

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 Weakest Preconditions

Given a program S and a postcondition φ , the weakest property of the input state that guarantees that S halts in a state satisfying φ is called the *weakest precondition* of S and φ and is denoted $\text{wp } S \varphi$. This says that

- $\text{wp } S \varphi$ implies that S terminates in a state satisfying φ ($\text{wp } S \varphi$ is a precondition of S and φ),
- if ψ is any other condition that implies that S terminates in a state satisfying φ , then $\psi \Rightarrow \text{wp } S \varphi$ ($\text{wp } S \varphi$ is the *weakest precondition* of S and φ).

As in the λ -calculus, juxtaposition represents function application, so wp can be viewed as a higher-order function that takes a program S and a postcondition φ and returns the weakest precondition of S and φ . The function wp can also be viewed as taking a program and returning a function that maps postconditions to preconditions. For this reason, axiomatic semantics is sometimes known as *predicate-transformer semantics*.

5 Guarded Commands

Dijkstra introduced the Guarded Command Language (GCL) with the grammar

$$\begin{aligned} S ::= & \text{skip} \mid x := E \mid S_1; S_2 \\ & \mid \text{if } B_1 \rightarrow S_1 \parallel B_2 \rightarrow S_2 \parallel \dots \parallel B_n \rightarrow S_n \text{ fi} \\ & \mid \text{do } B_1 \rightarrow S_1 \parallel B_2 \rightarrow S_2 \parallel \dots \parallel B_n \rightarrow S_n \text{ od} \end{aligned}$$

where the B_i are Boolean expressions. The B_i are called *guards* because they guard the corresponding statements S_i . The symbol \parallel is the *nondeterministic choice operator* and is not to be confused with $|$. In if and do statements, a clause $B_i \rightarrow S_i$ is said to be *enabled* if its guard B_i is true.

Informally, when executing the if statement, at least one of its clauses must be enabled, otherwise it is a runtime error. One of the enabled clauses $B_i \rightarrow S_i$ is chosen nondeterministically and the corresponding statement S_i is executed. The do statement works similarly, except that there is no requirement that at least one clause be enabled. If none are enabled, execution just falls through to the following statement. If at least one is enabled, then one of the enabled clauses is chosen nondeterministically for execution. After the clause is executed, the guards are reexamined, and the process is repeated. This process repeats until all guards become false.

6 Weakest Preconditions in GCL

We now show how to determine the weakest preconditions for each part of GCL, as well as provide generic examples and special cases of the wp function.

Skip

Since skip does not do anything, we have $\text{wp skip } \varphi \Leftrightarrow \varphi$. Examples:

- $\text{wp skip } (x = 1) \Leftrightarrow x = 1$
- $\text{wp skip false} \Leftrightarrow \text{false}$.

Assignments

For assignments $x := E$,

$$\text{wp } (x := E) \varphi \Leftrightarrow \varphi \{E/x\}.$$

Here $\varphi\{E/x\}$ denotes safe substitution of E for x in the formula φ , in which variables bound by quantifiers $\forall x$ or $\exists x$ are renamed if necessary to avoid capture. Note that the mapping $\varphi \mapsto \varphi\{E/x\}$ goes right-to-left; that is, $\varphi\{E/x\}$ is the precondition that must hold of the input state in order to ensure that the postcondition φ holds of the output state. Examples:

- **wp** $(x := 1) (x = 1) \Leftrightarrow \text{true}$. In words, $x = 1$ is true after $x := 1$ no matter what holds before execution.
- **wp** $(y := 1) (x = 1) \Leftrightarrow x = 1$.
- **wp** $(x := y) (x = 1) \Leftrightarrow (x = 1)\{y/x\} \Leftrightarrow y = 1$.
- **wp** $(x := x + 1) (x = 3)$
 - $\Leftrightarrow (x = 3)\{x + 1/x\}$ by definition of assignment
 - $\Leftrightarrow x + 1 = 3$ by substitution
 - $\Leftrightarrow x = 2$ by arithmetic,
 so **wp** $(x := x + 1) (x = 3) \Leftrightarrow x = 2$.

Sequential Composition

To determine the weakest precondition for which φ holds after executing $S_1; S_2$, we first find the weakest precondition for which φ holds after the execution of S_2 , and then determine the weakest precondition that ensures that property after S_1 :

$$\text{wp } (S_1; S_2) \varphi \Leftrightarrow \text{wp } S_1 (\text{wp } S_2 \varphi).$$

Examples:

- **wp** $(x := 1; y := 2) (x = 1 \wedge y = 2)$
 - $\Leftrightarrow \text{wp } (x := 1) (\text{wp } (y := 2) (x = 1 \wedge y = 2))$ by definition defn of ;
 - $\Leftrightarrow \text{wp } (x := 1) (x = 1 \wedge 2 = 2)$ by definition of assignment
 - $\Leftrightarrow 1 = 1 \wedge 2 = 2$ by definition of assignment
 - $\Leftrightarrow \text{true}$ by predicate calculus.
- **wp** $(x := x + 1; y := y - 1) (x \leq y)$
 - $\Leftrightarrow \text{wp } (x := x + 1) (\text{wp } (y := y - 1) (x \leq y))$ by definition defn of ;
 - $\Leftrightarrow \text{wp } (x := x + 1) (x \leq y - 1)$ by definition of assignment
 - $\Leftrightarrow x + 1 \leq y - 1$ by definition of assignment
 - $\Leftrightarrow y - x \geq 2$ by arithmetic.

If

In an if statement, at least one guard B_i must be true. This condition is expressed by the disjunction $B \triangleq \bigvee_i B_i$. The S_i that is chosen for execution is chosen nondeterministically among all enabled clauses, and in order to guarantee the postcondition φ , all enabled clauses had better guarantee φ . Thus

$$\text{wp } (\text{if } B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \text{ fi}) \varphi \Leftrightarrow B \wedge \bigwedge_i (B_i \Rightarrow \text{wp } S_i \varphi).$$

Example: The following program computes the maximum of two numbers.

$$\text{MAX} \triangleq \text{if } x \geq y \rightarrow z := x \parallel y \geq x \rightarrow z := y \text{ fi}$$

To prove that the program halts and correctly computes the maximum of x and y regardless of input state, it suffices to show that **true** is the weakest precondition corresponding to the postcondition $z = \max x, y$.

$$\begin{aligned}
& \mathbf{wp} \text{ MAX } (z = \max x, y) \\
& \Leftrightarrow x \geq y \vee y \geq x \\
& \quad \wedge (x \geq y \Rightarrow (\mathbf{wp} (z := x) (z = \max x, y))) \\
& \quad \wedge (y \geq x \Rightarrow (\mathbf{wp} (z := y) (z = \max x, y))) \quad \text{by definition of if} \\
& \Leftrightarrow \mathbf{true} \\
& \quad \wedge (x \geq y \Rightarrow x = \max x, y) \\
& \quad \wedge (y \geq x \Rightarrow y = \max x, y) \quad \text{by predicate calculus and the definition of assignment} \\
& \Leftrightarrow \mathbf{true} \quad \text{by predicate calculus.}
\end{aligned}$$

Do

Since **do** has the complication that it may not terminate, it is difficult to formalize its weakest precondition. In fact, over arbitrary structures, first-order predicate logic is not sufficiently expressive to formulate weakest preconditions for this construct. However, we can use infinitary logic (logic with infinite conjunctions and disjunctions). We can write

$$\mathbf{wp} (\mathbf{do} B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \mathbf{od}) \varphi \Leftrightarrow \bigvee_k P_k,$$

where informally P_k is the weakest precondition ensuring that the **do** statement terminates in exactly k iterations and satisfies φ upon termination. Formally, let E be the body of the **do** statement (thus the statement is **do** E **od**), and let $B = \bigvee_i B_i$ as above. Define inductively

$$\begin{aligned}
P_0 & \triangleq \neg B \wedge \varphi, & (1) \\
P_{k+1} & \triangleq B \wedge \mathbf{wp} (\mathbf{if} E \mathbf{fi}) P_k. & (2)
\end{aligned}$$

The basis condition (1) says that no clause is enabled and φ is true of the input state, which is equivalent to the condition that the body E of the **do** statement is executed exactly 0 times and terminates in a state satisfying φ . The inductive condition (2) says that B is true, thus the body E of the **do** statement is executed at least once, and after executing the body once, P_k will hold, implying that the **do** statement will execute exactly k more times and satisfy φ upon termination.

If you don't like infinitary logic, you can do the same thing with the μ -calculus predicate

$$\mu P. (\neg B \wedge \varphi) \vee (B \wedge \mathbf{wp} (\mathbf{if} E \mathbf{fi}) P),$$

which denotes the least fixpoint of the monotone map $\lambda P. (\neg B \wedge \varphi) \vee (B \wedge \mathbf{wp} (\mathbf{if} E \mathbf{fi}) P)$ on predicates.