

When separation logic met Java (by example)

Matthew Parkinson

Middlesex University

Abstract. Separation logic is a promising new approach to modular reasoning, but so far it has primarily been applied to low-level C-like languages. To extend separation logic to allow modular reasoning about object-oriented languages like Java, we must add behavioural subtyping to the logic. However, a naïve integration of behavioural subtyping and separation logic is too restrictive. In this paper we demonstrate how abstract predicate families provide an abstraction mechanism that addresses these restrictions, by mirroring dynamic dispatch in the logic. We demonstrate the utility of our approach with a series of examples, including the Visitor pattern.

1 Introduction

Modularity is key to building large software systems. Object-oriented languages like Java provide support for modularity through encapsulation and inheritance. This modularity allows portions of object-oriented programs to be modified without rewriting, or even recompiling, all the code. We need formal methods that reflect this modularity.

Separation logic [10, 14, 18] is a promising new approach to modular reasoning. It provides local reasoning for low level C-like languages. The essence of “local reasoning” is that in order 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”). The pre-condition of a method specifies its footprint: all the state (the fields of objects) the method touches during its execution. This contrasts standard Hoare logic: in Hoare logic $\{y = 7\} x = 3; \{y = 7\}$ is a valid judgement, but in separation logic $\{y.f \mapsto 7\} x.f = 3; \{y.f \mapsto 7\}$ is not a valid judgement, because the pre-condition does say the heap contains $x.f$. This constraint on pre-conditions means that anything disjoint from the pre-condition is guaranteed not to be changed. This allows method specifications to be adapted to different calling contexts, and avoids the need to reverify methods at every call site. This is essential for modular reasoning.

To achieve modular verification in object-oriented languages with dynamic dispatch, we must also enforce *behavioural subtyping* [2, 12]. Dynamic dispatch means that many methods can be potentially called at a single call site. Behavioural subtyping allows us to consider just a single call by enforcing compatibility between subtypes’ methods. We use a slightly generalised notion called *specification compatibility* [16], which allows the manipulation of logical variables¹ and intersection of specifications. Our generalisation is similar to *specification inheritance* [7].

¹ Sometimes called ghost or auxiliary variables.

However, a naïve combination of behavioural subtyping and separation logic does not work. The footprint of methods in subtypes are larger, in general than at supertypes, but behavioural subtyping requires that footprints are the same (or smaller). This is the *extended state problem* [11]. To address this problem, we use an abstraction called *abstract predicate families* [16].

An abstract predicate family is a predicate with a set of definitions indexed by class. In object-oriented programming a method body is selected based on the dynamic type of the first argument, the receiver. We reflect this in the logic: an abstract predicate family definition is selected based on the dynamic type of the first argument.

In this paper we show that abstract predicate families provide enough abstraction to introduce subtypes that are typically not considered to be *behavioural* subtypes. We demonstrate the utility of this approach with a series of examples. In particular, we present an example using the Visitor pattern [9]. The visitor pattern is an important test for program verification as it uses call-backs: the visitor’s `visit` method calls the data structure’s `accept` methods, and these `accept` methods call-back into the visitor’s `visit` methods. Call-backs can cause problems when reasoning using class invariants [4]. By using abstract predicate families, we do not encounter these complications.

In previous work with Bierman [16], we defined abstract predicate families and specification compatibility. This paper elaborates on the intuitions behind abstract predicate families and uses examples to demonstrate their utility. We also present more details on the integration with behavioural subtyping. In particular, we show how specification compatibility generalises behavioural subtyping to allow the intersection of specifications.

2 Separation logic for Java

Separation logic is an extension to Hoare logic that permits reasoning about shared mutable state. It extends Hoare logic by adding spatial connectives to the assertion language, which allow us to assert that two portions of the heap are disjoint. This separation provides the key feature of separation logic—*local reasoning*—specifications need only mention the state they access [14].

Previous works on separation logic have dealt with a low-level C-like language. In this section, we present our separation logic for a simple subset of Java. Space prevents us from giving the complete description of the separation logic, but it may be found in Parkinson’s thesis [15].

We use a simple subset of Java based on Middleweight Java (MJ) [5, 15]. We restrict MJ’s expressions to be stack variables and `null`,² and remove constructors. We present the full syntax in Figure 1. We write f and m to range over field and method names respectively. We use C, D to range over class names, and x, y to range over variable names. We will assume two functions: $fields(C)$, which returns the set of field names for the class, C ; and $method(C, m)$, which returns a four tuple of (1) the argument types, (2) the return type, (3) argument names, and (4) method body.

² This restriction simplifies the presentation of the rules.

Program	Expressions
<code>prog ::= cldf₁...cldf_n; \bar{s}</code>	<code>e ::= x null</code>
Class definition	Statements
<code>cldf ::= class C extends C' {$\overline{\text{fdef mdef}}$}</code>	<code>s ::= x=y.f; x=(C)y; x=new C();</code>
Method definition	<code> x.f=e; x=y.m(\bar{e}); C x;</code>
<code>mdef ::= Cm(C₁x₁, ..., C_nx_n){\bar{s} return x; }</code>	<code> {\bar{s}} ; if (e == e) s else s</code>
Field definition	
<code>fdef ::= Cf;</code>	

Fig. 1. Syntax of MJ subset

2.1 The storage model

Separation logic is a logic of partial heaps, that is, heaps with (potentially) dangling pointers. The spatial connectives are based on the composition of disjoint partial heaps, which allows the logic to specify the splitting, or separation, of the heap.

To reason about Java, we must extend this partial heap model to represent the structure of objects. In addition to the consideration of dangling pointers, we also consider heaps with partial objects, that is objects with fields missing. A method might not modify all the fields of an object. We consider the heap to store two forms of information: (1) the values of fields; and (2) the type of objects. Thus a heap, h , is a pair of functions: (1) a partial function from object identifier and field names to values (for simplicity we take Values to be the object identifiers and null); and (2) a partial function from object identifier to class.

$$\text{Heaps} \stackrel{\text{def}}{=} (\text{OIDS} \times \text{FieldNames} \rightarrow_{\text{fin}} \text{Values}) \times (\text{OIDS} \rightarrow_{\text{fin}} \text{Class})$$

We use h to range over heaps, and h_v and h_t for the first and second components of a heap respectively. We write $h * h'$ for the composition of two disjoint heaps:

$$(h_v, h_t) * (h'_v, h'_t) \stackrel{\text{def}}{=} \begin{cases} (h_v \cup h'_v, h_t) & (\text{dom}(h_v) \cap \text{dom}(h'_v) = \emptyset) \wedge (h_t = h'_t) \\ \text{undefined} & \text{otherwise} \end{cases}$$

We only consider the composition where the value heaps, h_v , contain disjoint pairs of object identifiers and field names. We do not split the typing information, because it is immutable and hence valid to share.

A stack is a function from (program) variables to values. In addition to program variables, we also use logical variables,³ which are variables that can be quantified over and used to relate values between pre- and post-conditions. In general, we will use upper case letters for logical variables to distinguish them from program variables. Unlike other presentations [10], we do not interpret logical variables using the stack because the fragment of Java does not have globally scoped variables. Instead, we define a logical interpretation that is a function from logical variable names to values

$$\begin{aligned} \text{Stacks} &\stackrel{\text{def}}{=} \text{ProgVarNames} \rightarrow \text{Values} \\ \text{Interpretations} &\stackrel{\text{def}}{=} \text{LogVarNames} \rightarrow \text{Values} \end{aligned}$$

³ Sometimes called ghost or auxiliary variables.

We define a state as a triple consisting of a stack, a heap and an interpretation.

$$\text{States} \stackrel{\text{def}}{=} \text{Heaps} \times \text{Stacks} \times \text{Interpretations}$$

2.2 The assertion language

A predicate is interpreted as a set of states, and formulae are given by the following grammar where e ranges over variables and null.

$$P, Q ::= \text{true} \mid \text{false} \mid \neg P \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \\ \mid P * Q \mid e.f \mapsto e' \mid e : C \mid \exists X. P \mid \forall X. P$$

The intuitionistic⁴ connectives ($\neg, \vee, \wedge, \Rightarrow$) and quantifiers (\forall, \exists) are interpreted in the usual way [10]. In addition to the intuitionistic connectives, we have the new spatial connective $*$, along with the predicate \mapsto , and a predicate for Java’s dynamic type information, $e : C$. Taking these in reverse order:

$e.f \mapsto e'$ consists of all the states where the heap consists of at least the single mapping from the field f of the object given by the meaning of e to the value given by the meaning of e' .

$P * Q$ 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 $*$ to join smaller heaps.

$e : C$ means that the result of e is an object whose dynamic type is precisely C .

We use $P(_)$ as a shorthand for $\exists X. P(X)$, for example, $e \mapsto _$ means $\exists X. e \mapsto X$.

2.3 The rules

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.

$$\text{L-FRAME} \quad \frac{\vdash \{P\} \bar{s} \{Q\}}{\vdash \{P * R\} \bar{s} \{Q * R\}}$$

where \bar{s} does not modify the free variables of R , i.e. $\text{mods}(\bar{s}) \cap FV(R) = \emptyset$.

Note 1. $\text{mods}(\bar{s})$ denotes the set of stack variables assigned by a given command, \bar{s} , e.g. $\text{mods}(x = e) = \{x\}$. However assignment through a stack variable to a field is not counted: $\text{mods}(x.f = e) = \emptyset$. See [15] for the full definition.

⁴ Intuitionistic separation logic [10] allows memory leaks; it admits weakening, $P * Q \Rightarrow Q$ holds. This is required for Java because it lacks an explicit “dispose” primitive. Note: we do not have to worry about memory leaks, because Java is a garbage collected language

The frame rule’s side-condition is inherited from Hoare logic. It is required because $*$ only describes the separation of heap locations and not variables; see [6, 17] for more details. Without the side-condition we could make the following incorrect derivation:

$$\frac{\{true\}x = \text{new } C() \{x.f \mapsto _ \}}{\{true * x.f \mapsto _ \}x = \text{new } C() \{x.f \mapsto _ * x.f \mapsto _ \}}$$

The two fields in the post-condition of the conclusion are of different objects, but without the side-condition we can wrongly infer them to be the same field.

By using the frame 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.

A judgement in our assertion language is written as follows:

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

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

$$\Gamma := \epsilon \mid \{P\} C.m(\bar{x}) \{Q\}, \Gamma$$

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

We present the rules and axioms for Java in Figure 2. The top section presents the axioms associated to particular Java statements. We give small axioms for field manipulation in the style of O’Hearn, Reynolds, and Yang [14]. Field write, L-FWRITE, 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 axiom specifies that no other fields are modified, and hence it can be extended by the frame rule to talk about additional state. The field access axiom, L-FREAD, simply requires the field’s state to be known and sets the variable equal to its contents. The logical 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. The cast axiom, L-CAST, ensures that a cast will complete successfully: the type of the object identifier being cast must be a subtype of the target type or null. The rule for constructing an object simply constructs all the fields of the class. The final axiom of the section, L-CALL, allows the call of a dynamically dispatched method without case analysis. In the next section we will place a constraint on method environments to make this axiom sound. The logical variable Y is used so that x and \bar{e} can mention y . If x and \bar{e} do not mention y , then we can remove the substitution $[Y/y]$ and equality $Y = y$ from the axiom.

The second section of Figure 2 defines the rules for introducing assumptions about methods. We define a new judgement $\Gamma_1 \Vdash \Gamma_2$ to allow the introduction of mutually recursive methods. $\Gamma_1 \Vdash \Gamma_2$ means that the bodies of the methods in Γ_2 can be verified

Java Commands	
L-FWRITE	$\Gamma \vdash \{x.f \mapsto _ \} x.f = e; \{x.f \mapsto e\}$
L-FREAD	$\Gamma \vdash \{X = x \wedge X.f \mapsto Y\} y = x.f; \{X.f \mapsto Y \wedge y = Y\}$
L-UPCAST	$\Gamma \vdash \{P[x/y] \wedge (x : C \vee x = \text{null})\} y = (C')x; \{P\}$ provided $C \prec C'$.
L-NEW	$\Gamma \vdash \{\text{true}\} x = \text{new } C(); \{x : C * x.f_1 \mapsto _ * \dots * x.f_n \mapsto _ \}$ provided $\text{fields}(C) = f_1, \dots, f_n$
L-CALL	$\{P\} C.m(\bar{w})\{Q\}, \Gamma \vdash \{\theta(P) \wedge Y = y \wedge x \neq \text{null}\} y = x.m(\bar{e}); \{\theta(Q)\}$ provided $\theta = [\bar{e}[Y/y], x[Y/y], y/\bar{w}, \text{this}, \text{ret}]$, and x has static type C
Method introduction	
L-DMETHOD	$\frac{\Gamma \vdash \{P \wedge \text{this} : C\} \bar{C} \bar{y}; \bar{y} = \bar{x}; \bar{s}[\bar{y}/\bar{x}] \{Q[x/\text{ret}]\}}{\Gamma \Vdash \{P\} C.m(\bar{x}) \{Q\}}$ provided $\text{method}(C, m) = (\bar{C}, _, \bar{x}, \bar{s} \text{ return } x)$ and \bar{y} is fresh.
L-DSPLIT	$\frac{\Gamma \Vdash \Gamma_1 \quad \Gamma \Vdash \Gamma_2}{\Gamma \Vdash \Gamma_1, \Gamma_2}$
L-DINTRO	$\frac{\Gamma, \Gamma' \Vdash \Gamma' \quad \Gamma, \Gamma' \vdash \{P\} \bar{s}\{Q\}}{\Gamma \vdash \{P\} \bar{s}\{Q\}}$

Fig. 2. Java rules and axioms

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 additional variable declarations are used to ensure that the call-by-value semantics is respected. The remaining two rules are used to introduce and manipulate these definitions. These rules are similar in style to those used by von Oheimb [19] to reason about mutually recursive procedures.

For completeness we present the standard rules from Hoare and separation logic in Figure 3.

3 Behavioural subtyping

Next we present *specification compatibility*: a generalised notion of behavioural subtyping. For the purpose of this exposition we focus on method specifications and ignore class invariants and abstraction functions [2, 12] as we, later, use abstract predicate families to deal with these issues.

A method specification, $\{P_D\} _ \{Q_D\}$, is said to be a subtype of another method specification, $\{P_C\} _ \{Q_C\}$, iff the pre-condition of the supertype implies the pre-condition of the subtype, $P_C \Rightarrow P_D$, and the subtype's post-condition implies the supertype's, $Q_D \Rightarrow Q_C$. We can see behavioural subtyping of specifications as simply the rule of consequence in Hoare logic, L-CONSEQUENCE in Figure 3. We can generalise this notion to allow the other structural rules: L-FRAME and L-VARELIM in Figure 3.

Standard commands	
L-ASSIGN	$\Gamma \vdash \{P[e/x] \ x=e; \{P\}$
L-BLOCK	$\frac{\Gamma \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{P\} \{\bar{s}\} \{Q\}}$
L-SKIP	$\Gamma \vdash \{P\}; \{P\}$
L-IF	$\frac{\Gamma \vdash \{P \wedge x=y\} s_1 \{Q\} \quad \Gamma \vdash \{P \wedge x \neq y\} s_2 \{Q\}}{\Gamma \vdash \{P\} \text{if } (x == y) s_1 \text{ else } s_2 \{Q\}}$
L-SEQ1	$\frac{\Gamma \vdash \{P\} s_1 \{R\} \quad \Gamma \vdash \{R\} s_2 \dots s_n \{Q\} \quad s_1 \neq C \ x;}{\Gamma \vdash \{P\} s_1 s_2 \dots s_n \{Q\}}$
L-SEQ2	$\frac{\Gamma \vdash \{P\} s_1 \dots s_n \{Q\} \quad x \notin FV(P) \cup FV(Q)}{\Gamma \vdash \{P\} C \ x; s_1 \dots s_n \{Q\}}$
<hr/> Structural rules	
L-CONSEQUENCE	$\frac{P \Rightarrow P' \quad \Gamma \vdash \{P'\} \bar{s} \{Q'\} \quad Q' \Rightarrow Q}{\Gamma \vdash \{P\} \bar{s} \{Q\}}$
L-VARELIM	$\frac{\Gamma \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{\exists X \cdot P\} \bar{s} \{\exists X \cdot Q\}} \quad \text{provided } X \text{ is not free in } \bar{s}.$
L-FRAME	$\frac{\Gamma \vdash \{P\} \bar{s} \{Q\}}{\Gamma \vdash \{P * R\} \bar{s} \{Q * R\}} \quad \text{provided } \text{mods}(\bar{s}) \cap FV(R) = \emptyset$

Fig. 3. Standard rules

Specification compatibility generalises behavioural subtyping in two important ways: it allows the introduction of new logical variables in subclasses' specifications; and it allows the intersection of specifications. We demonstrate these at the end of the section.

Definition 1 (Specification compatibility). *We define specification compatibility,*
 $\vdash \{P_D\} \bar{s} \{Q_D\} \Rightarrow \{P_C\} \bar{s} \{Q_C\}$, *iff*

$$\forall \bar{s}, \Gamma \cdot (\text{mods}(\bar{s}) \subseteq \{\text{ret}\}) \Rightarrow \frac{\Gamma \vdash \{P_D\} \bar{s} \{Q_D\}}{\Gamma \vdash \{P_C\} \bar{s} \{Q_C\}}$$

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 \bar{s} and Γ in the definition restricts the derivation to the structural rules in Figure 3. The frame rule, L-FRAME, can only introduce formulae that do not mention *ret*: all methods potentially alter this variable.

We will restrict method environments to *well-behaved environments*, where methods in subtypes have compatible specifications with their supertypes' specification.

Definition 2 (Well-behaved environments). Γ is well-behaved iff for each specification $\{P_C\}C.m(\bar{x})\{Q_C\}$ in Γ , if $D \prec C$, then $\{P_D\}D.m(\bar{x})\{Q_D\}$ is in Γ and $\vdash \{P_D\} - \{Q_D\} \implies \{P_C\} - \{Q_C\}$

For simplicity, we assume parameter names are not changed in subclasses.

As mentioned earlier, our generalised notion of behavioural subtyping allows the following additional ideas:

3.1 Logical Variable Manipulation

Consider the following specifications

$$\{\text{this}.f \mapsto _ \} - \{\text{this}.f \mapsto _ \} \quad (1)$$

$$\{\text{this}.f \mapsto X \} - \{\text{this}.f \mapsto X \} \quad (2)$$

If code expects the behaviour of (1), then clearly it would be satisfied with (2), after all (2) is a refinement of (1). However, the standard idea of behavioural subtyping would not permit this refinement as $\text{this}.f \mapsto _ \implies \text{this}.f \mapsto X$ does not hold. This is valid in specification compatibility as it is a use of the variable elimination rule, L-VARELIM.

3.2 Specification intersection

Consider two specifications:

$$\{P_C\} - \{Q_C\} \quad (3)$$

$$\{P_D\} - \{Q_D\} \quad (4)$$

How can a single specification capture both of the above? Consider

$$\left\{ \begin{array}{l} (P_C \wedge X = 1) \\ \vee (P_D \wedge X = 2) \end{array} \right\} - \left\{ \begin{array}{l} (Q_C \wedge X = 1) \\ \vee (Q_D \wedge X = 2) \end{array} \right\} \quad (5)$$

We use a fresh logical variable X to allow us to encode the intersection of the specification. Semantically, logical variables are universally quantified at the level of specifications, i.e.

$$\forall X. \left\{ \begin{array}{l} (P_C \wedge X = 1) \\ \vee (P_D \wedge X = 2) \end{array} \right\} - \left\{ \begin{array}{l} (Q_C \wedge X = 1) \\ \vee (Q_D \wedge X = 2) \end{array} \right\}$$

Hence, we can understand the specification as

$$(\{P_C\} - \{Q_C\}) \bigwedge (\{P_D\} - \{Q_D\})$$

We use this notation as a shorthand in the examples later. We can show that (5) is specification compatible with both (3) and (4). We present the derivation for (3).

$$\begin{array}{c}
\frac{\vdash \left\{ \begin{array}{l} (P_C \wedge X = 1) \\ \vee (P_D \wedge X = 2) \end{array} \right\} - \left\{ \begin{array}{l} (Q_C \wedge X = 1) \\ \vee (Q_D \wedge X = 2) \end{array} \right\}}{\text{L-FRAME}} \\
\frac{\vdash \left\{ \begin{array}{l} X = 1 * \\ \left(\begin{array}{l} (P_C \wedge X = 1) \\ \vee (P_D \wedge X = 2) \end{array} \right) \end{array} \right\} - \left\{ \begin{array}{l} X = 1 * \\ \left(\begin{array}{l} (Q_C \wedge X = 1) \\ \vee (Q_D \wedge X = 2) \end{array} \right) \end{array} \right\}}{\text{L-CONSEQUENCE}} \\
\frac{\vdash \{P_C \wedge X = 1\} - \{Q_C \wedge X = 1\}}{\text{L-VARELIM}} \\
\frac{\vdash \{\exists X \cdot P_C \wedge X = 1\} - \{\exists X \cdot Q_C \wedge X = 1\}}{\text{L-CONSEQUENCE}} \\
\vdash \{P_C\} - \{Q_C\}
\end{array}$$

The key to the proof is the L-FRAME rule, which lets us preserve the value of the logical variable across the call. In standard Hoare logic the invariance rule [3] could be used.

Specification inheritance of Dhara and Leavens [7] also allows the intersection of specifications.

4 Abstract predicate families

So far, we have presented a separation logic for Java and enforced behavioural subtyping. As it stands, it is very hard to create any valid subtypes, because the specifications are too concrete. In this section, we define our abstraction mechanism, *abstract predicate families*.

We use α to range over abstract predicate family names. One can think of an abstract predicate family as a predicate that existentially hides the type:

$$e.\alpha(\bar{e}) \approx \exists C \cdot \alpha_C(e; \bar{e}) \wedge e : C \quad (6)$$

where α_C is the definition of α for class C . This intuition mirrors dynamic dispatch in object-oriented languages, where $e.m(\dots)$ is a call to a method m , but the body is selected by the dynamic type of the receiver, e . Abstract predicate families reflect this in the logic, by allowing predicates whose definition is selected by the dynamic type of the “receiver”.

We define abstract predicate family definitions, Λ , with the following syntax.

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

Λ 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, $x; \bar{x}$. The first argument is distinguished, as it is used to index by class. We use a semi-colon, $;$, to separate this distinguished argument, and commas to separate the remaining arguments. We treat Λ 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.

This semantic intuition, (6), leads to the following pair of axioms:

$$\text{OPEN} \quad \Lambda \models (x : C \wedge x.\alpha(\bar{x})) \Rightarrow \Lambda(\alpha, C)[x; \bar{x}]$$

$$\text{CLOSE} \quad \Lambda \models (x : C \wedge \Lambda(\alpha, C)[x; \bar{x}]) \Rightarrow x.\alpha(\bar{x})$$

where $\alpha, C \in \text{dom}(\Lambda)$.

To OPEN or CLOSE a predicate we must know which class contains the definition, and must have that definition in scope.

In our examples later, we find it useful to consider the predicate to have many different arities. We add the following pair of axioms to alter a predicate's arity.

$$\begin{array}{l} \text{WIDEN} \\ \text{NARROW} \end{array} \quad \frac{}{\begin{array}{l} A \models x.\alpha(\bar{x}) \Rightarrow \exists \bar{y} \cdot x.\alpha(\bar{x}, \bar{y}) \\ A \models \exists \bar{y} \cdot x.\alpha(\bar{x}, \bar{y}) \Rightarrow x.\alpha(\bar{x}) \end{array}}$$

As we allow changes in arity, we must provide an operation to provide fewer or more arguments than expected. We use square brackets to denote this unusual argument application operation.

$$(\alpha_C(x; \bar{x}) \stackrel{\text{def}}{=} P)[e; \bar{e}] \stackrel{\text{def}}{=} \begin{cases} P[e/x, \bar{e}_1/\bar{x}] & |\bar{e}_1| = |\bar{x}| \text{ and } \bar{e} = \bar{e}_1, \bar{e}_2 \\ \exists \bar{x}' \cdot P[e/x, (\bar{e}, \bar{x}')/\bar{x}] & |\bar{e}, \bar{x}'| = |\bar{x}| \end{cases}$$

If we give a predicate more arguments than its definition requires, it ignores them. If too few, it existentially quantifies the missing arguments. This definition of substitution is used in the definition of OPEN and CLOSE.

We present the following rule for introducing abstract predicate families

$$\text{ABSTRACT WEAKENING} \quad \frac{A; \Gamma \vdash \{P\} \bar{s} \{Q\}}{A, A'; \Gamma \vdash \{P\} \bar{s} \{Q\}} \quad \text{provided } \text{dom}(A') \text{ and } \text{dom}(A) \text{ are disjoint.}$$

This rule allows us to add more predicate definition without affecting a proofs validity. For details of the modularity and scoping of abstract predicate families see [15, 16].

5 Examples

For legibility, in the following examples we use full Java statements and expressions and a primitive type of integer, `int`. These can be encoded in the obvious way [15].

5.1 Cell/Recell

We begin with a simple example of subtyping: `Cell` and `Recell` [1]. We give the source code in Figure 4. We ignore `NullCell` until the next subsection.

If we give the obvious specification to `Cell`

```
{this.cnts↦_}Cell.set(n){this.cnts↦n}
```

then there is no specification that `Recell` satisfies and is a behavioural subtype. This is the extended state problem: the subclass is modifying more state. Therefore, we need to introduce an abstract predicate family into the specification. We present the correct specifications in Figure 5.

We have to validate four methods: the `set` and `get` methods of both `Cell` and `Recell`. Details of these proofs can be found in [15, 16]. Even though the bodies of

```

class Cell extends Object {
  Object cnts;
  void set(Object x)
  {this.cnts = x;}
  Object get()
  {return this.cnts;}
}
class NullCell extends Cell {
  Object get() {return null;}
}
class Recell extends Cell {
  Object bk;
  void set(Object o) {
    temp = this.cnts;
    this.bk = temp;
    this.cnts = x;
  }
}

```

Fig. 4. Source code for Cell classes

Method	Specification
<code>Cell.set(y)</code>	$\{this.Val(-)\} - \{this.Val(y)\}$
<code>Cell.get()</code>	$\{this.Val(X)\} - \{this.Val(X) \wedge ret = X\}$
<code>Recell.set(y)</code>	$\{this.Val(X, -)\} - \{this.Val(y, X)\}$
<code>Recell.get()</code>	$\{this.Val(X, Y)\} - \{this.Val(X, Y) \wedge ret = X\}$
<code>Nullcell.set(y)</code>	$(\{this.Val(-)\} - \{this.Val(y)\})$ $\wedge (\{this.NVal()\} - \{this.NVal()\})$
<code>Nullcell.get()</code>	$(\{this.Val(X)\} - \{this.Val(X) \wedge ret = X\})$ $\wedge (\{this.NVal()\} - \{this.NVal() \wedge ret = null\})$
$Val_{Cell}(x; y) \stackrel{def}{=} x.cnts \mapsto y$	$Val_{Recell}(x; y, z) \stackrel{def}{=} x.cnts \mapsto y * x.bk \mapsto z$
$Val_{Nullcell}(x;) \stackrel{def}{=} false$	$NVal_{Nullcell}(x;) \stackrel{def}{=} x.cnts \mapsto -$

Fig. 5. Specifications and predicate definitions for Cell classes

`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.

We must prove that the method specifications are compatible, defined in the sense of definition 2. The two `get` methods have the same specification, so they are compatible. The compatibility of the `set` method follows as

$$\frac{\frac{\vdash \{this.Val(X, -)\} - \{this.Val(n, X)\}}{\vdash \{this.Val(-, -)\} - \{this.Val(n, -)\}} \text{L-VARELIM}}{\vdash \{this.Val(-)\} - \{this.Val(n)\}} \text{L-CONSEQUENCE}$$

Above L-CONSEQUENCE uses WIDEN on the pre-condition, and NARROW on the post-condition.

5.2 NullCell

Normally, we would not consider the cell that always returns null, `NullCell`, to be a behavioural subtype of `Cell`. However, we may wish it to inherit from `Cell`. Abstract

```

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); }
}

class Visitor {
  void visitC(Const x) {;}
  void visitP(Plus x) {;}
}
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);
  }
}

```

Fig. 6. Source code for formulae and visitor classes.

predicate families allow us to see `NullCell` as a behavioural subtype (using specification compatibility), while preventing it being used in place of a `Cell`.

In Figure 5, we define the predicate family *Val* for `NullCell` to be false. This means that it is never possible to get a satisfied instance of *Val* where its first parameter is a `NullCell`, and hence we cannot use a `NullCell` in place of a `Cell`. We give the `NullCell` a new predicate family, *NVal*, to represent its internals, and specify its methods as the intersection, as defined in section 3.2, of its supertype’s method’s specification and its own specification with *NVal*.

Abstract predicate families allow implementations to be inherited, while not preserving behavioural subtyping. They distinguish implementation inheritance from specification inheritance.

5.3 Visitor pattern

Now let us consider an extended example using the visitor design pattern [9]. We consider a visitor over a very simple syntax of formulae with just constants and addition of two terms. We present the source code in Figure 6. 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. We define three classes to represent an abstract syntax tree of these formulae: `Ast` is used as a common parent for the term constructors; `Const` represents constant terms; and `Plus` 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.

Before we can formally specify the visitor, we must extend the logic with the formulae’s syntax:

$$\begin{aligned}
\tau &::= n \mid \tau \oplus \tau \\
\mu &::= \bullet \mid \tau \oplus \mu \mid \mu \oplus \tau
\end{aligned}$$

Method	Specification
<code>visitC(x)</code>	$\{x : \text{Const} \wedge x.\text{Ast}(\tau) * \text{this}.\text{Visitor}(\mu)\} - \{\text{this}.\text{Visited}(x, \tau, \mu)\}$
<code>visitP(x)</code>	$\{x : \text{Plus} \wedge x.\text{Ast}(\tau) * \text{this}.\text{Visitor}(\mu)\} - \{\text{this}.\text{Visited}(x, \tau, \mu)\}$
<code>accept(x)</code>	$\{\text{this}.\text{Ast}(\tau) * x.\text{Visitor}(\mu)\} - \{x.\text{Visited}(\text{this}, \tau, \mu)\}$
	$\text{Visitor}_{\text{calc}}(x; \mu) \stackrel{\text{def}}{=} x.\text{amount} \mapsto \text{lcalc}(\mu)$
	$\text{Visited}_{\text{calc}}(x; y, \tau, \mu) \stackrel{\text{def}}{=} x.\text{amount} \mapsto (\text{lcalc}(\mu) + \text{calc}(\tau)) * \text{Ast}(y; \tau)$
where	
$\text{calc}(n) \stackrel{\text{def}}{=} n$	$\text{lcalc}(\bullet) \stackrel{\text{def}}{=} 0$
$\text{calc}(\tau_1 \oplus \tau_2) \stackrel{\text{def}}{=} \text{calc}(\tau_1) + \text{calc}(\tau_2)$	$\text{lcalc}(\mu_1 \oplus \tau) \stackrel{\text{def}}{=} \text{lcalc}(\mu_1)$
	$\text{lcalc}(\tau \oplus \mu_1) \stackrel{\text{def}}{=} \text{calc}(\tau) + \text{lcalc}(\mu_1)$

Fig. 7. Specifications for accept and visit methods.

We use τ to represent the structure of an abstract syntax tree, and μ for a context: a tree with a hole, \bullet , in it. We use, \oplus , to distinguish it from the arithmetic operation of addition, and use $\mu[\mu']$ to mean replace \bullet by μ' in μ . We define the *Ast* predicate family for the three classes.

$$\begin{aligned}
\text{Ast}_{\text{Ast}}(x; \tau) &\stackrel{\text{def}}{=} \text{false} \\
\text{Ast}_{\text{Const}}(x; \tau) &\stackrel{\text{def}}{=} \exists n \cdot x.v \mapsto n \wedge \tau = n \\
\text{Ast}_{\text{Plus}}(x; \tau) &\stackrel{\text{def}}{=} \exists i j \tau_l \tau_r \cdot x.l \mapsto i * x.r \mapsto j * \text{Ast}(i, \tau_l) * \text{Ast}(j, \tau_r) \wedge \tau = \tau_l \oplus \tau_r
\end{aligned}$$

Note 2. Defining the predicate for the `Ast` class as false prevents any invocation of its methods.

We give the specifications for the `visit` and `accept` methods in Figure 7. 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 accumulate the value of the expression as we traverse it, rather than calculating it in a bottom up fashion. We can remove the μ parameter from the predicates if we do not want context sensitive visitors, that is, if the state of the visitor did not depend on its context. One might have expected the methods all to have post-conditions of the form `this.Visited(...)` $*$ $x.\text{Ast}(\tau)$. However, this kind of specification prevents us altering the structure of the expression.

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.

$$\left. \begin{aligned}
&\{\text{this}.\text{Ast}(\tau) * x.\text{Visitor}(\mu) \wedge \text{this} : \text{Plus}\} \\
&\left. \left. \begin{aligned}
&\{(x.\text{Ast}(\tau) * \text{this}.\text{Visitor}(\mu) \wedge x : \text{Plus})[\text{this}, x/x, \text{this}]\} \\
&x.\text{visitP}(\text{this}); \\
&\{(\text{this}.\text{Visited}(x, \tau, \mu))[\text{this}, x/x, \text{this}]\} \\
&\{x.\text{Visited}(\text{this}, \tau, \mu)\}
\end{aligned} \right\} \left. \begin{array}{l} \text{L-CALL} \\ \text{L-CONSEQUENCE} \end{array} \right\}
\end{aligned}
\right.$$

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 selection 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 calculates the value of the expression. We define the *Visitor* and *Visited* predicates for this class in Figure 7. The *calc* function calculates the value of the tree, and 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 as

$$\left. \begin{array}{l} \{x : \text{Plus} \wedge x.\text{Ast}(\tau) * \text{this}.\text{Visitor}(\mu) \wedge \text{this} : \text{Calc}\} \\ \left\{ \begin{array}{l} \{\exists i, j, \tau_1, \tau_2 \cdot Q * i.\text{Ast}(\tau_1) * j.\text{Ast}(\tau_2) * \text{this}.\text{Visitor}(\mu)\} \\ \{Q * i.\text{Ast}(\tau_1) * j.\text{Ast}(\tau_2) * \text{this}.\text{Visitor}(\mu)\} \\ x.l.\text{accept}(\text{this}); x.r.\text{accept}(\text{this}); \\ \{Q * i.\text{Ast}(\tau_1) * \text{this}.\text{Visited}(j, \tau_2, \mu[\tau_1 \oplus \bullet])\} \\ \{\exists i, j, \tau_1, \tau_2 \cdot Q * i.\text{Ast}(\tau_1) * \text{this}.\text{Visited}(j, \tau_2, \mu[\tau_1 \oplus \bullet])\} \end{array} \right\} \subseteq \end{array} \right\} \left. \begin{array}{l} \text{L-VAR-ELIM} \\ \text{L-CONSEQUENCE} \end{array} \right\}$$

where $P \stackrel{\text{def}}{=} x : \text{Plus} \wedge \text{this} : \text{Calc} \wedge \tau = \tau_1 \oplus \tau_2$

and $Q \stackrel{\text{def}}{=} P * x.l \mapsto i * x.r \mapsto j$

Above L-CONSEQUENCE uses OPEN with the definition of *Visitor* for Calc in the implication between the pre-conditions, and for the post-conditions uses both OPEN and CLOSE with the definition of *Visited* and $lcalc(\mu) + calc(\tau) = lcalc(\mu[\tau \oplus \bullet])$ which can be shown by induction on μ .

We present an outline of the proof for the two function calls:

$$\left. \begin{array}{l} \{P * x.l \mapsto i * x.r \mapsto j * i.\text{Ast}(\tau_1) * j.\text{Ast}(\tau_2) * \text{this}.\text{Visitor}(\mu)\} \\ \left\{ \begin{array}{l} \{x.l \mapsto i * i.\text{Ast}(\tau_1) * \text{this}.\text{Visitor}(\mu, \bullet)\} \\ x.l.\text{accept}(\text{this}); \\ \{x.l \mapsto i * \text{this}.\text{Visited}(i, \tau_1, \mu)\} \end{array} \right\} \subseteq \end{array} \right\} \left. \begin{array}{l} \text{L-FRAME} \\ (9) \end{array} \right\}$$

To prove (8), we expand the code and actually prove $t = x.l ; t.\text{accept}(\text{this}) ;$. We omit the details of the proof.

$$\left. \begin{array}{l} \{P * x.l \mapsto i * x.r \mapsto j * j.\text{Ast}(\tau_2) * i.\text{Ast}(\tau_1) * \text{this}.\text{Visitor}(\mu[\tau_1 \oplus \bullet])\} (10) \\ \left\{ \begin{array}{l} \{x.r \mapsto j * j.\text{Ast}(\tau_2) * \text{this}.\text{Visitor}(\mu[\tau_1 \oplus \bullet])\} \\ x.r.\text{accept}(\text{this}); \\ \{x.r \mapsto j * \text{this}.\text{Visited}(j, \tau_2, \mu[\tau_1 \oplus \bullet])\} \end{array} \right\} \subseteq \end{array} \right\} \left. \begin{array}{l} \text{L-FRAME} \\ (11) \end{array} \right\}$$

(11) is verified in the same way. We can combine these two proofs to prove 7, provided (9) \Rightarrow (10) holds, which follows from using CLOSE with *Visitor* and OPEN with *Visited*, and $lcalc(\mu) + calc(\tau) = lcalc(\mu[\tau \oplus \bullet])$.

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. Parkinson's thesis [15] presents another subclass that removes all additions of zero from the term. As the proof is modular, we only need to verify the new classes we write: we know all the other classes interact correctly.

6 Related and future work

In this paper, we have built on our previous work [16]. Our examples have demonstrated the flexibility of separation logic and abstract predicate families. Although we have enforced behavioural subtyping, this has not placed undue constraints on subtypes. Additionally, we have shown that this methodology does not have difficulties with call-backs, because we do not need class invariants.

One drawback of this methodology is that it requires the rechecking of every inherited method. This is because the predicate definitions can be changed completely in subclasses. We are currently investigating ways of structuring abstract predicate families to allow the inheritance of methods without rechecking their specifications.

Our original inspiration for abstract predicate families came from Leino’s work on the extended state problem [11]. He proposed *data groups* to abstract modifies clauses, the set of fields modified by each method. The clauses were given in terms of abstract groups of fields. Each class was free to add its new fields to the abstract groups. Abstract predicate families generalise data groups by allowing constraints to be placed on these fields.

Closely related to our work is Boogie [4], which takes a different approach to the class invariant and call-backs problem. Boogie uses an auxiliary field, *st*, to indicate if the invariant holds (or even how much of the invariant holds). The invariant can be seen as $(this.st = Valid) \Rightarrow P(this)$, where P is the property that should hold of the object. `pack` and `unpack` operations are used to validate and invalidate the *st* field, respectively. These operations roughly correspond to the OPEN and CLOSE axioms in our logic. Boogie also represents similar encapsulation to abstract predicate families using an auxiliary *owner* field. Owned objects can only be unpacked, if its owner is unpacked. For comparison, consider an example from [4]:

```

class T {
  rep U f;
  invariant 0 ≤ f.g;
  void m() requires st=Valid
  { unpack this;
    f.n();...
    pack this; }
}

class U {
  int g;
  void n() requires st=Valid
  { ... }
}

```

In this example `T` owns its `f` field. We could encode this in our methodology using a single predicate family

$$\begin{aligned}
BoogieInv_T(x;) &\stackrel{\text{def}}{=} \exists i, g \cdot x.f \mapsto i * i.BoogieInv(g) \wedge g \geq 0 \\
BoogieInv_U(x; g) &\stackrel{\text{def}}{=} x.g \mapsto g
\end{aligned}$$

and making the pre- and post-conditions of each method `this.BoogieInv`. In the Boogie proof, the `unpack` operation in `T.m()` exposes the object in the `f` field. In our proof, we simply use OPEN on the predicate family to replace `this.BoogieInv()` with $\exists i, g \cdot x.f \mapsto i * i.BoogieInv(g) \wedge g \geq 0$. This allows us to meet the pre-condition

of the call to `f.n()`. In the Boogie proof, the `pack` operation in `T.m()` hides the object in the `f` field, and must also check the invariant holds. Similarly in our proof, we use `CLOSE`, which requires the predicate family for the field, and the invariant is true. There are many similarities between the two proof systems. In the future, we plan to explore the correspondence between the two systems in more detail.

Another closely related work is `Typestates` by Fähndrich and DeLine [8], which allows specification to be given with respect to named externally visible states. These states correspond to abstracted predicates in a similar way to that used here. They impose a strong structure on the predicates, which allows them to inherit methods without re-verifying code.

Finally, in this paper we have only considered sequential Java. Recently, O’Hearn has shown how to extend separation logic with rules to reason about concurrency primitives [13]. They allow state to be stored in a semaphore, and by manipulating this semaphore the state can be transferred between threads. Unfortunately the semaphore is statically scoped, which prevents reasoning about heap allocated semaphores including, for example, Java’s `synchronised` primitive. We are currently investigating how to extend O’Hearn’s system for concurrency to allow for reasoning about semaphores in the heap, and hence Java with threads.

Acknowledgments I should like to thank Gavin Bierman, Sophia Drossopoulou, Alisdair Wren and the anonymous referees for their comments on this work. This work was supported by an EPSRC DTA at the University of Cambridge, EPSRC grant EP/C523997/1, and Intel Research Cambridge.

Bibliography

- [1] M. Abadi and L. Cardelli. *A theory of objects*. Springer-Verlag, 1996.
- [2] P. America. Designing an object-oriented programming language with behavioural subtyping. In *Proceedings of the REX School/Workshop on Foundations of Object-Oriented Languages*, pages 60–90, 1991.
- [3] K. R. Apt. Ten years of Hoare’s logic: A survey: Part I. *ACM TOPLAS*, 3(4):431–483, 1981.
- [4] M. Barnett, R. DeLine, M. Fähndrich, K. R. M. Leino, and W. Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6):27–56, 2004.
- [5] G. M. Bierman and M. J. Parkinson. Effects and effect inference for a core Java calculus. In *Proceedings of WOOD*, volume 82 of *ENTCS*, 2004.
- [6] R. Bornat. Variables as resources in separation logic. In *Proceedings of MFPS*, pages 125–146, 2005.
- [7] K. K. Dhara and G. T. Leavens. Forcing behavioral subtyping through specification inheritance. In *Proceedings of ICSE*, pages 258–267, 1996.
- [8] M. Fähndrich and R. DeLine. Typestates for objects. In *Proceedings of ECOOP*, pages 465–490, 2004.
- [9] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, 1994.
- [10] S. Ishtiaq and P. W. O’Hearn. BI as an assertion language for mutable data structures. In *Proceedings of POPL*, pages 14–26, 2001.
- [11] K. R. M. Leino. Data groups: Specifying the modification of extended state. In *Proceedings of OOPSLA*, pages 144–153, 1998.
- [12] B. H. Liskov and J. M. Wing. A behavioral notion of subtyping. *ACM TOPLAS*, 16(6): 1811–1841, 1994.
- [13] P. W. O’Hearn. Resources, concurrency and local reasoning. In *Proceeding of CONCUR*, volume 3170, pages 49–67, 2004.
- [14] P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Proceedings of CSL*, pages 1–19, 2001.
- [15] M. J. Parkinson. *Local Reasoning for Java*. PhD thesis, Computer Laboratory, University of Cambridge, 2005. UCAM-CL-TR-654.
- [16] M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *Proceedings of POPL*, pages 247–258, 2005.
- [17] M. J. Parkinson, R. Bornat, and C. Calcagno. Variables as resource in Hoare logic. In *Proceedings of LICS*, 2006.
- [18] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings of LICS*, pages 55–74, 2002.
- [19] D. von Oheimb. Hoare logic for mutual recursion and local variables. In *Proceedings of FSTTCS*, pages 168–180, 1999.