

1 Axiomatic Semantics

So far we have focused on *operational semantics*, which are natural for modeling computation or talking about how state changes from one step of the computation to the next. In operational semantics, there is a well-defined notion of *state*. We take great pains to say exactly what a state is and how it is manipulated by a program.

In *axiomatic semantics*, on the other hand, we do not so much care what the states actually are, but only the properties that we can observe about them. This approach emphasizes the relationship between the properties of the input (preconditions) and properties of the output (postconditions). This approach is useful for specifying what a program is supposed to do and talk about a program's correctness with respect to that specification.

2 Preconditions and Postconditions

The *preconditions* and *postconditions* of a program say what is true before and after the program executes, respectively. Often the correctness of the program is specified in these terms. Typically this is expressed as a contract: as long as the caller guarantees that the initial state satisfies some set of preconditions, then the program will guarantee that the final state will satisfy some desired set of postconditions. Axiomatic semantics attempts to say exactly what preconditions are necessary for ensuring a given set of postconditions.

3 An Example

Consider the following program to compute x^p :

```

y := 1;
q := 0;
while q < p {
  y := y · x;
  q := q + 1;
}

```

The desired postcondition we would like to ensure is $y = x^p$; that is, the final value of the program variable y is the p th power of x . We would also like to ensure that the program halts. One essential precondition needed to ensure halting is $p \geq 0$, because the program will only halt and compute x^p correctly if that holds. Note that $p > 0$ will also guarantee that the program halts and produces the correct output, but this is a stronger condition (is satisfied by fewer states, has more logical consequences).

$$\underbrace{p > 0}_{\text{stronger}} \Rightarrow \underbrace{p \geq 0}_{\text{weaker}}$$

The weaker precondition is better because it is less restrictive of the possible starting values of p that ensure correctness. Typically, given a postcondition expressing a desired property of the output state, we would like to know the *weakest precondition* that guarantees that the program halts and satisfies that postcondition upon termination.

4 Partial vs Total Correctness

Two approaches to program verification are:

- *Partial correctness*: check if program is correct when it terminates. This is characterized by wlp and the Hoare logic we will define shortly. The termination issue is handled separately.
- *Total correctness*: ensure both that the program terminates and that it is correct. This is characterized by wp.

Partial correctness is the more common approach, since it separates the two issues of correctness and termination. These two verification tasks use very different methods, and it is helpful to separate them. Often partial correctness is easier to establish, and once this is done, the correctness conditions can be used in conjunction with a well-founded relation to establish termination.

5 Syntax of Hoare Logic

To define Hoare logic, we need to define the well-formed formulas in the logic. Hoare logic is built on top of another conventional logic, such as first-order logic. For now, let us take first-order logic as our base logic. Let φ, ψ, \dots denote first-order formulas. The formulas of Hoare logic are *partial correctness assertions* (PCA's), also known as *Hoare triples*. They look like

$$\{\varphi\}c\{\psi\}.$$

Informally, this means, “if φ holds before execution of c , and if c terminates, then ψ will hold upon termination.” This is equivalent to

$$\varphi \Rightarrow \text{wlp } c \psi.$$

5.1 Proof Rules

We will discuss the semantics of Hoare logic later. For now, we just give the deduction rules for the language IMP with programs

$$c ::= \text{skip} \mid x := a \mid c_0 ; c_1 \mid \text{if } b \text{ then } c_1 \text{ else } c_2 \mid \text{while } b \text{ do } c$$

The rules are

(skip)	$\{\varphi\}\text{skip}\{\varphi\}$
(assignment)	$\{\varphi\{a/x\}\}x := a\{\varphi\}$
(sequential composition)	$\frac{\{\varphi\}c_1\{\psi\} \quad \{\psi\}c_2\{\sigma\}}{\{\varphi\}c_1 ; c_2\{\sigma\}}$
(conditional)	$\frac{\{b \wedge \varphi\}c_1\{\psi\} \quad \{\neg b \wedge \varphi\}c_2\{\psi\}}{\{\varphi\}\text{if } b \text{ then } c_1 \text{ else } c_2\{\psi\}}$
(while)	$\frac{\{b \wedge \varphi\}c\{\varphi\}}{\{\varphi\}\text{while } b \text{ do } c\{\varphi \wedge \neg b\}}$
(weakening)	$\frac{\varphi \Rightarrow \varphi' \quad \{\varphi'\}c\{\psi'\} \quad \psi' \Rightarrow \psi}{\{\varphi\}c\{\psi\}}.$

In the assignment rule, $\varphi\{a/x\}$ denotes the safe substitution of the arithmetic expression a for the variable x in φ . As with the λ -calculus, there may be bound variables in φ bound by quantifiers \forall and \exists , and these may have to be renamed to avoid capturing the free variables of a . In the weakening rule, the operator \Rightarrow is implication in the underlying logic. Note the parallels between these rules and the definitions of wlp.

6 Soundness and Completeness

A deduction system defines what it means for a formula to be *provable*, whereas a semantics defines what it means for a formula to be *true*. Given a logic with a semantics and a deduction system, two desirable properties are

- *Soundness*: Every provable formula is true.
- *Completeness*: Every true formula is provable.

Soundness is a basic requirement of any logical system. A logic would not be good for much if its theorems were false! With respect to the small-step or big-step semantics of IMP, Hoare logic is sound.

Completeness, on the other hand, is a much more difficult issue. Hoare logic, as presented, is not complete in general. However, it is *relatively complete* given an oracle for truth in the underlying logic, provided that logic is expressive enough to express weakest preconditions. This is a famous result of Stephen Cook. Although first-order logic is not expressive enough to express weakest preconditions over arbitrary domains of computation, it is expressive enough over \mathbb{N} or \mathbb{Z} . Therefore Hoare logic is relatively complete for IMP programs over the integers.