ROPAS

Research On Program Analysis System
National Creative Research Initiative Center
Korea Advanced Institute of Science and Technology

ROPAS MEMO
2001-14
November 13, 2001

A Semantic Basis for Local Reasoning*

Hongseok Yang
ROPAS, KAIST

Peter W. O'Hearn
Queen Mary, University of London

November 13, 2001

Abstract

We present a semantic analysis of a recently proposed formalism for local reasoning, where a specification (and hence proof) can concentrate on only those cells that a program accesses. Our main results are the soundness and, in a sense, completeness of a rule that allows frame axioms, which describe invariant properties of portions of heap memory, to be inferred automatically, and thus avoided when writing specifications.

1 Introduction

The need to say what memory cells or other resources are not changed, along with those that are, has always been a vexing problem in program specification.

Consider a specification of a program to copy a tree:

$$\{\mathtt{tree}\,\tau\,p\}\,\mathtt{CopyTree}(p;q)\,\{(\mathtt{tree}\,\tau\,p)*(\mathtt{tree}\,\tau\,q)\}.$$

Here, the parameter list indicates that p is a value and q a reference parameter. The predicate $\mathtt{tree}\,\tau\,p$ says that p is, or points to, a data structure representing a binary tree τ, and (anticipating the work to come) $(\mathtt{tree}\,\tau\,p)*(\mathtt{tree}\,\tau\,q)$ says that p and q both represent this tree, but that their representations utilize disjoint storage.

This specification certainly captures part of what we intend to say about the procedure. But a Hoare triple typically only describes the effects an action has on the portion of program store it explicitly mentions; it does not say what cells among those not mentioned remain unchanged. As a result, the specification of $\mathtt{CopyTree}(p;q)$ leaves open the possibility that the procedure alters a cell not in the data structure described in the precondition. To make a stronger specification we need to, in one way or another, take into account the notorious "frame axioms", which describe cells that remain unchanged.

It might seem that this problem is just a nuisance, that we should be content for practical purposes to prove weak properties and not worry about frame axioms.

*Yang was supported by Creative Research Initiatives of the Korean Ministry of Science and Technology and by the US NSF under grant INT-9813854. O'Hearn was supported by the EPSRC under the "Local Reasoning about State" project.

This viewpoint is untenable, for the following reason. At a call site for `CopyTree`$(p; q)$ there will often be more cells active than those in p's data structure. In that case the specification is not strong enough to use, unless we somehow take the frame axioms into account. A particular example of such a call site is in the body of a recursive definition of `CopyTree`$(p; q)$, which uses recursive calls for each of two subtrees. As explained in [7], if we do not have some way of showing that each recursive call doesn't affect the other, then the specification will not be strong enough to use as an induction hypothesis when proving the program.

It might alternatively be thought that the problem can be easily solved, simply by listing the variables that a program might alter as part of the specification. This viewpoint is untenable when there are storage cells other than those directly named by variables, typically when there are pointers of one form or another. This solution is thus not applicable to realistic imperative languages.

Nonetheless, although much more than a nuisance, this frame problem is still irritating. It seems unfortunate to have to think up a general formula to function as a description of the cells left unchanged, covering all potential call sites, whenever we write a specification. And intuitively the specification of `CopyTree`$(p; q)$, for instance, seems to already carry enough information. The hitch is that for this intuition to be realized we need to somehow require that any state alteration not explicitly mandated by the specification is excluded.

Unfortunately, this last part, that "any state alteration not explicitly mandated by the specification is excluded" is difficult to make precise, and that difficulty has spawned hundreds of papers in AI and in program specification. Twice unfortunately, no completely convincing solution has emerged.

In this paper we study the semantics of an approach recently developed in an extension of Hoare's logic for reasoning about mutable data structures [7]. The scope of the approach is modest in intent, in that it is not put forward as a general solution to the frame problem. Rather, the suggestion is that, when certain assumptions are met, there is a natural and simple way to avoid frame axioms. Although, as we further demonstrate here, these assumptions are met in some natural models of imperative languages, there is no claim that they are universally applicable in reasoning about action.

The general idea is that by focusing on the idea of the memory footprint of a program we can get a concrete handle on the resources that a specification of a program needs to describe. More specifically, there are two components to the approach.

1. We interpret a specification $\{P\}\, C\, \{Q\}$ so that, when C is run in a state satisfying P, it must dereference only those cells guaranteed to exist by P or allocated during execution.

2. An inference rule, the Frame Rule, lets us obtain $\{P * R\}\, C\, \{Q * R\}$ from the initial specification $\{P\}\, C\, \{Q\}$ of a procedure or command, where $P * R$ is true just when P and R are true of separate areas of the current heap memory.

Point 1 is reminiscent of the old informal idea of a tight interpretation of specifications: we assume that a specification mentions all the resources relevant to understanding a program, that other resources are automatically unaffected. With it the specification

of the `CopyTree` procedure above implies that any active cell not in p's tree before execution will remain unchanged; this allows us to avoid explicit frame axioms and instead work with the simple specification. Point 2 then gives us a proof method that enables us to infer the invariant properties supported by tightness.

We stress that the first of these points does not depend on the second. Given the tight interpretation, a host of invariant properties are simply true, and this is independent of the language used to describe pre and postconditions. The $*$ connective just gives us a direct way to exploit tightness, in a program logic.

The purpose of this paper is to provide a semantic analysis of these ideas. The basic conception of the interplay between points 1 and 2 above as a basis for local reasoning is due to the second author. Many of the semantic foundations, particularly those related to completeness described later, were first worked out thoroughly by the first author. We refer to the precursor papers [9, 3, 7] for examples of reasoning with the spatial formalism, for program logic axioms for specific heap-altering and accessing commands, and for references to the literature on program logic and on the frame problem.

The first problem we tackle is the soundness of the Frame Rule. This turns out to be surprisingly delicate, and it is not difficult to find situations where the rule doesn't work. So a careful treatment of soundness, appealing to the semantics of a specific language, is essential. We phrase our argument here in terms of a model devised by Reynolds [9, 7], for an extension of the language of `while` programs with operations for pointer manipulation, including address arithmetic.

In the course of proving soundness we will attempt to isolate the properties on which it relies. Once these properties are established, much of the work in this paper can be carried out at a more abstract level and we sketch the appropriate definitions. But a detailed abstract account will be left for other work; here our goal is to remain concrete and provide a thorough analysis of a single language.

After showing soundness we move on to prove a completeness result, which shows a sense in which no frame axioms are missing. Our approach to completeness follows a line of work which goes under the name of sharpness or adaptation completeness [6]. There one of the main issues is always to conclude certain invariant properties of variables not free in a command. Our completeness result can be seen as extending such results to handle the heap as well, where there are other store locations than those associated with program variables. The key idea for proving completeness is to view commands as predicate transformers satisfying a locality property.

2 A Programming Language

The programming language is an extension of the language of `while` programs, with operations for manipulating pointers in a heap. The syntax and domains for the language are in Table 1.

The model has two components, the store and the heap. The store is a mapping from variables to integers, and the heap is a finite mapping from natural numbers (addresses) to integers. The heap is accessed using indirect addressing $[E]$ where E is an arithmetic expression.

Syntax

$$
\begin{array}{rcl}
C & ::= & x := E \mid x := [E] \mid [E] := E \mid x := \mathtt{cons}(E_1, ..., E_n) \mid \mathtt{dispose}(E) \\
 & \mid & C; C \mid \mathtt{while}\ B\ C \mid \mathtt{if}\ B\ \mathtt{then}\ C\ \mathtt{else}\ C \\
E & ::= & x, y, ... \mid 0 \mid 1 \mid E + E \mid E \times E \mid E - E \\
B & ::= & \mathtt{false} \mid B \Rightarrow B \mid E = F \mid E < F
\end{array}
$$

Domains

$$
\begin{array}{rcl@{\qquad}rcl}
\mathtt{Nats} & \triangleq & \{0, 1, ..., 17, ...\} & \mathtt{Ints} & \triangleq & \{..., -17, ..., -1, 0, 1, ..., 17, ...\} \\
\mathtt{Variables} & \triangleq & \{x, y, ...\} & \mathtt{Stores} & \triangleq & \mathtt{Variables} \to \mathtt{Ints} \\
\mathtt{Heaps} & \triangleq & \mathtt{Nats} \rightharpoonup_{fin} \mathtt{Ints} & \mathtt{States} & \triangleq & \mathtt{Stores} \times \mathtt{Heaps}
\end{array}
$$

Functionality of Expressions

$$[\![E]\!]s \in \mathtt{Ints} \qquad [\![B]\!]s \in \{true, false\}.$$

Table 1: Syntax and Domains

We assume the standard denotational semantics of integer and boolean expressions. Note that expressions are heap-independent.

The crucial operation on heaps is disjoint combination. We write $h \# h'$ to indicate that the domains $dom(h)$ and $dom(h')$ are disjoint. When $h \# h'$ holds, $h * h'$ is the heap obtained by taking the union of disjoint partial functions. When $h \# h'$ does not hold, $h * h'$ is undefined. The empty heap [] is the unit of $*$. The notation $[n \mapsto m]$ describes the singleton heap which maps n to m and which is undefined everywhere else.

The operational semantics of commands defines a relation $\rightsquigarrow$ on configurations. Configurations include terminal configurations s, h, triples C, s, h, and a special configuration *fault* indicating a memory fault. The command $x := [E]$ reads the value at address E in the heap and places it in x. $[E] := F$ updates address E so that its contents is F. $x := \mathtt{cons}(E_1, ..., E_n)$ allocates a sequence of n contiguous heap cells, initializes them to $E_1, ..., E_n$, and places the address of the first cell in the segment in x. $\mathtt{dispose}(E)$ removes address E from the heap. The commands $x := [E]$, $[E] := F$ and $\mathtt{dispose}(E)$ generate a memory fault, a particular kind of error, if E is not an active address. Notice that the number m is chosen non-deterministically in the rule for `cons`. An example of a command that always faults is $\mathtt{dispose}(x)$; $[x] := 42$. In typical implementations an attempt to dereference a disposed address might not always lead immediately to a fault. Generating these faults early in the semantics is a device

$$\frac{}{x := E, s, h \rightsquigarrow (s \mid x \mapsto [\![E]\!]s), h}$$

$$\frac{[\![E]\!]s = n}{\texttt{dispose}(E), s, h * [n \mapsto m] \rightsquigarrow s, h} \qquad \frac{[\![E]\!]s = n \quad n \notin dom(h)}{\texttt{dispose}(E), s, h \rightsquigarrow fault}$$

$$\frac{[\![E]\!]s = n \in dom(h) \quad h(n) = m}{x := [E], s, h \rightsquigarrow (s \mid x \mapsto m), h} \qquad \frac{[\![E]\!]s \notin dom(h)}{x := [E], s, h \rightsquigarrow fault}$$

$$\frac{[\![E]\!]s = n \in dom(h)}{[E] := F, s, h \rightsquigarrow s, (h \mid n \mapsto [\![F]\!]s)} \qquad \frac{[\![E]\!]s \notin dom(h)}{[E] := F, s, h \rightsquigarrow fault}$$

$$\frac{m, ..., n + m - 1 \notin dom(h) \quad v_1 = [\![E_1]\!]s, ..., v_n = [\![E_n]\!]s}{x := \texttt{cons}(E_1, ..., E_n), s, h \rightsquigarrow (s \mid x \mapsto m), (h * [m \mapsto v_1, ..., n + m - 1 \mapsto v_n])}$$

$$\frac{C_1, s, h \rightsquigarrow C_1', s', h'}{(C_1; C_2), s, h \rightsquigarrow (C_1'; C_2), s', h'} \qquad \frac{C_1, s, h \rightsquigarrow s', h'}{(C_1; C_2), s, h \rightsquigarrow C_2, s', h'}$$

$$\frac{C_1, s, h \rightsquigarrow fault}{(C_1; C_2), s, h \rightsquigarrow fault}$$

$$\frac{[\![B]\!]s = true}{\texttt{if } B \texttt{ then } C \texttt{ else } C', s, h \rightsquigarrow C, s, h} \qquad \frac{[\![B]\!]s = false}{\texttt{if } B \texttt{ then } C \texttt{ else } C', s, h \rightsquigarrow C', s, h}$$

$$\frac{[\![B]\!]s = false}{\texttt{while } B \texttt{ do } C \texttt{ od}, s, h \rightsquigarrow s, h} \qquad \frac{[\![B]\!]s = true \quad (C; \texttt{while } B \texttt{ do } C \texttt{ od}), s, h \rightsquigarrow K}{\texttt{while } B \texttt{ do } C \texttt{ od}, s, h \rightsquigarrow K}$$

Table 2: The Programming Language: Syntax and Semantics

that allows us to arrange the formalism in a conservative manner, where well-specified programs will never try to dereference a disposed address.

In the semantics we use $(f \mid i \mapsto j)$ for the (perhaps partial) function like f except that i goes to j. This notation is used both when i is and is not in the domain of f.

3 Specifications

We treat predicates semantically in this paper, so a predicate is just a subset of the set of states.

$$\mathsf{Pred} \triangleq \mathcal{P}(\mathsf{States})$$

To evoke the semantics of the assertion languages from [9, 3, 7], we sometimes use the satisfaction notation

$$s, h \models p$$

as an alternative to $(s, h) \in p$.

To define the semantics of Hoare triples first recall point 1 from the Introduction, where we guarantee that if $\{p\}C\{q\}$ holds then C must not access any cells not guaranteed to exist by p. We can formalize this by observing that if C did guarantee to

access such a cell, then it could be made to fault by running it in a state in which that cell is not active. This is the role of the following notion of safety.

- "C, s, h is safe" when $C, s, h \not\rightsquigarrow^* \mathit{fault}$.

For partial correctness the interpretation of triples is the standard one, with an additional safety requirement.

Partial Correctness. $\{p\}C\{q\}$ is true just when, for all s, h,

if $s, h \models p$ then

- C, s, h is safe, and
- if $C, s, h \rightsquigarrow^* s', h'$ then $s', h' \models q$.

This fault-avoiding interpretation of triples is not new, but the connection to the intuitive notion of tight specification does not seem to have been observed before (excepting [3, 7]). To describe this, suppose $E \hookrightarrow F$ is a predicate saying that F is the contents of the active address E in the current heap. Suppose, further, that the triple

$$\{x \hookrightarrow 5\}C\{x \hookrightarrow 6\}$$

holds and that C does not alter any variables (though it might alter heap cells). Then we claim that, just from this information, we can infer that C does not modify any heap cell existing in the starting state, other than the one denoted by x. To see the reason, suppose that C does modify such a cell and call it y. Then the specification says that the program will not fault if it is run in the singleton state where x is the only active address (with contents 5). But we just said that C alters the address y, and so this attempt to dereference y must generate a fault starting from the singleton state.

We will also consider a total correctness form of specification. For total correctness, we do not need an explicit safety assumption, because total correctness is about "must termination" which itself includes a safety requirement.

- "C, s, h must terminate normally" when C, s, h is safe and there is no infinite $\rightsquigarrow$-sequence starting from C, s, h.

Total Correctness. $\{p\}C\{q\}$ is true just when, for all s, h

if $s, h \models p$ then

- C, s, h must terminate normally, and
- if $C, s, h \rightsquigarrow^* s', h'$ then $s', h' \models q$.

To formulate the Frame Rule we will need the $*$ connective; if p and q are predicates then

$$p * q \quad \triangleq \quad \{(s, h * h') \mid (s, h) \in p \wedge (s, h') \in q \wedge h \# h'\}$$

The statement of the Frame Rule from [3, 7] is as follows.

FRAME RULE, SYNTACTIC VERSION

$$\frac{\{P\}\,C\,\{Q\}}{\{P * R\}\,C\,\{Q * R\}}\ \text{Modifies}(C) \cap \text{Free}(R) = \{\ \}$$

where Modifies(C) denotes the set of variables updated in the command C, i.e., those appearing as the destination of an assignment statement in C. To be precise, the Modifies set of each of $x := \cdots$ is $\{x\}$, while for $[E] := F$ and $\texttt{dispose}(E)$ it is empty; these latter two statements affect the heap but not the values of variables in the store. Since we are working semantically with predicates in this paper, we need to reformulate the rule slightly and replace the reference to the free variables of R in the side condition. We will precede this with a short discussion.

The condition for variables in the rule is straightforward; it simply checks whether any variables in R are modified by a command C. However, the condition for heap cells is more elaborate. The solution in the rule is to take an indirect route via the precondition P. With the spatial conjunction, it says that for every state satisfying $P * R$, the current heap can be split into two subheaps so that P holds for the one and R for the other; then, the tight interpretation of the Hoare triple $\{P\}\,C\,\{Q\}$ says that the command can only access the first part, i.e. the part for P, consequently making R an invariant during execution.

As an example, starting from the specification of `CopyTree` in the Introduction we can infer that copying p's tree does not affect a cell not in it's data structure.

$$\frac{\{\texttt{tree}\,\tau\,p\}\,\texttt{CopyTree}(p;q)\,\{(\texttt{tree}\,\tau\,p) * (\texttt{tree}\,\tau\,q)\}}{\{(\texttt{tree}\,\tau\,p) * (x \mapsto y)\}\,\texttt{CopyTree}(p;q)\,\{(\texttt{tree}\,\tau\,p) * (\texttt{tree}\,\tau\,q) * (x \mapsto y)\}}$$

Here, the Modifies set of $\texttt{CopyTree}(p;q)$ is assumed to be $\{q\}$.

To describe a version of the rule which refers to semantic rather than syntactic predicates we utilize a notion $X\#p$ of independence of a predicate from a set of variables. This can be formulated simply in terms of quantification. If X is a set of variables then

$$\forall X.p \ \stackrel{\Delta}{=}\ \{(s,h) \mid \forall s_X \in [X \to \texttt{Ints}].\,(s[s_X], h) \in p\}$$

where $s[s']$ denotes the update of s by s' defined by:

$$s[s'](y) \ \stackrel{\Delta}{=}\ \begin{cases} s'(y) & \text{if } y \in dom(s') \\ s(y) & \text{otherwise} \end{cases}$$

Then

- $X\#p$ holds just if $p = \forall X.p$.

FRAME RULE, SEMANTIC VERSION

$$\frac{\{p\}\,C\,\{q\}}{\{p * r\}\,C\,\{q * r\}}\ \text{Modifies}(C)\#r$$

4 Soundness of the Frame Rule

The Frame Rule codifies a notion of local behaviour, and in this section we undertake to describe that notion in terms of the operational semantics.

It will be helpful to first consider a plausible property that does not hold.

If $C, s, h_0 \rightsquigarrow^* s', h_0'$ and $h_0 \# h_1$, then $C, s, h_0 * h_1 \rightsquigarrow^* s', h_0' * h_1$.

The intuition behind this property is just that we can add on extra state, and any execution that works for a smaller state can still go ahead. This property fails because of the behaviour of `cons`. An address that is allocated during an execution from a small state cannot be allocated starting in a bigger state where it is already active. For example,

$$x := \mathtt{cons}(2,3), [x \mapsto m], [\,] \rightsquigarrow [x \mapsto 0], [0 \mapsto 2, 1 \mapsto 3]$$

but if we run $x := \mathtt{cons}(2,3)$ in a heap where 0 is already active, then a different address than 0 must be chosen by $\rightsquigarrow$.

This is an example of where an action on a little state can be disabled when moving to bigger states. Such behaviour might make us doubt that the Frame Rule could be sound at all. However, it only indicates that the language does not behave locally with respect to "may" properties. With Hoare triples we are interested in a form of "must" property. If $\{p\}C\{q\}$ holds and, starting from a state satisfying p, the command terminates, then the final state must satisfy q. Put another way, it is not individual computations that we will judge local, but properties of classes of computations.

The correct property says that if a command is safe in a given state, then the result of executing it in a larger state can be tracked to *some* execution computation on the little state.

Lemma 1 (Safety and Termination Monotonicity) *1. If C, s, h is safe and $h \# h'$, then $C, s, h * h'$ is safe.*

2. If C, s, h must terminate normally and $h \# h'$, then $C, s, h * h'$ must terminate normally.

Lemma 2 (Frame Property) *Suppose*

C, s, h_0 *is safe, and* $C, s, h_0 * h_1 \rightsquigarrow^* s', h'$.

Then there is h_0' where

$C, s, h_0 \rightsquigarrow^* s', h_0'$, *and* $h' = h_0' * h_1$.

Proof: For $x := \mathtt{cons}(E_1, ..., E_n)$, consider $m, ..., n + m - 1 \notin dom(h_0 * h_1)$, The operational rule gives us $s' = (s \mid x \mapsto m)$ and $h' = h_0 * h_1 * [m \mapsto v_1, ..., m+n-1 \mapsto v_n]$. Then since $m, ..., m+n-1 \notin dom(h_0)$, this segment may be selected by the operational rule for `cons` applied in the smaller heap h_0. So $h_0' = h_0 * [m \mapsto v_1, ..., m+n-1 \mapsto v_n]$ gives us the desired result.

For $[E] := F$, since $[E] := F, s, h_0$ is safe by assumption we know $[\![E]\!]s \in dom(h_0)$ and therefore $[\![E]\!]s \notin dom(h_1)$. Thus, the assignment leaves h_1 unchanged, and taking $h_0' = (h_0 \mid [\![E]\!]s \mapsto [\![F]\!]s)$ gives the result.

For `dispose`(E), if $[\![E]\!]s = n$ then $n \in dom(h_0)$ by safety, and h_0 decomposes as $h'_0 * [n \mapsto m]$ for some m. This h'_0 satisfies the requirement of the theorem.

For $x := [E]$, safety ensures that $[\![E]\!]s \in dom(h_0)$, and taking $h'_0 = h_0$ gives the result.

To cover $C; C'$, `while` and `if`, we need to prove a slightly stronger result. We state the property, which is needed to get the right induction hypothesis, but omit the detailed proof. Consider a variant of the theorem which considers non-terminal configurations; specifically, where the terminal configurations s', h' and s', h'_0 are replaced by C', s', h' and C', s', h'_0. The terminal and non-terminal variants are proven simultaneously, by induction on the derivation of $C, s, h_0 * h_1 \rightsquigarrow^* s', h'$ or $C, s, h_0 * h_1 \rightsquigarrow^* C', s', h'$. The rules for atomic commands above were already considered above, and each of the rules for `while` and $C; C'$ has an immediate proof. □

Theorem 3 (Soundness) *The* Frame Rule *is sound for both partial and total correctness.*

Proof: The proof uses Frame Property, and a locality property for variables.

If $C, s, h \rightsquigarrow^* s', h'$, x is not assigned in C and $x \in dom(s)$, then $s(x) = s'(x)$.

For partial correctness, suppose the premise of the Frame Rule holds, and that $s, h_0 \models p$ and $s, h_1 \models r$. Then premise gives us that C, s, h_0 is safe, and safety of $C, s, h_0 * h_1$ follows from Safety Monotonicity. If $C, s, h_0 * h_1 \rightsquigarrow^* s', h$ then Frame Property yields h'_0 where $h = h'_0 * h_1$ and $C, s, h_0 \rightsquigarrow^* s', h'_0$. The premise then ensures $s', h'_0 \models q$. The variable locality property implies that s' agrees with s on all variables not in Modifies(C), so $s', h_1 \models r$ follows since we know $s, h_1 \models r$ and Modifies$(C)\#r$. The semantics of $*$ then yields $s, h \models q * r$.

The argument for total correctness appeals additionally to Termination Monotonicity. □

Some remarks on the delicacy of this soundness result are in order. First, the non-deterministic nature of `cons` is essential for the result. If we had interpreted it so that, say, the smallest possible free address was always chosen for allocation, then adding memory would change what this new address was, and the difference could be detected with address arithmetic; this would invalidate the Frame Rule. Non-deterministic allocation is used to force a program proof not to depend on details of how the allocator might work. (In a language without address arithmetic, we could use invariance under location renaming rather than non-determinism in allocation to ensure this sort of independence.)

Second, suppose we were to add an operation for trapping a memory fault to our language. If we did this, without changing $*$, then the Frame Rule would be invalid; the reason is that we could branch on whether or not a cell is active, and this would even contradict Safety Monotonicity. This does not necessarily mean that fault trapping is incompatible with the Frame Rule. We could perhaps change the interpretation of $*$ so that the undefinedness it introduces is regarded as introducing the possibility of a further kind of error, different from memory fault. (The way Calcagno puts it, we need a notion that is detectable in the program logic, but not in the programming language.)

CONSEQUENCE

$$\frac{p' \subseteq p \quad \{p\} - \{q\} \quad q \subseteq q'}{\{p'\} - \{q'\}}$$

FRAME RULE

$$\frac{\{p\} - \{q\}}{\{p * r\} - \{q * r\}} \; X \# r$$

Table 3: Proof system for specification statements with Modifies set X

These examples indicate that the Frame Rule is not something we expect to be automatically valid, in the way that, say, we expect the rule of Consequence always to be. Close attention is needed to the interplay between the definition of $*$ and the kinds of operation present in a language.

5 Completeness of the Frame Rule

Suppose we are given a Hoare triple specification $\{p\}C\{q\}$ but we are not told exactly what C is. The question we are concerned with in this section is whether we can derive all other specifications that follow from it, without making use of knowledge of C.

To formulate this we consider Hoare triples $\{p\} - \{q\}$ with an unspecified command. Following [5, 11, 1], we call such a Hoare triple with a hole a *specification statement*. Throughout this section we assume a given set X of variables, regarded as the Modifies set. To derive one specification statement from the other we use the usual rule of Consequence from Hoare logic, and the Frame Rule; see Table 3.

The completeness question, which is variously called sharpness or adaptation completeness, is whether this system lets us derive one specification statement from another just when this inference holds semantically. Since the rule of Consequence is itself treated semantically, in that it uses inclusion between predicates-as-sets rather than a provable implication, all of the stress in this question is placed on the Frame Rule: it is essentially asking whether we obtain enough frame axioms.

To show completeness we need to define a notion of semantic consequence between specification statements. We use predicate transformers to interpret specification statements because, as often noticed from the work on refinement calculus, predicate transformers provide a good setting where the essence of a proof rule for Hoare triples can be expressed directly.

So the plan in this section is to define a class of "local" predicate transformers, which satisfy a condition that ensures soundness of the Frame Rule, then to show the connection between this condition and the locality properties of the operational semantics discussed in Section 4, and finally to prove the completeness result. We refer the reader to Yang's thesis [10] for further results on when a local predicate transformer corresponds to a relation on states, and on when the completeness result can be formulated in terms of relations rather than transformers.

5.1 Local Predicate Transformers

In a predicate transformer interpretation, a command C is interpreted as a mapping from a postcondition to a precondition. For instance, the predicate transformer induced by $x := 2$ maps $y = x$ to $y = 2$.

Mathematically, predicate transformers are monotone maps from predicates to predicates.

$$\mathsf{PT} \triangleq \mathsf{Pred} \rightarrow_{monotone} \mathsf{Pred}$$

Given a predicate transformer t the Hoare triple "$\{p\}\, t\, \{q\}$" corresponds to the property $p \subseteq t(q)$. Then, the monotonicity requirement is equivalent to saying that the rule of Consequence is valid.

The domain PT contains predicate transformers which, when viewed operationally, don't exhibit the local behavior of commands as described in Section 4. For instance, the predicate transformer $\lambda q.\, \{(s,h) \in q \mid |dom(h)| \leq 3\}$ corresponds to a command which generates a memory fault when there are more than 3 active heap cells and skips otherwise; this command doesn't satisfy Safety Monotonicity. So, it is not surprising that the Frame Rule is not valid for PT since the rule is closely related to the locality properties of Section 4.

We refine PT by requiring predicate transformers to satisfy a locality condition. Recall that in this section we are assuming a given set X of modifiable variables.

$$\text{Locality for } X\text{: } \forall r, q \in \mathsf{Pred}.\ X\#r \Longrightarrow t(q) * r \subseteq t(q * r)$$

The condition is equivalent to saying that t satisfies the Frame Rule, which in predicate transformer terms is

$$\forall p, r, q \in \mathsf{Pred}.\ X\#r \wedge p \subseteq t(q) \Longrightarrow p * r \subseteq t(q * r)$$

The domain $\mathsf{LPT}(X)$ of *local predicate transformers with Modifies set* X is defined as:

$$\mathsf{LPT}(X) \triangleq \{t \in \mathsf{PT} \mid t \text{ satisfies locality for } X\}.$$

$\mathsf{LPT}(X)$ inherits the ordering from PT, which is just a pointwise ordering induced by the subset ordering of Pred. With such an ordering, $\mathsf{LPT}(X)$ forms a complete lattice, in fact, a complete sublattice of PT.

Proposition 4 *$\mathsf{LPT}(X)$ is a complete sublattice of PT. Therefore, given a set $\{t_i\}_{i\in I}$ of elements in $\mathsf{LPT}(X)$, its least upper bound is given by $\lambda q.\ \bigcup_{i\in I} t_i(q)$, and its greatest lower bound is given by $\lambda q.\ \bigcap_{i\in I} t_i(q)$.*

Proof: Let r be a predicate such that $X\#r$. It suffices to show that $\bigcup_{i\in I} t_i(q * r)$ and $\bigcap_{i\in I} t_i(q * r)$ include $(\bigcup_{i\in I} t_i(q)) * r$ and $(\bigcap_{i\in I} t_i(q)) * r$, respectively. The first inclusion follows because $\bigcup_{i\in I}(t_i(q) * r) = (\bigcup_{i\in I} t_i(q_i)) * r$, and the second inclusion holds since $*$ is monotone with respect to the subset ordering. □

5.2 Operational Sensibility

The locality condition in $\mathsf{LPT}(X)$ has an operational explanation via mappings, *wp* and *wlp*, from relations to predicate transformers, which correspond to total correctness and partial correctness. Let R be a relation of type $\mathsf{States} \leftrightarrow (\mathsf{States} \cup \{\bot, \mathit{fault}\})$ and q a predicate. Then we define the weakest, and weakest liberal, precondition predicate transformers as follows.[1]

$$\begin{array}{rcl} wp(R)(q) & \triangleq & \{(s,h) \mid (R,s,h) \text{ must terminate normally} \\ & & \quad \wedge \forall (s',h').\, (s,h)\, R\, (s',h') \Longrightarrow (s',h') \in q\} \\ wlp(R)(q) & \triangleq & \{(s,h) \mid (R,(s,h)) \text{ is safe} \\ & & \quad \wedge \forall (s',h').\, (s,h)\, R\, (s',h') \Longrightarrow (s',h') \in q\} \end{array}$$

We can now establish the operational sensibility of local predicate transformers as follows. First, given a command C we can obtain a relation

$$[\![C]\!] : \mathsf{States} \leftrightarrow (\mathsf{States} \cup \{\bot, \mathit{fault}\})$$

where

- $(s,h)[\![C]\!](s',h')$ iff $C,s,h \rightsquigarrow^* (s,h)$;
- $(s,h)[\![C]\!]\mathit{fault}$ iff $C,s,h \rightsquigarrow^* \mathit{fault}$;
- $(s,h)[\![C]\!]\bot$ iff there is an infinite $\rightsquigarrow$ sequence starting from C,s,h.

Then, if $X \supseteq \mathrm{Modifies}(C)$, we obtain that both $wp[\![C]\!]$ and $wlp[\![C]\!]$ are local predicate transformers in $\mathsf{LPT}(X)$.

Proposition 5 *For a relation R: $\mathsf{States} \leftrightarrow \mathsf{States} \cup \{\bot, \mathit{fault}\}$, both $wp(R)$ and $wlp(R)$ satisfy locality for X iff the following four conditions hold for R:*

1. *Variable Locality: if R,s,h is safe and $(s,h)\, R\, (s',h')$, then $s(y) = s'(y)$ for all $y \in \mathsf{Variables} - X$;*
2. *Safety Monotonicity: if R,s,h is safe and $h\#h'$, then $R,s,h*h'$ is safe;*
3. *Termination Monotonicity: if R,s,h must terminate normally and $h\#h'$, then $R,s,h*h'$ must terminate normally; and*
4. *Frame Property: if R,s,h_1 is safe, $h_1\#h_2$ and $(s,h_1*h_2)\, R\, (s',h')$, then there exists h_1' such that $(s,h_1)\, R\, (s',h_1')$ and $h' = h_1'*h_2$.*

Notice that Safety Monotonicity, Termination Monotonicity and Frame Property in the lemma are exactly those in Section 4, and Variable Locality was essentially used in the proof of Theorem 3. This proposition establishes as close a correspondence as we would expect with the locality properties of Section 4; since both *wp* and *wlp* completely ignore all unsafe configurations we can not obtain any properties of unsafe configurations with predicate transformers.

[1] We abuse the terminologies for commands here: for a relation R, we say that R,s,h is safe just if not $((s,h)\, R\, \mathit{fault})$, and that R,s,h must terminate normally iff it is safe and not $((s,h)\, R\, \bot)$.

Proof: For if direction, let r, q be predicates such that $X\#r$. Pick $(s,h) \in wp(R)(q)*r$. Then, there exist h_1, h_2 such that $h_1\#h_2$, $h = h_1*h_2$, $(s,h_1) \in wp(R)(q)$ and $(s,h_2) \in r$. By Termination Monotonicity, $R, s, h_1 * h_2$ must terminate normally. For all $(s, h_1 * h_2)\, R\, (s', h')$, Variable Locality says that $s' = s[s_X]$ for some $s_X \in [X \to \mathsf{Ints}]$, and Frame Property implies that $h' = h'_1 * h_2$ for some h'_1 with $(s,h_1)\, R\, (s', h'_1)$; therefore, $(s', h') \in q*r$ because $(s[s_X], h_2) \in r$ and $(s', h'_1) \in q$. The case for wlp can be handled similarly; the only difference is to use Safety Monotonicity instead of Termination Monotonicity.

For the other direction, we handle each item separately:

1. *Variable Locality*: suppose that R, s, h is safe and $(s,h)\, R\, (s', h')$. Let $r = \{(s[s_X], []) \mid s_X \in [X \to \mathsf{Ints}]\}$ and let q be $\{(s'', h'') \mid (s,h)\, R\, (s'', h'')\}$. Since $X\#r$, we have that $wlp(R)(q*r) \supseteq wlp(R)(q)*r$. Note that $(s,h) \in wlp(R)(q)*r$. Therefore, it also belongs to $wlp(R)(q*r)$, so $(s', h') \in q*r$. The conclusion follows since $(s', []) \in r$.

2. *Safety Monotonicity*: suppose that R, s, h is safe and $h\#h'$. Let $r = \{(s[s_X], h') \mid s_X \in [X \to \mathsf{Ints}]\}$ and let q be $\{(s'', h'') \mid (s,h)\, R\, (s'', h'')\}$. Since $X\#r$, we have that $wlp(R)(q*r) \supseteq wlp(R)(q)*r$. Note that $(s, h*h') \in wlp(R)(q)*r$. Therefore, $(s, h*h') \in wlp(R)(q*r)$, so $R, s, h*h'$ is safe.

3. *Termination Monotonicity*: the proof is essentially the same as the case of Safety Monotonicity. It can be obtained just replacing "is safe" by "must terminate normally" and wlp by wp.

4. *Frame Property*: suppose that R, s, h_1 is safe, $h_1\#h_2$ and $(s, h_1 * h_2)\, R\, (s', h')$. Let $r = \{(s[s_X], h_2) \mid s_X \in [X \to \mathsf{Ints}]\}$ and let q be $\{(s'', h'') \mid (s,h_1)\, R\, (s'', h'')\}$. Since $X\#r$, we have that $wlp(R)(q*r) \supseteq wlp(R)(q)*r$. Since $(s, h_1*h_2) \in wlp(R)(q)*r$, the state (s, h_1*h_2) belongs to $wlp(R)(q*r)$. So, $(s', h') \in q*r$. Then, there are h'_1 and $s_X \in [X \to \mathsf{Ints}]$ such that $(s,h_1)\, R\, (s[s_X], h'_1)$, $s' = s[s_X]$ and $h' = h'_1 * h_2$.

□

5.3 Proof of Completeness

We have already given a proof system for specification statements in Table 3. This determines a notion of consequence between specification statements.

$$\{p\} - \{q\} \vdash_X \{p'\} - \{q'\} \quad \text{iff} \quad \{p'\} - \{q'\} \text{ can be derived from } \{p\} - \{q\} \text{ using the rules in Table 3.}$$

The semantic interpretation of each specification statement is given by a satisfaction relation between local predicate transformers in $\mathsf{LPT}(X)$ and specification statements. For t in $\mathsf{LPT}(X)$

$$t \models_X \{p\} - \{q\} \quad \text{iff} \quad p \subseteq t(q).$$

We can now define the semantic consequence relation $\models$ by requiring that any transformer satisfying the antecedent also satisfies the consequent.

$$\{p\} - \{q\} \models_X \{p'\} - \{q'\} \quad \text{iff} \quad \begin{array}{l}\text{for all } t \in \mathsf{LPT}(X),\\ \quad t \models_X \{p\} - \{q\} \text{ implies } t \models_X \{p'\} - \{q'\}.\end{array}$$

The proof system is sound with respect to this interpretation since the monotonicity condition in PT ensures that the rule of Consequence is sound and the locality condition for X guarantees that the Frame Rule is sound.

Proposition 6 (Soundness, II) *If* $\{p\} - \{q\} \vdash_X \{p'\} - \{q'\}$, *then* $\{p\} - \{q\} \models_X \{p'\} - \{q'\}$.

In the remainder of this section we concentrate on the proof of the converse.

Theorem 7 (Completeness) *If* $\{p\} - \{q\} \models_X \{p'\} - \{q'\}$, *then* $\{p\} - \{q\} \vdash_X \{p'\} - \{q'\}$.

The proof of completeness proceeds by finding the worst (smallest) local predicate transformer t in $\mathsf{LPT}(X)$, amongst those that satisfy a given specification. To describe this transformer it will be convenient to use a notation for the spatial implication $-\!\!*$ of predicates in pointer logic [3, 7], which is itself taken from the logic of Bunched Implications [8].

$$p \mathrel{-\!\!*} q \quad \stackrel{\Delta}{=} \quad \{(s,h) \mid \forall h'.\, h' \# h \wedge (s,h') \in p \Longrightarrow (s, h * h') \in q\}$$

In words, $p \mathrel{-\!\!*} q$ holds of a given heap if, whenever we are given new or fresh heap satisfying p, the combined new and current heap satisfies q.

Lemma 8 *Given predicates* $p, q \in \mathsf{Pred}$, *the predicate transformer* $t = \lambda r.\, p * \forall X.\, (q \mathrel{-\!\!*} r)$ *is in* $\mathsf{LPT}(X)$ *and satisfies* $p \subseteq t(q)$. *Moreover, it is the worst (smallest) such in* $\mathsf{LPT}(X)$ *with respect to the pointwise ordering.*

Proof: It is straightforward to see that t is monotone. To see that t satisfies the locality condition for X, pick predicate r, r' with $X \# r'$. Then, we have:

$$\begin{array}{lll} t(r * r') & = & p * (\forall X.\, (q \mathrel{-\!\!*} r * r')) \\ & \supseteq & p * (\forall X.\, (q \mathrel{-\!\!*} r)) * (\forall X.\, r') \end{array} \quad \begin{array}{lll} & \supseteq & p * (\forall X.\, (q \mathrel{-\!\!*} r)) * r' \\ & = & p * (\forall X.\, (q \mathrel{-\!\!*} r)) * r' \\ & = & t(r) * r' \end{array}$$

The condition $p \subseteq t(q)$ is also easily verified as follows:

$$\begin{array}{lll} t(q) & = & p * (\forall X.\, (q \mathrel{-\!\!*} q)) \\ & \supseteq & p * \mathsf{emp} \end{array} \quad \begin{array}{lll} & \supseteq & p * (\forall X.\, \mathsf{emp}) \\ & \supseteq & p \end{array}$$

where $\mathsf{emp} = \{(s, []) \mid s \in \mathtt{Stores}\}$ is the unit of $*$. Finally, t is the worst satisfying $p \subseteq t(q)$ because for any other $t' \in \mathsf{LPT}(X)$ with $p \subseteq t'(q)$,

$$t'(q * \forall X.\, (q \mathrel{-\!\!*} r)) \;\supseteq\; t'(q) * \forall X.\, (q \mathrel{-\!\!*} r) \;\supseteq\; p * \forall X.\, (q \mathrel{-\!\!*} r).$$

□

We denote the worst local predicate transformer in Lemma 8 by $worst(p, q, X)$.

To prove completeness, first note that for all predicates p, q, r, the precondition $worst(p, q, X)(r)$ can be calculated with the Frame Rule and the rule of Consequence:

$$\dfrac{\dfrac{\{p\} - \{q\}}{\{p * \forall X.(q \mathrel{-\!\!*} r)\} \ - \ \{q * \forall X.(q \mathrel{-\!\!*} r)\}} \qquad q * \forall X.(q \mathrel{-\!\!*} r) \subseteq r}{\{p * \forall X.(q \mathrel{-\!\!*} r)\} \ - \ \{r\}}$$

The right-hand inclusion is just an application of the usual rule of $\forall$-elimination, together with a monotonicity rule for $*$ and the version of modus ponens that connects $*$ and $\mathrel{-\!\!*}$.

Now, let $\{p'\} - \{q'\}$ be a specification statement with $\{p\} - \{q\} \models_X \{p'\} - \{q'\}$. Then, we have $p' \subseteq worst(p, q, X)(q')$ by Lemma 8, and $\{p'\} - \{q'\}$ can be derived from $\{p\} - \{q\}$ as follows:

$$\dfrac{p' \subseteq p * \forall X.(q \mathrel{-\!\!*} q') \qquad \dfrac{\{p\} - \{q\}}{\{p * \forall X.(q \mathrel{-\!\!*} q')\} \ - \ \{q'\}}}{\{p'\} \ - \ \{q'\}}$$

This finishes the proof of Theorem 7.

An interesting question is whether the most general invariant $\forall X.(q \mathrel{-\!\!*} r)$ in the proof of completeness could be utilized in automatic or semi-automatic program verification. When verifying programs with procedures a key step is often choosing an invariant for use with the Frame Rule [3, 7], and the above argument shows a theoretical sense in which the guesswork can be taken out of such steps. Of course, the price is that we are left with an implication $p' \subseteq p * \forall X.(q \mathrel{-\!\!*} q')$ that might be much more difficult to establish than if the invariant were chosen differently. So it is natural to ask under what circumstances this price might be controlled, in particular by taming the quantification in $\forall X.(q \mathrel{-\!\!*} r)$, and perhaps incorporated into the decidable sublogic of [2].

References

[1] R.-J. R. Back. *On the Correctness of Refinement Steps in Program Development.* PhD thesis, Department of Computer Science, University of Helsinki, 1978. Report A-1978-4.

[2] C. Calcagno, H. Yang, and P. O'Hearn. Computability and complexity results for a spatial assertion language for data structures. To appear in the Proceedings of FSTTCS, 2001.

[3] S. Ishtiaq and P. O'Hearn. BI as an assertion language for mutable data structures. In *Principles of Programming Languages*, pages 14–26, January 2001.

[4] C. Morgan and T. Vickers, editors. *On the Refinement Calculus.* Springer-Verlag, 1992.

[5] C. C. Morgan. The specification statement. *ACM Transactions on Programming Languages and Systems*, 10(3), Jul 1988. Reprinted in [4].

[6] D. Naumann. Calculating sharp adaptation rules. *Information Processing Letters*, 2000. To appear.

[7] P. O'Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *Proceedings of 15th Annual Conference of the European Association for Computer Science Logic: CSL 2001*, pages 1–19. Springer-Verlag. LNCS 2142.

[8] P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, June 99.

[9] J. C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In Jim Davies, Bill Roscoe, and Jim Woodcock, editors, *Millennial Perspectives in Computer Science*, pages 303–321, Houndsmill, Hampshire, 2000. Palgrave.

[10] H. Yang. *Local Reasoning for Stateful Programs.* Ph.D. thesis, University of Illinois, Urbana-Champaign, Illinois, USA, 2001.

[11] H. Yang and U. S. Reddy. On the semantics of refinement calculi. In *Foundations of Software Science and Computation Structures*, pages 359–374. Springer-Verlag, 2000.