _ * p.r \mapsto _ \right\}$$

By weakening

$$\left\{ \text{Visited}(\text{this}; p, \tau, \mu) \right\}$$

Figure 6.7: Part of RemoveZeros' proof

environment as

$$\Delta_f : \mathcal{A} \times \mathbf{Classes} \rightarrow (\mathbb{N}^+ \rightarrow \mathcal{P}(\mathcal{H}))$$

This is a partial function from pairs of predicate, \mathcal{A} , and class name, $\mathbf{Classes}$, to semantic definitions. An abstract predicate family is defined for all arities, so the semantic definition must be a function from *all* tuples of non-zero arity. This semantically supports the change in arity required by WIDEN and NARROW.

We can now give the semantics of the new abstract predicate family assertion as follows

$$VS, H, I \models_{\Delta_f} \alpha(e; \bar{e}) \Leftrightarrow H \in (\Delta_f(\alpha, C))[\llbracket e; \bar{e} \rrbracket_{VS, I}] \wedge H(\llbracket e \rrbracket_{VS, I}) = C$$

The assertion $\alpha(e; \bar{e})$ holds for some heap H , iff $\llbracket e \rrbracket_{VS, I}$ has class C , and H satisfies the predicate definition for C , given arguments $\llbracket e, \bar{e} \rrbracket_{VS, I}$, in the predicate family α .

To ensure that WIDEN and NARROW hold we restrict our attention to *argument refineable* environments.

Definition 6.4.1 (Argument refineable). A semantic predicate family environment is said to be argument refineable if adding an argument cannot increase, or decrease, the set of accepting states, i.e.

$$AR(\Delta_f) \Leftrightarrow \forall \alpha, n, \bar{n}. \Delta_f(\alpha)[n; \bar{n}] = \bigcup_{n'} \Delta_f(\alpha)[n; \bar{n}, n']$$

Next we prove this semantics condition allows the arity change provided by WIDEN and NARROW.

Proposition 6.4.2. *Argument refinement coincides precisely with WIDEN and NARROW:* $\forall VS, H, I, \alpha, n, \bar{n}$.

$$AR(\Delta_f) \iff (VS, H, I \models_{\Delta_f} \alpha(n; \bar{n}) \Leftrightarrow \exists n'. \alpha(n; \bar{n}, n'))$$

Proof. Straight from the definitions of argument refineable, and the semantics of abstract predicate families and existential quantifiers. \square

Again the semantic predicate family environment is a complete lattice. The lattice has the following order:

$$\Delta_f \sqsubseteq \Delta'_f \stackrel{\text{def}}{=} \forall \alpha, C, n, \bar{n}. (\alpha, C) \in \text{dom}(\Delta_f) \Rightarrow \Delta_f(\alpha, C)[n; \bar{n}] \subseteq \Delta'_f(\alpha, C)[n; \bar{n}]$$

Again, the least upper bound of the order is written \sqcup . Lemmas 4.3.1 and 4.3.4 can be extended to semantic predicate family environments as follows:

Lemma 6.4.3. *Argument refineable predicate family environments form a complete lattice with respect to \sqsubseteq .*

Proof. We must show that the least upper bound of a set of argument refineable environments is also argument refineable.

$$(\forall i. AR(\Delta_f^i)) \Rightarrow AR(\bigsqcup_i \Delta_f^i)$$

That is shown by

$$\begin{aligned} H \in (\bigsqcup_i \Delta_f^i)(\alpha)[n; \bar{n}] &\Leftrightarrow \exists i. H \in \Delta_f^i(\alpha)[n; \bar{n}] \\ &\Leftrightarrow \exists i, n'. H \in \Delta_f^i(\alpha)[n; \bar{n}, n'] \\ &\Leftrightarrow \exists n'. H \in (\bigsqcup_i \Delta_f^i)(\alpha)[n; \bar{n}, n'] \\ &\Leftrightarrow H \in \bigcup_{n'} (\bigsqcup_i \Delta_f^i)(\alpha)[n; \bar{n}, n'] \end{aligned}$$

\square

Lemma 6.4.4. *Positive formulae are monotonic with respect to semantic predicate family environments*

$$\Delta_f \sqsubseteq \Delta'_f \wedge VS, H, I \models_{\Delta_f} P \Rightarrow VS, H, I \models_{\Delta'_f} P$$

Proof. By induction on P . □

However extending Lemma 4.3.2 is less straightforward, as it is not possible to tell which predicate name, class name pairs are used in a formula.

Lemma 6.4.5. *Formulae only depend on the abstractions they mention. If Δ_f contains all the abstractions in P , and $\text{dom}_a(\Delta_f) \cap \text{dom}_a(\Delta'_f) = \emptyset$, then*

$$\forall VS, H, I. VS, H, I \models_{\Delta_f} P \Leftrightarrow VS, H, I \models_{\Delta_f \sqcup \Delta'_f} P$$

Proof. By induction on P . □

Now let us consider the construction of semantic predicate family environments from their abstract syntactic counterparts. We define a new function, stepf , that accounts for the first argument's type and uses the special substitution,

$$\text{stepf}_{(\Lambda^f, \Delta_f)}(\Delta'_f)(\alpha, C)[n; \bar{n}] \stackrel{\text{def}}{=} \{H \mid H(n) = C \wedge VS, H, I \models_{\Delta_f \sqcup \Delta'_f} \Lambda^f(\alpha, C)[n; \bar{n}]\}$$

This function is monotonic on predicate family environments, because of Lemma 6.4.4 and that all the predicate definitions are positive. Hence by Lemma 6.4.3 and Tarski's theorem we know a fixed point must always exist. We write $\llbracket \Lambda^f \rrbracket_{\Delta_f}$ as the least fixed point of $\text{stepf}_{(\Lambda^f, \Delta_f)}$.

Lemma 6.4.6 (The fixed point is argument refineable). *If $AR(\Delta_f)$ then $AR(\llbracket \Lambda^f \rrbracket_{\Delta_f})$*

Proof. We show any fixed point is argument refineable. We know the following two equations hold:

$$\begin{aligned} VS, H, I \models_{\Delta_f} \Lambda^f(\alpha, C)[n; \bar{n}] \wedge n : C &\iff H \in \Delta_f(\alpha, C)[n; \bar{n}] \\ VS, H, I \models_{\Delta_f} \Lambda^f(\alpha, C)[n; \bar{n}] &\iff \exists n'. \Lambda^f(\alpha, C)[n; \bar{n}, n'] \end{aligned}$$

The first follows from the definition of the fixed point, and the second follows from the definition of substitution on predicate family definitions. Combining these we get: for some n'

$$\begin{aligned} VS, H, I \models_{\Delta_f} \Lambda^f(\alpha, C)[n; \bar{n}] \wedge n : C &\iff H \in \Delta_f(\alpha, C)[n; \bar{n}] \\ &\iff \\ VS, H, I \models_{\Delta_f} \Lambda^f(\alpha, C)[n; \bar{n}, n'] \wedge n : C &\iff H \in \Delta_f(\alpha, C)[n; \bar{n}, n'] \end{aligned}$$

Hence any fixed point is argument refineable. □

Consider the following set of solutions:

$$\text{close}(\Lambda^f) \stackrel{\text{def}}{=} \{\llbracket \Lambda^f \rrbracket_{\Delta_f} \sqcup \Delta_f \mid ((\mathcal{A} \times \mathbf{Classes}) \setminus \text{dom}(\Delta_f)) = \text{dom}(\Lambda^f) \wedge AR(\Delta_f)\}$$

This satisfies the analogues of Lemmas 4.3.6 and 4.3.7.

Lemma 6.4.7. *Adding new predicate definitions refines the set of possible semantic predicate environments.*

$$\text{close}(\Lambda^f) \supseteq \text{close}(\Lambda^f, \Lambda^{f'})$$

Proof. Identical to proof of Lemma 4.3.6. \square

Lemma 6.4.8. *The removal of predicate definitions does not affect predicates that do not use them, i.e. given Λ^f which is disjoint from $\Lambda^{f'}$ and does not mention predicates in its domain; we have*

$$\forall \Delta \in \text{close}(\Lambda^f). \exists \Delta'_f \in \text{close}(\Lambda^f, \Lambda^{f'}). \Delta_f \upharpoonright \text{dom}(\Lambda^{f'}) = \Delta'_f \upharpoonright \text{dom}(\Lambda^{f'})$$

where $f \upharpoonright S$ is $\{a \mapsto b \mid a \mapsto b \in f \wedge a \notin \text{dom}(S)\}$

Proof. Identical to proof of Lemma 4.3.7 with Lemma 6.4.5 replacing Lemma 4.3.2. \square

Validity is defined identically to the previous section, i.e.

$$\Lambda^f \models P \stackrel{\text{def}}{=} \forall VS, H, I, \Delta_f \in \text{close}(\Lambda^f). VS, H, I \models_{\Delta_f} P$$

Theorem 6.4.9. OPEN and CLOSE are valid, i.e.

$$\begin{aligned} \Lambda^f \models \alpha(e; \bar{e}) \wedge e : C &\Rightarrow \Lambda^f(\alpha, C)[e, \bar{e}/x, \bar{x}] \\ \Lambda^f \models \Lambda^f(\alpha, C)[e, \bar{e}/x, \bar{x}] \wedge e : C &\Rightarrow \alpha(e; \bar{e}) \end{aligned}$$

where $(\alpha, C) \in \text{dom}(\Lambda^f)$.

Proof. Direct from definition of validity and the definition of the fixed point. \square

Theorem 6.4.10. WIDEN and NARROW are valid, i.e.

$$\begin{aligned} \Lambda^f \models \alpha(e; \bar{e}) &\Rightarrow \exists X. \alpha(e; \bar{e}, X) \\ \Lambda^f \models \alpha(e; \bar{e}, e') &\Rightarrow \alpha(e; \bar{e}), \end{aligned}$$

Proof. Direct from the definition of validity, Lemma 6.4.2 and Lemma 6.4.6. \square

We are now in a position to define the semantics for our reasoning system. We write $\Lambda^f; \Gamma \models \{P\}C\{Q\}$ to mean that if every specification in Γ is true of a method environment, and every abstract predicate family in Λ^f is true of a predicate family environment, then so is $\{P\}C\{Q\}$, i.e.

$$\Lambda^f; \Gamma \models \{P\}C\{Q\} \stackrel{\text{def}}{=} \forall \Delta_f \in \text{close}(\Lambda^f). (\Delta_f \models \Gamma) \Rightarrow \Delta_f \models \{P\}C\{Q\}$$

Given this definition we can show that the two new rules for abstract predicate families are sound.

Theorem 6.4.11. Abstract weakening is sound.

Proof. Direct consequence of the definition of judgements and Lemma 6.4.7. \square

Theorem 6.4.12. Abstract elimination is sound.

Proof. Follows from Lemmas 6.4.5 and 6.4.8 \square

6.5 Reverification for inherited methods

We conclude this chapter by briefly discussing the issues of reverifying inherited methods. We will not give details of these ideas as proper consideration would require `super` calls to methods (calling a parent’s version of a method), which neither our operational semantics or programming logic supports.

In the development given so far we must verify every method of every class, even if the method is inherited. This is due to the method bodies being verified with respect to the current class’s type:

$$\text{L-DMETHOD} \quad \frac{\Lambda^f; \Gamma \vdash \{P \wedge \text{this} : C\} \bar{s}\{Q\}}{\Lambda^f; \Gamma \Vdash \{P\} C.m(\bar{x})\{Q\}} \quad \text{where } \text{mbody}(C, m) = (\bar{x}, \bar{s})$$

Prior to this chapter we do not actually require `this : C` in method verification, so we could alter the proof rules and hence allow methods to be inherited without reverification. However, the examples in this chapter depend completely on it. This assertion is required to `OPEN` and `CLOSE` the predicate families. Without it, the proofs simply fail.

We need to restrict changes to abstract predicate family definitions to allow the reuse of proofs in subtypes. We begin by presenting a generalized notion with respect to an abstract relation, and then provide two concrete relations.

We assume a relation, \prec , that constrains the changes of abstract predicate families between types. Rather than proving each method with the defining class’s precise type as in `L-DMETHOD`, we will consider proving for all types that are related:

$$\forall D \prec C. \{P \wedge \text{this} : D\} \vdash \{Q\} \quad (6.3)$$

We use \mathcal{C}, \mathcal{D} for variables over class names.

In particular the relation must be transitive and reflexive. Reflexive means (6.3) contains the original proof required for the method call. Transitive means that if we have a proof for C of the form (6.3), and D is related to C , then we also have a proof of the form (6.3) for D . The programmer must prove that all classes related by \prec can inherit this method.

The subtype must prove it is contained in the relation, to inherit a method without reverifying the method body.

We briefly consider two possible relations: (1) if the predicate family definitions are unchanged in the subtype; and (2) if the predicate family definitions are “proper” extensions.

(1) If we do not change any of the definitions a method uses then it should be valid to inherit the method. We define the relation $D \prec_{\{\alpha_1, \dots, \alpha_n\}} C$ as holding only if all the predicates $\alpha_1, \dots, \alpha_n$ are defined the same for both classes, that is

$$\Lambda^f \vdash D \prec_{\{\alpha_1, \dots, \alpha_n\}} C \iff \forall \alpha_i \in \{\alpha_1, \dots, \alpha_n\}. \Lambda^f(\alpha_i, C) = \Lambda^f(\alpha_i, D)$$

This relation is useful when a method uses no predicate families, because $\prec_{\{\}} is the universal relation. This allows any subtype to inherit that method. This is useful for the template method pattern [40]. This is where a class defines the outline or skeleton of an algorithm, leaving the details of the specific implementations to the subclasses. The method specifying the algorithm uses methods to manipulate the underlying data, so it does not depend on the representation, and hence the subclass can change it. Consider adding the following method into the `Cell` class.$

```
Object swap(Object new) {
    Object old;
    old = this.get();
```

```

    this.set(new);
    return old;
}

```

It is always valid to inherit this method into a subtype of `Cell` as it does not depend on the cell's representation. In particular, in a subclass, `twiceCell`, whose representation is:

$$\text{Value}(\text{this}; x) \stackrel{\text{def}}{=} \text{this.contents} \mapsto x * \text{this.contents2} \mapsto x$$

It is perfectly valid to inherit the `swap` method, but not the `set` or `get` method.

The first relation does not allow extension without reverifying methods. We have made some progress considering the following relation:

$$\begin{aligned} \Lambda^f \vdash D \prec_{\{\alpha_1, \dots, \alpha_n\}} C &\iff \\ \exists P. \Lambda^f(\alpha_i, C) = \lambda(x; \bar{x}). P_C \wedge \Lambda^f(\alpha_i, D) = \lambda(x; \bar{x}, \bar{y}). P_D \wedge & \\ P_C * P = P_D \wedge FV(P) \subseteq \{x, \bar{y}\} & \end{aligned}$$

This relation expresses that a class's definition contains its parent's definition, $P_C * P = P_D$, and that the extra specification does not depend on the free variables of the parent's definition, $FV(P) \subseteq \{x, \bar{y}\}$. The second condition is required so the parent cannot invalidate the subclass's predicate. For example in the `twiceCell` example its definition contained the `Cell`'s definition, but treating it like the `Cell`'s definition would break the property that `contents` and `contents2` have the same value. Using this relation one must alter the definition of `OPEN` and `CLOSE` to account for the extra state that the subtype has introduced.

These ideas are related to the stacks of specifications of Fähndrich and DeLine [36]. Their work allows extensions in a similar form to this relation. To properly express modularity, they make calls to the superclass's methods, which our semantics does not support.

Abstract predicate families help in allowing classes to change their implementation, and we leave open how best to tame them to reduce the verification burden.

7

Conclusion

This thesis has presented a logic, based on the local reasoning approach of O’Hearn, Reynolds, and Yang [65], to reason about Java. We have demonstrated several extensions to separation logic designed to allow reasoning about languages with class-based modularity. In particular, we have extended separation logic with three key features: abstract predicates; read-only points-to predicates; and abstract predicate families.

The first, abstract predicates, allows reasoning about encapsulation. They are relevant not only to class-based languages, but also to any language with dynamically hidden encapsulated state, such as abstract datatypes. They can be used to prove a `struct` in C is used as an abstract datatype, by disallowing code that depends on the `struct`’s layout. Abstract predicates do not just allow data encapsulation, but also protocol specification. They can be used to enforce a call sequence of methods in a similar fashion to `Typestates` [36].

The second extension, the read-only points-to predicate, allows datastructures read access to state, which is essential for many common idioms. Although important for datatypes and classes, the read-only predicate is essential for many concurrent algorithms [13, 14].

The third and final extension is the notion of an abstract predicate family. This concept allows abstract predicates to have multiple definitions indexed by class. This allows each class to define how it represents a datatype. We have shown the utility of this concept by reasoning about the Visitor design pattern.

7.1 Open questions

We conclude this thesis by outlining a series of open issues:

Integrating hypothetical frame rule with abstract predicates There are currently two ways to deal with information hiding in separation logic: abstract predicates and the hypothetical frame rule. These two approaches have different strengths and weaknesses. We conjecture the two approaches could successfully be combined in a logic, as they appear to be orthogonal concepts. However, the semantic development in this thesis is very different to the work on the hypothetical frame rule [66], so further work is required to prove this conjecture.

Higher-order separation logic and abstract predicates Birkedal and Torp-Smith [10] have shown how proofs similar to those expressed in Chapter 4 can be expressed in higher-order separation logic [8]. However, they do not encode the rules in this thesis, but simply show similarly

structured proofs. It would prove useful to encode the rules given in this thesis directly into higher-order separation logic in order to show that existential predicates do indeed capture all of the required ideas. Birkedal and Torp-Smith also remark that abstract predicate families can be encoded in higher-order separation logic, but provide no details of this encoding. A detailed exploration of possible encodings of abstract predicate families into higher order separation logic should be conducted.

Higher-order functions and code pointers The semantics in this thesis has avoided the view that objects use code pointers to represent dynamic dispatch on methods. Instead we have relied on the class hierarchy and behavioural subtyping. It would be interesting to explore the use of abstract predicate families to reason about code pointers or higher-order functions.

Concurrency O'Hearn [63] has shown how to reason about statically scoped concurrency primitives in separation logic. His rules use the same method of information hiding as the hypothetical frame rule: this style of reasoning cannot deal with Java's concurrency primitives as these are dynamically scoped. We are currently exploring integrating the information hiding aspects of abstract predicates with O'Hearn's concurrency rules to allow reasoning about Java's concurrency primitives.

Nomenclature

- * The separating conjunction, page 47
- \multimap The separating implication (magic wand), page 47
- $\otimes_{=}$ The iterated separating conjunction, page 72
- \prec The subtype relation, page 22
- \prec_1 The immediate subtype relation, page 22
- \prec_{θ} The substitution relation on stacks, page 59
- \square An empty stack, page 28
- α An abstract predicate name, page 64
- Δ A semantic predicate environment, page 69
- Δ_f A semantic predicate family environment, page 103
- δ A class table, page 22
- δ_C A constructor environment, page 22
- δ_F A field environment, page 22
- δ_M A method environment, page 22
- η A method or constructor name, page 49
- Γ A specification environment, page 48
- γ A typing environment, page 24
- Λ An abstract predicate environment, page 64
- Λ^f An abstract predicate family environment, page 92
- μ A method type, page 21
- π An MJ program, page 20
- τ A return or statement type, page 20
- θ A substitution, page 50
- ζ A return statement or skip, page 56

- A* The set of all abstract predicate names, page 69
- BS* A block scope, page 28
- C, D* A class name, page 20
- cd* A class definition, page 20
- CF* A closed frame, page 28
- cmd* A constructor definition, page 20
- config* A configuration, page 28
- e* An expression, page 20
- e* An expression, page 44
- EF* An exception frame, page 28
- f* A field name, page 20
- F* A frame, page 28
- fd* A field definition, page 20
- FS* A frame stack, page 28
- H* A heap, page 45
- H* A heap, page 28
- \mathcal{H} The set of all heaps, page 45
- I* A ghost stack, page 46
- m* A method name, page 20
- M** The abstract permissions partial commutative semi-group, page 85
- m* An element of the abstract permissions semi-group **M**, page 85
- md* A method definition, page 20
- md* A method definition, page 44
- MS* A method scope, page 28
- m_w The abstract write permission, page 86
- o* An object identifier, page 28
- OF* An open frame, page 28
- p** A countably infinite subset of \mathbb{P} , page 89
- p* A non-empty subset of \mathbb{P} , page 84
- P, Q, R* A separation logic formula, page 47

- \mathbb{P} The set of all permissions, page 84
- pe* A promotable expression, page 20
- ret* The return variable for specifications, page 49
- s* A statement, page 20
- s* A statement, page 44
- v* A value, page 28
- VS* A variable stack, page 28
- x* A program variable, page 20
- X, Y, ... A ghost variable, page 46

Bibliography

- [1] Martin Abadi and Luca Cardelli. *A theory of objects*. Springer, 1996. (Cited on pages 16, 17, and 62.)
- [2] Martin Abadi and K. Rustan M. Leino. A logic of object-oriented programs. In *Proceedings of CAAP/FASE*, volume 1214 of *LNCS*, pages 682–696. Springer, 1997. (Cited on page 17.)
- [3] Pierre America. Designing an object-oriented programming language with behavioural subtyping. In *Proceedings of the REX School/Workshop on Foundations of Object-Oriented Languages*, volume 489 of *LNCS*, pages 60–90. Springer, 1991. (Cited on pages 12 and 13.)
- [4] Krzysztof R. Apt. Ten years of Hoare’s logic: A survey: Part I. *ACM TOPLAS*, 3(4):431–483, 1981. (Cited on pages 14 and 53.)
- [5] Michael Barnett, Robert DeLine, Manuel Fähndrich, K. Rustan M. Leino, and Wolfrum Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6):27–56, 2004. (Cited on page 17.)
- [6] Michael Barnett, K. Rustan M. Leino, and Wolfrum Schulte. The Spec[#] programming system: An overview. In *Proceedings of CASSIS*, volume 3362 of *LNCS*, pages 49–69. Springer, 2005. (Cited on page 17.)
- [7] Martin Berger, Kohei Honda, and Nobuko Yoshida. A logical analysis of aliasing in imperative higher-order functions. In *Proceedings of ICFP*, volume 40(9) of *SIGPLAN Not.*, pages 280–293. ACM, 2005. (Cited on page 18.)
- [8] Bodil Biering, Lars Birkedal, and Noah Torp-Smith. BI hyperdoctrines and separation logic. In *Proceedings of ESOP*, volume 3444 of *LNCS*, pages 233–247. Springer, 2005. (Cited on pages 77 and 109.)
- [9] Gavin M. Bierman, Matthew J. Parkinson, and Andrew M. Pitts. MJ: An imperative core calculus for Java and Java with effects. Technical Report 563, University of Cambridge Computer Laboratory, 2004. (Cited on pages 27 and 44.)
- [10] Lars Birkedal and Noah Torp-Smith. Higher order separation logic and abstraction. Unpublished manuscript, 2005. (Cited on pages 77, 109, and 110.)
- [11] Lars Birkedal, Noah Torp-Smith, and John C. Reynolds. Local reasoning about a copying garbage collector. In *Proceedings of POPL*, volume 39(1) of *SIGPLAN Not.*, pages 220–231. ACM, 2004. (Cited on page 43.)
- [12] Richard Bornat. Proving pointer programs in Hoare logic. In *Proceedings of MPC*, volume 1837 of *LNCS*, pages 102–126. Springer, 2000. (Cited on page 15.)

- [13] Richard Bornat. Variables as resources in separation logic. In *Proceedings of MFPS*, 2005. to appear in ENTCS. (Cited on pages 15, 48, and 109.)
- [14] Richard Bornat, Cristiano Calcagno, Peter W. O’Hearn, and Matthew J. Parkinson. Permission accounting in separation logic. In *Proceedings of POPL*, volume 40(1) of *SIGPLAN Not.*, pages 259–270. ACM, 2005. (Cited on pages 68, 79, 80, 84, 85, 88, and 109.)
- [15] Chandrasekhar Boyapati, Barbara H. Liskov, and Liuba Shrira. Ownership types for object encapsulation. In *Proceedings of POPL*, volume 38(1) of *SIGPLAN Not.*, pages 213–223. ACM, 2003. (Cited on page 14.)
- [16] John Boyland. Checking interference with fractional permissions. In *Proceedings of SAS*, volume 2694 of *ACM*, pages 55–72. Springer, 2003. (Cited on page 88.)
- [17] Kim B. Bruce. Some challenging typing issues in object-oriented languages. *ENTCS*, 82(7), 2003. (Cited on page 97.)
- [18] Peter Buchlovsky and Hayo Thielecke. A type-theoretic reconstruction of the visitor pattern. In *Proceedings of MFPS*, 2005. to appear in ENTCS. (Cited on page 97.)
- [19] Lilian Burdy, Yoonsik Cheon, David Cok, Michael Ernst, Joe Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *ENTCS*, 80:73–89, 2003. (Cited on page 17.)
- [20] Rodney M. Burstall. Some techniques for proving correctness of programs which alter data structures. *Machine Intelligence*, 7:23–50, 1972. (Cited on page 15.)
- [21] Cristiano Calcagno and Peter W. O’Hearn. On garbage and program logic. In *Proceeding of FOSSACS*, volume 2030 of *LNCS*, pages 137–151. Springer, 2001. (Cited on page 47.)
- [22] David G. Clarke. *Object ownership and containment*. PhD thesis, University of New South Wales, Australia, 2001. (Cited on page 14.)
- [23] David G. Clarke and Sophia Drossopoulou. Ownership, encapsulation and the disjointness of type and effect. In *Proceedings of OOPSLA*, volume 37(11) of *SIGPLAN Not.*, pages 292–310. ACM, 2002. (Cited on pages 14 and 15.)
- [24] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *Proceedings of OOPSLA*, volume 33(10) of *SIGPLAN Not.*, pages 48–64. ACM, 1998. (Cited on page 14.)
- [25] Pierre-Jacques Courtois, F. Heymans, and David L. Parnas. Concurrent control with “readers” and “writers”. *Communications of the ACM*, 14(10):667–668, 1971. (Cited on page 79.)
- [26] Frank S. de Boer. A WP-calculus for OO. In *Proceeding of FOSSACS*, volume 1578 of *LNCS*, pages 135–149. Springer, 1999. (Cited on page 18.)
- [27] Frank S. de Boer and Cees Pierik. How to cook a complete Hoare logic for your pet OO language. In *Proceedings of FMCO*, volume 3188 of *LNCS*, pages 111–133. Springer, 2003. (Cited on page 18.)
- [28] Robert DeLine and Manuel Fähndrich. The Fugue protocol checker: Is your software baroque. Technical Report MSR-TR-2004-07, Microsoft Research, 2004. (Cited on page 18.)

- [29] David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and J. B. Saxe. Extended static checking. Technical Report 159, Compaq SRC, Palo Alto, CA, 1998. (Cited on pages 9 and 17.)
- [30] Krishna K. Dhara and Gary T. Leavens. Forcing behavioral subtyping through specification inheritance. In *Proceedings of ICSE*, pages 258–267. IEEE Computer Society, 1996. (Cited on page 13.)
- [31] Krishna K. Dhara and Gary T. Leavens. Forcing behavioral subtyping through specification inheritance. Technical Report 95-20c, Department of Computer Science, Iowa State University, March 1997. (Cited on page 13.)
- [32] Werner Dietl and Peter Müller. Universes: Lightweight ownership for JML. *Journal of Object Technology*, 2005. To appear. (Cited on pages 14 and 17.)
- [33] Sophia Drossopoulou, Susan Eisenbach, and Sarfraz Khurshid. Is the Java type system sound? *Theory and Practice of Object Systems*, 7(1):3–24, 1999. (Cited on pages 41 and 44.)
- [34] Sophia Drossopoulou, Tanya Valkevych, and Susan Eisenbach. Java type soundness revisited, 2000. Unpublished. (Cited on page 41.)
- [35] Jon Ellis and Linda Ho. JDBC 3.0 specification, 2001. <http://java.sun.com/products/jdbc/download.html>. (Cited on page 65.)
- [36] Manuel Fähndrich and Robert DeLine. Typestates for objects. In *Proceedings of ECOOP*, volume 3086 of *LNCS*, pages 465–490. Springer, 2004. (Cited on pages 17, 77, 81, 107, and 109.)
- [37] Robert B. Findler and Matthias Felleisen. Contract soundness for object-oriented languages. In *Proceedings of OOPSLA*, volume 36(11) of *SIGPLAN Not.*, pages 1–15. ACM, 2001. (Cited on page 13.)
- [38] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. A programmer’s reduction semantics for classes and mixins. Technical Report TR-97-293, Rice University, 1997. Corrected June, 1999. (Cited on pages 19, 41, and 44.)
- [39] Robert W. Floyd. Assigning meanings to programs. In *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 19–32. American Mathematical Society, 1967. (Cited on page 10.)
- [40] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, 1994. (Cited on pages 17, 91, 96, and 106.)
- [41] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification*. Addison Wesley, second edition, 2000. (Cited on pages 19, 21, 24, 25, and 31.)
- [42] Mark Grand. *Patterns in Java*. Wiley, second edition, 2002. (Cited on pages 16 and 65.)
- [43] John V. Guttag and James J. Horning. *Larch: languages and tools for formal specification*. Springer, 1993. (Cited on page 14.)
- [44] C. A. R. Hoare. Assertions: a personal perspective. In *Software pioneers: contributions to software engineering*, pages 356–366. Springer, 2002. (Cited on page 9.)

- [45] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969. (Cited on pages 10 and 13.)
- [46] C. A. R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1(4):271–281, 1972. (Cited on page 12.)
- [47] A. Igarashi, B.C. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *ACM TOPLAS*, 23(3):396–450, 2001. (Cited on pages 19, 24, and 41.)
- [48] Samin Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In *Proceedings of POPL*, volume 36(3) of *SIGPLAN Not.*, pages 14–26. ACM, 2001. (Cited on pages 15, 43, 44, 46, and 47.)
- [49] Bart Jacobs, K. Rustan M. Leino, and Wolfrum Schulte. Verification of multithreaded object-oriented programs with invariants. In *Proceedings of SAVCBS*, pages 2–9, November 2004. Published in Technical Report #04-09, Department of Computer Science, Iowa State University. (Cited on page 17.)
- [50] Brian W. Kernighan and Dennis M. Ritchie. *The C Programming Language, Second Edition*. Prentice-Hall, 1988. (Cited on page 73.)
- [51] Gerwin Klein and Tobias Nipkow. A machine-checked model for a Java-like language, virtual machine and compiler. Technical Report 0400001T.1, National ICT Australia, Sydney, March 2004. (Cited on page 41.)
- [52] John Lamping. Typing the specialization interface. In *Proceedings of OOPSLA*, volume 28(10) of *SIGPLAN Not.*, pages 201–214. ACM, 1993. (Cited on page 63.)
- [53] Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: A behavioral interface specification language for Java. Technical Report 98-06v, Iowa State University, Department of Computer Science, May 2003. (Cited on page 17.)
- [54] K. Rustan M. Leino. Data groups: Specifying the modification of extended state. In *Proceedings of OOPSLA*, volume 33(10) of *SIGPLAN Not.*, pages 144–153. ACM, 1998. (Cited on pages 14 and 63.)
- [55] Barbara H. Liskov and Jeannette M. Wing. A behavioral notion of subtyping. *ACM TOPLAS*, 16(6):1811–1841, 1994. (Cited on pages 12 and 13.)
- [56] Ronald Middlekoop. A proof system for object oriented programming using separation logic. Master’s thesis, Technische Universiteit Eindhoven, Eindhoven, the Netherlands, November 2003. (Cited on page 18.)
- [57] Ronald Middlekoop, Kees Huizing, and Ruurd Kuiper. A separation logic proof system for a class-based language. In *Proceedings of LRPP*, 2004. (Cited on page 18.)
- [58] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM TOPLAS*, 10(3):470–502, 1988. (Cited on page 77.)
- [59] Joseph M. Morris. A general axiom of assignment. In *Theoretical Foundations of Programming Methodology*, pages 25–51. D. Reidel, 1982. (Cited on pages 14 and 18.)
- [60] Peter Müller. *Modular Specification and Verification of Object-Oriented Programs*, volume 2262 of *LNCIS*. Springer, 2002. PhD thesis, FernUniversität Hagen. (Cited on pages 14 and 17.)

- [61] Peter Müller, Arnd Poetzsch-Heffter, and Gary T. Leavens. Modular specification of frame properties in JML. Technical Report 02-02a, Department of Computer Science, Iowa State University, October 2002. (Cited on pages 14 and 17.)
- [62] Tobias Nipkow and David von Oheimb. Java_{light} is type-safe—definitely. In *Proceedings of POPL*, pages 161–170. ACM, 1998. (Cited on page 41.)
- [63] Peter W. O’Hearn. Resources, concurrency and local reasoning. In *Proceedings of CONCUR*, volume 3170 of *LNCS*, pages 49–67. Springer, 2004. Invited Talk. (Cited on pages 68 and 110.)
- [64] Peter W. O’Hearn and David J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999. (Cited on page 15.)
- [65] Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Proceedings of CSL*, volume 2142 of *LNCS*, pages 1–19. Springer, 2001. (Cited on pages 10, 43, 44, 49, 72, and 109.)
- [66] Peter W. O’Hearn, Hongseok Yang, and John C. Reynolds. Separation and information hiding. In *Proceedings of POPL*, volume 39(1) of *SIGPLAN Not.*, pages 268–280. ACM, 2004. (Cited on pages 16, 63, 67, 74, 75, 76, and 109.)
- [67] Matthew J. Parkinson and Gavin M. Bierman. Separation logic and abstraction. In *Proceedings of POPL*, volume 40(1) of *SIGPLAN Not.*, pages 247–258. ACM, 2005. (Cited on pages 63 and 72.)
- [68] David L. Parnas. The secret history of information hiding. In *Software Pioneers: Contributions to Software Engineering*, pages 399–411. Springer, 2002. (Cited on page 75.)
- [69] Cees Pierik and Frank S. de Boer. A syntax-directed Hoare logic for object-oriented programming concepts. In *Proceedings of FMOODS*, volume 2884 of *LNCS*, pages 64–78. Springer, 2003. (Cited on page 18.)
- [70] Cees Pierik and Frank S. de Boer. Modularity and the rule of adaptation. In *Proceedings of AMAST*, volume 3116 of *LNCS*, pages 394–408. Springer, 2004. (Cited on page 18.)
- [71] Arnd Poetzsch-Heffter and Peter Müller. A programming logic for sequential Java. In *Proceedings of ESOP*, volume 1576 of *LNCS*, pages 162–176. Springer, 1999. (Cited on pages 17 and 60.)
- [72] Arnd Poetzsch-Heffter and Peter Müller. Logical foundations for typed object-oriented languages. In *Proceedings of PROCOMET*, volume 125 of *IFIP Conference Proceedings*, pages 404–423. Chapman & Hall, 1998. (Cited on pages 13, 17, and 53.)
- [73] Uday S. Reddy. Objects and classes in Algol-like languages. *Information and Computation*, 172(1):63–97, 2002. (Cited on page 77.)
- [74] Bernhard Reus. Class-based versus object-based: A denotational comparison. In *Proceedings of AMAST*, volume 2422 of *LNCS*, pages 473–488. Springer, 2002. (Cited on page 17.)
- [75] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings of LICS*, pages 55–74. IEEE Computer Society, 2002. (Cited on pages 43, 44, 49, 50, and 72.)

- [76] John C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In *Millennial Perspectives in Computer Science*, pages 303–321. Palgrave, 2000. (Cited on pages 15, 43, 44, and 47.)
- [77] John C. Reynolds. *The Craft of Programming*. Prentice-Hall International, London, 1981. (Cited on pages 14 and 53.)
- [78] John C. Reynolds. Idealized Algol and its specification logic. In *Tools and Notions for Program Construction*, pages 121–161. Cambridge University Press, 1982. (Cited on page 77.)
- [79] Matthew Smith and Sophia Drossopoulou. Cheaper reasoning with ownership types. In *Proceedings of IWACO*, 2003. Published in technical report UU-CS-2003-030, Institute of Information and Computing Sciences, Utrecht University. (Cited on page 14.)
- [80] Raymie Stata. Modularity in the presence of subclassing. Technical Report 145, Digital Systems Research Center, Palo Alto, CA, April 1997. (Cited on page 63.)
- [81] Robert E. Strom and Shaula Yemini. Tpestate: A programming language concept for enhancing software reliability. *IEEE Trans. Softw. Eng.*, 12(1):157–171, 1986. (Cited on pages 17 and 77.)
- [82] Donald Syme. Proving Java type soundness. Technical Report 427, University of Cambridge Computer Laboratory, 1997. (Cited on page 41.)
- [83] David von Oheimb. Hoare logic for mutual recursion and local variables. In *Proceedings of FSTTCS*, volume 1738 of *LNCS*, pages 168–180. Springer, 1999. (Cited on pages 50 and 60.)
- [84] David von Oheimb and Tobias Nipkow. Hoare logic for NanoJava: Auxiliary variables, side effects and virtual methods revisited. In *Proceedings of FME*, volume 2391 of *LNCS*, pages 89–105. Springer, 2002. (Cited on page 17.)
- [85] Glynn Winskel. *The Formal Semantics of Programming Languages: an introduction*. MIT press, 1993. (Cited on page 70.)
- [86] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994. (Cited on page 35.)
- [87] Hongwei Xi and Frank Pfenning. Eliminating array bound checking through dependent types. In *Proceedings of PLDI*, volume 33(5) of *SIGPLAN Not.*, pages 249–257. ACM, 1998. (Cited on page 9.)
- [88] Hongseok Yang. *Local reasoning for stateful programs*. PhD thesis, University of Illinois, July 2001. (Cited on pages 43, 56, and 57.)