

# 1 Program verification

We have covered a number of general topics so far in the course: propositional logic, predicate logic, recursively defined sets, and modal logic. Each of these topics might also be covered in a course on mathematical logic and you may be wondering about the connection to computer science.

In fact, one of the most important—perhaps the single most important—applications of formal logic is to computer science and related areas like electrical engineering. You have already seen examples of using formal methods to solve computational problems in sports tournament scheduling and network security analysis. These applications were quite specific; here we will focus on applying formal methods to a much more general problem: *program verification*.

The goal of program verification is to prove that some program of interest has certain behavior. We have already seen some simple examples of this, such as when we used induction to prove that the addition function on naturals is commutative. That kind of proof is based on *operational semantics*—that is, the execution behavior of the code (*e.g.*, the add function). Proofs based on operational semantics are quite useful but hard to generalize to larger programs.

Our goal now is to introduce an alternative technique for program verification called *Hoare logic*. Hoare logic is a kind of modal logic that is specialized for reasoning about program verification. We proceed by example; first we explain a simple programming language and provide a model for computational state. Next, we define a modal logic of assertions and provide an *axiomatic semantics* for our language by defining series of *Hoare rules*. Finally, we show how to apply those rules to reason about a number of example programs.

# 2 Programming language

We start by defining a very simple programming language. The language is too brain-dead to make programming in it much fun, since it lacks lots of features that make a programmer's task easier (*e.g.*, function calls). Still, it is Turing complete—that is, it is capable of doing any computation that any other machine is capable of doing, if the programmer is willing to work hard enough.

**Expressions.** We start with the notions of numerical and boolean expressions:

$$\begin{array}{ll} \text{numerical expressions } e & \equiv z \mid x \mid e_1+e_2 \mid e_1-e_2 \mid e_1*e_2 \\ \text{boolean expressions } b & \equiv e_1<=e_2 \mid b_1||b_2 \mid !b \end{array}$$

Numerical expressions  $e$  are inductive types with two kinds of base elements: integers  $z$  or program variables  $x$ . The recursive cases of  $e$  are for the sum, difference, and product of two other numerical expressions  $e_1$  and  $e_2$ , respectively. Boolean expressions have only one kind of base element, the comparison of two numerical expressions  $e_1$  and  $e_2$ ; in addition they have two recursive cases, for disjunction and negation. Notice that we are missing certain other

obvious boolean expressions; it is convenient to define the following macros:

$$\begin{aligned}
 \mathbf{true} &\equiv 0 <= 0 \\
 \mathbf{false} &\equiv \mathbf{!true} \\
 b_1 \&\& b_2 &\equiv \mathbf{!((!b_1) \vee (!b_2))} \\
 e_1 == e_2 &\equiv (e_1 <= e_2) \&\& (e_2 <= e_1) \\
 e_1 != e_2 &\equiv \mathbf{!(e_2 == e_1)} \\
 e_1 < e_2 &\equiv (e_1 <= e_2) \&\& (e_1 != e_2)
 \end{aligned}$$

These macros are as you would expect; notice that conjunction uses deMorgan.

**Expression evaluation.** Once we have an expression, the natural thing to do is to evaluate it into some (numerical or boolean) value. For simple numerical expressions like the following, this is not so hard:

$$3 + (2 * (0 - 2))$$

Clearly this expression evaluates to  $-1$ . The tricky point is that numerical expressions can contain program variables, like this one (let's call it  $e_a$ ):

$$e_a \equiv m * (2 + a)$$

Here  $m$  and  $a$  are two program variables. The value of the expression  $e_a$  clearly depends on the values of  $m$  and  $a$ : for example, if  $m = 2$  and  $a = 3$ , then the expression  $e_a$  evaluates to 10. To help, we introduce the notion of a *context*:

$$\text{context } \rho \equiv \mathbb{X} \rightarrow \mathbb{Z}$$

That is, a context  $\rho$  is a function<sup>1</sup> from the set from the set of variable names  $\mathbb{X}$  to the set of integers  $\mathbb{Z}$ . The context above (we'll call it  $\rho_a$ ) is defined as:

$$\rho_a(x) \equiv \begin{cases} 2 & \text{when } x = m \\ 3 & \text{when } x = a \\ \text{undefined} & \text{otherwise} \end{cases}$$

Given this notion of context, it is easy to define the numeric evaluator function  $\text{neval}(\rho, e)$  that evaluates  $e$  in the context  $\rho$  by:

$$\text{neval}(\rho, e) \equiv \begin{cases} z & \text{when } e = z \\ \rho(x) & \text{when } e = x \\ \text{neval}(\rho, e_1) + \text{neval}(\rho, e_2) & \text{when } e = e_1 + e_2 \\ \text{neval}(\rho, e_1) - \text{neval}(\rho, e_2) & \text{when } e = e_1 - e_2 \\ \text{neval}(\rho, e_1) \times \text{neval}(\rho, e_2) & \text{when } e = e_1 * e_2 \end{cases}$$

Hopefully our “overloading” of the arithmetical symbols is clear. When we write  $\text{neval}(\rho, e_1) + \text{neval}(\rho, e_2)$  we mean to use the standard addition function

---

<sup>1</sup>We use the symbol “ $\rightarrow$ ” here to indicate the set of functions.

on integers, whereas when we write  $e = e_1 + e_2$  we are indicating one of the constructors of the inductively defined set of numerical expressions. Although usually it is not hard to understand which symbol we mean in a given context, we also use different fonts to help distinguish which meaning is indicated. Using our previously defined expression  $e_a$  and context  $\rho_a$ , it should be clear that

$$\text{neval}(\rho_a, e_a) = 10$$

Boolean expressions are evaluated in much the same way; evaluation here also needs a context so that numerical comparisons can be properly evaluated:

$$\text{beval}(\rho, b) \equiv \begin{cases} \text{neval}(\rho, e_1) \leq \text{neval}(\rho, e_2) & \text{when } b = e_1 \leq e_2 \\ \text{beval}(\rho, b_1) \vee \text{beval}(\rho, b_2) & \text{when } b = b_1 \vee b_2 \\ \neg \text{beval}(\rho, b') & \text{when } b = !b' \end{cases}$$

Hopefully our notation here is straightforward. The `beval` function turns boolean expressions into truth values like  $\top$  and  $\perp$ .

**Commands.** Now that we have defined the expressions (and their evaluation) in our language, we are ready to define its commands. Here they are:

Syntax	Command	Description
<code>skip</code>	<code>skip</code>	Do nothing; <code>nop</code>
<code>x := e</code>	<code>assign</code>	Store $e$ into the variable $x$
<code>c<sub>1</sub>; c<sub>2</sub></code>	<code>sequence</code>	Run $c_1$ followed by $c_2$
<code>if (b) {c<sub>1</sub>} else {c<sub>2</sub>}</code>	<code>if</code>	If $b$ is <b>true</b> then run $c_1$ else run $c_2$
<code>while (b) {c}</code>	<code>loop</code>	As long as $b$ is <b>true</b> then run $c$

We have only five commands; the result is an imperative language like C or Java, but hugely simplified. One (very useful) way to think about commands is that they transform a given context  $\rho_1$  into a succeeding context  $\rho_2$ .

The first command, `skip`, has no effect—the succeeding context is equal to the starting context. The second command, `x := e` (assignment), evaluates the expression  $e$  and stores the result into the program variable  $x$ —the succeeding context has a new value for variable  $x$ . The third command, `c1; c2` (sequence), runs the command  $c_1$  and then runs the command  $c_2$ . The succeeding context of  $c_1$  becomes the preceding context of  $c_2$ , and the succeeding context of the entire sequence is the succeeding context of  $c_2$ . The fourth command, `if (b) {c1} else {c2}` (if/conditional), evaluates  $b$ ; if the result is **true** then it runs  $c_1$ ; otherwise it runs  $c_2$ . The succeeding context of the conditional is the succeeding context of whichever instruction ran. Finally, the fifth command, `while (b) {c}` (while), evaluates  $b$  and if the result is **true** runs  $c$ , and then loops around (evaluates  $b$  again, which may have a different value after running  $c$ , and if  $b$  still evaluates to **true** then runs  $c$  again, before looping around again and testing  $b$ ...). All of these commands should be familiar to you (except maybe `skip`; if you want you can substitute the assignment  $x := x$ ).

**Example 1.** To illustrate our language, consider this simple program:

```
y := x;
while (2<=y)
{
  y := y-2
}
```

What does this program (let us call it “Parity”) do? Parity first copies the value of  $x$  into  $y$  and then sets up a loop that keeps going until the program variable  $y$  is less than 2; until this is the case, it keeps decrementing  $y$  by 2. The intention of the Parity program is that if the program variable  $x$  starts as an even number, then when the program ends  $y$  will be equal to 0. On the other hand, when  $x$  starts as an odd number, then when the program ends  $y$  will be equal to 1. Thus, Parity seems to determine if a number is even or odd.

**Example 2.** Consider also the slightly more complicated “Factorial” program:

```
y := 1;
z := 0;
while (z!=x)
{
  z := z+1;
  y := y*z
}
```

Here we start by setting the program variables  $y$  and  $z$  to 1 and 0, respectively. We then enter a loop where we increment  $z$  and multiply the value of  $y$  by the new value of  $z$ ; we exit the loop when  $z==x$ . The purpose of the Factorial program is to compute  $x!$  into  $y$ —*i.e.*, after the program runs we should have

$$y = 1 \times 2 \times \dots \times (x - 1) \times x$$

**Exercise 1.** (\*) Prove that our language is Turing complete. The major challenge is representing an infinite tape with a finite number of named variables.

### 3 Hoare Logic

Now that we have an understanding of our programming language, we turn to the problem of program verification. The first question must be, what do we mean by “proving that a program has certain behavior”? This is actually a pretty complicated question. For example, quicksort and mergesort are two standard sorting algorithms, although they produce the same result given the same input, they are not the same algorithm (*e.g.*, they may use different amounts of memory, or have different running times).

We have to determine which properties of programs we are interested in verifying. The most fundamental property of a program is neither its running time nor its memory requirements, but rather whether it produces the output you expect for a given input. Another way of saying that is, given some starting context  $\rho_\alpha$  that satisfies some initial property  $P_\alpha$ , once the program terminates, the ending context  $\rho_\omega$  satisfies some final property  $P_\omega$ . For example, in the Parity program, we might like to verify that given any starting context, the ending context must have  $y$  equal to 0 if  $x$  is even, and  $y$  equal to 1 if  $x$  is odd.

**Assertions.** You may be noticing a connection to previous material. The truth or falsehood of a property **depends on the current context**. In other words, we have a modal logic in which our set of worlds are our program contexts. It is entirely reasonable to borrow notation used in modal logic as well and write:

$$\rho_\alpha \Vdash P_\alpha$$

to say that the initial property  $P_\alpha$  holds in the initial context  $\rho_\alpha$ , and

$$\rho_\omega \Vdash P_\omega$$

to say that the final property  $P_\omega$  holds in the final context  $\rho_\omega$ .

To define the logical operators in our modal logic we do our usual trick of lifting operators from the metalogic. When we did this for propositional and predicate logic, our metalogic was *English*, which was good because it provided intuition but troublesome because natural language is often ambiguous. Now that we understand propositional and predicate logic, however, they are perfectly reasonable candidates for our metalogic. We proceed as follows:

$\rho \Vdash \top$	$\equiv$	$\top$	truth constant falsehood constant conjunction disjunction implication negation
$\rho \Vdash \perp$	$\equiv$	$\perp$	
$\rho \Vdash P \wedge Q$	$\equiv$	$(\rho \Vdash P) \wedge (\rho \Vdash Q)$	
$\rho \Vdash P \vee Q$	$\equiv$	$(\rho \Vdash P) \vee (\rho \Vdash Q)$	
$\rho \Vdash P \rightarrow Q$	$\equiv$	$(\rho \Vdash P) \rightarrow (\rho \Vdash Q)$	
$\rho \Vdash \neg P$	$\equiv$	$P \rightarrow \perp$	

On the left-hand side we give the new operators of our modal logic, and on the right-hand side we provide their definitions using the operators of our metalogic (predicate logic)<sup>2</sup>. To avoid massive symbol overload we use the same symbols on both sides of the definitions but of course they are not the same.

We also want to have universal  $\forall$  and existential  $\exists$  quantifiers in our modal logic. Fortunately, these too can be lifted straight from our metalogic:

$\rho \Vdash \forall x P$	$\equiv$	$\forall x (\rho \Vdash P)$	universal existential
$\rho \Vdash \exists x P$	$\equiv$	$\exists x (\rho \Vdash P)$	

Here we assume that  $x$  is free in  $P$  so that the quantifiers can provide meaning.

<sup>2</sup>This is exactly the same technique that was used in the most recent Coq homework and you may find it useful to review how it was done there.

It is often very useful to embed predicate logic formulas into modal formulas. We define the *plain lift* operator  $\langle \cdot \rangle$  to do just that:

$$\rho \Vdash \langle P \rangle \quad \equiv \quad P \quad | \quad \text{plain lift}$$

That is,  $\langle P \rangle$  holds in any context  $\rho$  if and only if  $P$  is true in predicate logic.

All of the operators defined so far are either independent of the context  $\rho$  ( $\top$ ,  $\perp$ ,  $\langle \cdot \rangle$ ) or simply “pass on” the given context to their subformulas. What about formulas whose truth critically depends on the context? The simplest of these allow us to lift numeric and boolean expressions into our assertion logic:

$$\begin{array}{l} \rho \Vdash e \Downarrow z \quad \equiv \quad \text{neval}(\rho, e) = z \\ \rho \Vdash [b] \quad \equiv \quad \text{beval}(\rho, b) = \top \end{array} \quad \left| \begin{array}{l} \text{numeric evaluation} \\ \text{boolean evaluation} \end{array} \right.$$

When  $e \Downarrow z$  holds on some context  $\rho$ , then the evaluation of  $e$  in  $\rho$  is equal to  $z$ ; similarly, when  $[b]$  holds on some context  $\rho$ , then  $b$  evaluates to  $\top$  in  $\rho$ .

You probably remember that when we define a modal logic we usually define not only a set of worlds but also a relation between worlds. Here we will generalize this notion by allowing many different kinds of relations between worlds—in other words, we have a *multimodal* logic. One important family of world relations is the update relation  $U_{(x,n)}$  defined as follows:

$$\rho U_{(x,n)} \rho' \quad \equiv \quad \rho' = [x \mapsto n]\rho$$

In other words, two contexts  $\rho$  and  $\rho'$  are related by  $U_{(x,n)}$  if they are equal everywhere other than at variable  $x$ , and the value of  $\rho'$  at  $x$  is equal to  $n$ .

We can define the modal operator  $\Box_{(x,n)}^U$  in the usual way by universally quantifying over reachable worlds via this relation:

$$\rho \Vdash \Box_{(x,n)}^U P \quad \equiv \quad \forall \rho' ((\rho U_{(x,n)} \rho') \rightarrow (\rho' \Vdash P)) \quad | \quad \text{store update modality}$$

In other words, if  $\Box_{(x,n)}^U P$  holds on some context  $\rho$ , then  $P$  would hold if  $\rho$  **were updated so that variable  $x$  had value  $n$** . We can also define the related:

$$\rho \Vdash [x \mapsto e]P \quad \equiv \quad \rho \Vdash \Box_{(x, \text{neval}(\rho, e))}^U P \quad | \quad \text{store expression update modality}$$

In this case,  $P$  would hold if  $\rho$  were updated so that variable  $x$  had the result of evaluating  $e$  in  $\rho$ . We will see how this modality is used shortly.

What about other kinds of relations between worlds? There is a second very important class of relations. Notice that **the commands of our language** move from one world to the next by computing the subsequent context from the preceding one. For now we will avoid formal definitions of this modality, but we will return to it next week when we study the soundness of Hoare logic.

**Hoare rules.** Now that we have defined a modal logic of properties/assertions, we can return to the considering our goals for the verification. We are trying to find a method for proving that from a context that satisfies some initial property  $P_\alpha$ , we reach a final context that satisfies some final property  $P_\omega$ .

This is exactly the idea behind the *Hoare triple*. A Hoare triple, written

$$\{P\} c \{Q\}$$

is a way of saying that given any starting context that satisfies the property  $P$ , after running the command  $c$ , we will be in an ending context that satisfies the property  $Q$ . Of course, the truth of a particular Hoare tuple depends very much on the particular  $P$ ,  $c$ , and  $Q$ ! Some tuples seem pretty reasonable, such as:

$$\{x \Downarrow 5\} x := x+1 \{x \Downarrow 6\}$$

That is, starting from any context in which the expression  $x$  evaluates to 5, after incrementing  $x$ , we will be in a context in which  $x$  evaluates to 6. In contrast, the following tuple does not seem so reasonable:

$$\{x \Downarrow 5\} x := x+1 \{\perp\}$$

In other words, this Hoare tuple is claiming that, given a starting context in which  $x$  evaluates to 5, by incrementing  $x$  we reach no context at all (that is, the increment command “blows up”—or maybe increment never terminates). This seems pretty unlikely, since a reasonable intuition about the action of the increment command indicates that such a context should exist.

Unfortunately, we are not quite in a position to prove that the second Hoare tuple is invalid, since we have not pinned down the formal behavior of our commands, or given a formal definition for Hoare tuples. These important topics are the focus of next week’s lecture. Until then, we will proceed axiomatically: first, we give a simple series of rules (axioms) justified by our intuition; second, we show how to combine those axioms into proofs about whole programs.

Our first axiom is for `skip`; since `skip` is very simple (do nothing!), you would expect that it would have a very simple axiom, and it does:

$$\frac{}{\{P\} \text{skip} \{P\}} \text{Skip}$$

In other words, the `Skip` axiom says that for any starting property  $P$ , given an initial context  $\rho$  satisfying  $P$ , after running the command `skip` (which does nothing), we reach a concluding context (which is just  $\rho$  again) that satisfies  $P$ .

The axiom for  $x := e$  (assignment) uses the store expression update modality:

$$\frac{}{\{[x \mapsto e] P\} x := e \{P\}} \text{Assignment}$$

In other words, supposing that some property  $P$  would hold if the initial context were updated so that variable  $x$  mapped to (the evaluated) expression  $e$ . Since the assignment command does exactly that—transform the initial context into a new context in which  $x$  now contains the (evaluated) expression  $e$ , it makes perfect sense that  $P$  would hold in the concluding context.

You might notice that our Hoare axiom `Assignment` does not quite fit the pattern of our initial (reasonable) example of a Hoare triple above. To validate

that program, we will need another Hoare axiom called **Consequence**:

$$\frac{P \rightarrow P' \quad Q' \rightarrow Q \quad \{P'\} c \{Q'\}}{\{P\} c \{Q\}} \text{Consequence}$$

In other words, suppose we have proved some Hoare tuple  $\{P'\} c \{Q'\}$ , and further suppose that we have some (possibly) stronger precondition  $P$  and some (possibly) weaker postcondition  $Q$ . In that case, we are justified in concluding  $\{P\} c \{Q\}$ : if we know that our initial satisfies the (stronger)  $P$ , we know that it must also satisfy  $P'$ ; thus we know that the concluding state must satisfy the (stronger)  $Q'$ , which implies that it also satisfies  $Q$ . Let's see this in action:

$$\frac{\begin{array}{c} (x \Downarrow 5) \rightarrow ([x \mapsto x+1](x \Downarrow 6)) \\ (x \Downarrow 6) \rightarrow (x \Downarrow 6) \end{array} \quad \frac{}{\{[x \mapsto x+1](x \Downarrow 6)\} x := x+1 \{x \Downarrow 6\}} \text{Assignment}}{\{x \Downarrow 5\} x := x+1 \{x \Downarrow 6\}} \text{Consequence}$$

We have successfully verified the desired specification (precondition/postcondition) pair. Notice that in order to use the **Consequence** rule we had to:

1. Prove that our preferred precondition  $(x \Downarrow 5)$  is stronger than the precondition required by **Assignment**  $([x \mapsto x+1](x \Downarrow 6))$ . This proof is not contained in the above derivations, but fortunately it is not very difficult. If we are in a context where  $x$  has value 5, then if we move to a new context in which  $x$  contains the value is one greater than the old value of  $x$ —then we will be in a context in which  $x$  has value 6.
2. Prove that our desired postcondition is implied by the postcondition we get from the implication rule; since they are the same that is very simple.
3. Apply the **Assignment** rule to get the underlying Hoare triple.

All of the previously rules only apply to program fragments with a single command. One of the “big wins” of Hoare logic is that the rules can be composed in a modular way. The first composition rule is **Sequence**:

$$\frac{\{P\} c_1 \{Q\} \quad \{Q\} c_2 \{R\}}{\{P\} c_1; c_2 \{R\}} \text{Sequence}$$

That is, from specifications of the subcommands  $c_1$  and  $c_2$ , we can derive specifications for the combined sequence  $c_1; c_2$ . We apply **Consequence** if the obvious postcondition for  $c_1$  does not exactly match the obvious precondition for  $c_2$ . Let's take a look at an example of verifying a sequence<sup>3</sup>:

$$\frac{\frac{\dots \text{exactly as above} \dots}{\{x \Downarrow 5\} x := x+1 \{x \Downarrow 6\}} \text{Conseq.} \quad \frac{\dots \text{similar to above} \dots}{\{x > 5\} x := x+1 \{x > 6\}} \text{Conseq.}}{\frac{\{x \Downarrow 6\} x := x+1 \{x > 6\}}{\{x \Downarrow 5\} x := x+1; x := x+1 \{x > 6\}} \text{Sequence}}$$

<sup>3</sup>Here we write, *e.g.*,  $x > 6$  as shorthand for  $\exists n ((x \Downarrow n) \wedge \langle n > 6 \rangle)$ .

To avoid running into the margins in the preceding derivation we omitted the following side conditions of the lower-right application of **Consequence**:

$$(x \Downarrow 6) \rightarrow (x > 5) \quad \text{and} \quad (x > 6) \rightarrow (x > 6)$$

Of course, these are very easy. Still, it is easy to observe that our derivations are growing rapidly. We need to introduce a better notation to handle the problem. The idea is that we will write the assertions in between the lines of the program, applying the natural rules to get us from one assertion to the next. For example, here is the same verification in the new format:

```

{x \Downarrow 5}
{[x \mapsto x+1](x \Downarrow 6)}
x := x+1;
{x \Downarrow 6}
{x > 5}
{[x \mapsto x+1](x > 6)}
x := x+1
{x > 6}

```

Clearly this kind of format is more compact than full proof-tree derivations. In fact, it is (usually) not very hard to decode what is going on, since there are only two possibilities for getting from assertion to the next.

In the first case, applicable whenever there are two assertions in a row, the rule of consequence has been applied to get from one assertion to the next. If there are more than two assertions in a row, then consequence has been applied repeatedly, usually because the verifier wants to provide additional clarity in his proof. For example, the first two lines of the verification above are assertions, and to get from the first to the second we must apply **Consequence**. This means that a side proof, usually omitted since it is usually very easy, must show that the preceding assertion must imply the succeeding assertion.

In the second case, applicable whenever two assertions are separated by a single command, the underlying proof rule must be applied. Here, since both commands are assignments, the **Assignment** rule has been applied. The **Sequence** rule is implied by the way the notation proceeds from one command to the next.

There are only a few axioms remaining. Let us examine the rule for **if**:

$$\frac{\{P \wedge [b]\} c_1 \{Q\} \quad \{P \wedge \neg[b]\} c_2 \{Q\}}{\{P\} \text{if } (b) \{c_1\} \text{else } \{c_2\} \{Q\}} \text{if}$$

This rule is somewhat similar to the disjunction-elimination rule from propositional logic. That is, if we can show that starting from a context that satisfies  $P$ , then running either  $c_1$  or  $c_2$  leaves us in a context that satisfies  $Q$ , then we can conclude that the entire conditional (if statement) has postcondition  $Q$ . The only wrinkle is that while verifying  $c_1$  we are allowed to assume that the boolean was true ( $[b]$ ) when verifying the postcondition  $Q$ ; while verifying  $c_2$  we are allowed to assume that the boolean was not true ( $\neg[b]$ ).

To demonstrate the `If` rule, consider the following verification:

```

{(x ↓ 0) ∨ (x ↓ 1)}
if (x==0)
{
  {((x ↓ 0) ∨ (x ↓ 1)) ∧ [x==0]}
  {x ↓ 0}
  {[x ↦ x+2](x ↓ 2)}
  x := x+2
  {x ↓ 2}
} else {
  {((x ↓ 0) ∨ (x ↓ 1)) ∧ ¬[x==0]}
  {x ↓ 1}
  {[x ↦ x+1](x ↓ 2)}
  x := x+1
  {x ↓ 2}
}
{x ↓ 2}

```

Notice we carefully distributed the precondition (plus  $[b]/\neg[b]$ ) to both branches of the `if`, and ensured that both branches reached the same postcondition.

Finally, let us consider the `while` command. Here we have some decisions to make. Let us first present the following (perfectly reasonable) rule:

$$\frac{\{P \wedge [b]\} c \{P\}}{\{P\} \text{while } (b) \{c\} \{P \wedge [\neg b]\}} \text{While}$$

Suppose the loop body  $c$  can be verified with precondition  $P$  (plus the fact that the boolean  $b$  evaluated to true) and postcondition  $P$ . In that case, it is reasonable to conclude that if we start the `while` from precondition  $P$ , then if/when it finally terminates, it will end with postcondition  $P$  plus the fact that the boolean  $b$  finally evaluated to false. Here is why: the loop body cannot take us “out of”  $P$ ; in other words,  $P$  is an *invariant* of the loop. Therefore, no matter how many times we go around the loop,  $P$  will remain true.

This raises all kinds of tricky questions. The first is, how do we discover the “right” loop invariant  $P$  for a given situation? In general, this is a very hard problem and is an active research problem decades after Hoare logic was developed. We suggest an iterative approach: try some invariant that you think might work and perform the verification. If the Hoare proof breaks somewhere, then take a guess at how the invariant might be changed to fix the break, and start the verification over. Continue until you have found something that works.

To see this rule in action, let’s try to verify the Parity function with precondition  $\top$  (make no restrictions on the starting context) and postcondition

$$((y \downarrow 0) \wedge \text{even}(x)) \vee ((y \downarrow 1) \wedge \text{odd}(x))$$

Note that to avoid large formulas, we will start to use abbreviations like  $x = y$

$$\begin{array}{l}
\{\top\} \\
\{[y \mapsto x](x = y)\} \\
y := x; \\
\{x = y\} \\
\{\exists n (x = y + (2 \times n))\} \\
\text{while } (2 \leq y) \\
\{ \\
\quad \{(\exists n (x = y + (2 \times n))) \wedge [2 \leq y]\} \\
\quad \{\exists n (x = y + (2 \times n))\} \\
\quad \{[y \mapsto y-2](\exists n (x = y + (2 \times n)))\} \\
\quad y := y-2 \\
\quad \{\exists n (x = y + (2 \times n))\} \\
\} \\
\{(\exists n (x = y + (2 \times n))) \wedge \neg[2 \leq y]\} \\
\{(\exists n (x = y + (2 \times n))) \wedge (y < 2)\} \\
\{\dots \text{uh oh} \dots\} \\
\{((y \Downarrow 0) \wedge \text{even}(x)) \vee ((y \Downarrow 1) \wedge \text{odd}(x))\}
\end{array}$$

Figure 1: Attempted verification for Parity

instead of  $\exists n (x \Downarrow n \wedge y \Downarrow n)$ ,  $\text{even}(x)$  instead of  $\exists n (x \Downarrow n \wedge \exists m (m + m = n))$ , and so forth. We give the attempted verification in Figure 1.

The verification goes through just fine until the very end, when we reach the “uh oh”. The problem is that the preceding assertion **does not imply** the desired postcondition. A little examination leads to the reason:  $y$  is an integer; what happens if  $y$  is **negative**? In that case the final conclusion will be false, since  $y$  will neither equal 0 or 1. It turns out that  $y$  will be negative if the initial value of  $x$  is negative. This is exactly the kind of mistake that is very easy to make when thinking about a program informally, and exactly the kind of mistake that a formal technique like Hoare logic will help you find.

We need to fix our specification by strengthening the precondition (*e.g.*, to  $x \geq 0$ ) or weakening the postcondition (*e.g.*, to  $\top$ ). Here the right choice is to strengthen the precondition, yielding the following flawless given in Figure 2. We have put boxes around the formulas that are different in the revised proof; these track the fact that  $y$  is nonnegative during the program’s execution. At the penultimate step we are able to combine the fact that  $y$  is nonnegative and the fact that  $y < 2$  to conclude that  $y$  is either 0 or 1. Note we carefully ensure we reach the loop invariant at the end of the loop body.

**Verifying Factorial.** To give another example, let us reconsider the Factorial program. We would like to show that given any starting context, after the

$$\begin{aligned}
& \{\boxed{x \geq 0}\} \\
& \{[y \mapsto x](x = y \wedge \boxed{y \geq 0})\} \\
& y := x; \\
& \{x = y \wedge \boxed{y \geq 0}\} \\
& \{\exists n (x = y + (2 \times n) \wedge \boxed{y \geq 0})\} \\
& \text{while } (2 \leq y) \\
& \{ \\
& \quad \{(\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 0}) \wedge [2 \leq y]\} \\
& \quad \{\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 2}\} \\
& \quad \{[y \mapsto y-2](\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 0})\} \\
& \quad y := y-2 \\
& \quad \{\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 0}\} \\
& \} \\
& \{(\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 0}) \wedge \neg[2 \leq y]\} \\
& \{(\exists n (x = y + (2 \times n)) \wedge \boxed{y \geq 0}) \wedge (y < 2)\} \\
& \{(y \Downarrow 0) \wedge \text{even}(x) \vee ((y \Downarrow 1) \wedge \text{odd}(x))\}
\end{aligned}$$

Figure 2: Correct verification for Parity

program terminates the value of  $y$  will be equal to the value of  $x!$ —that is,

$$y = 1 \times 2 \times \dots \times (x - 1) \times x$$

This is a little bit trickier to precisely specify than it might appear. The problem is that the factorial function is usually defined on **naturals**, but our program variables are **integers**. What do we mean when we say  $y = x!$  if  $x$  is negative?

A little thought indicates several reasonable possibilities:

1. Define  $z!$  to be  $-1$  if  $z < 0$
2. Define  $z!$  to be equal to  $(-z)!$  when  $z < 0$  (that is, take the absolute value before calculating factorial)
3. Define  $z!$  to be equal to  $-((-z)!)!$  when  $z < 0$  (that is, take the absolute value before calculating factorial, then take the inverse)
4. Say that  $z!$  is undefined when  $z < 0$ , and so when we say  $y = x!$  we must include a side condition that  $y \geq 0$

```

{⊤}
{[y ↦ 1](y ↓ 1)}
y := 1;
{y ↓ 1}
{[z ↦ 0](y ↓ 1 ∧ z ↓ 0)}
z := 0;
{y ↓ 1 ∧ z ↓ 0}
{y = z!}
while (z! = x)
{
  {(y = z!) ∧ [z! = x]}
  {(y × (z + 1)) = (z + 1)!}
  {[z ↦ z+1]((y × z) = z!)}
  z := z+1;
  {(y × z) = z!}
  {[y ↦ y*z](y = z!)}
  y := y*z
  {y = z!}
}
{(y = z!) ∧ ¬[z! = x]}
{(y = z!) ∧ z = x}
{y = x!}

```

Figure 3: Verification for Factorial

The key point here is that if we want to provide a precise specification for Factorial then we need to do something. In this case, any of the above solutions would be fine—as long as you clearly communicate your choice and what you mean by your notation. On a whim we will pick option (2), formalized as follows:

$$z! = \begin{cases} 1 & \text{when } z = 0 \\ z \times ((z - 1)!) & \text{when } z > 0 \\ (-z)! & \text{when } z < 0 \end{cases}$$

Using this definition, we can formalize our desired postcondition as follows:

$$y = x! \equiv \exists n (y \downarrow n \wedge x \downarrow (n!))$$

With this background material out of the way, in Figure 3 we verify Factorial given the precondition  $\top$  and postcondition  $y = x!$ .

```

{⊤}
while (true)
{
  {⊤ ∧ [true]}
  skip
  {⊤ ∧ [true]}
  {⊤}
}
{⊤ ∧ ¬[true]}
{⊥}

```

Figure 4: Verifying that an infinite loop can have postcondition  $\perp$ .

**Exercise 2.** (★) The verification proof in Figure 3 is a standard proof from the literature used to teach Hoare logic, and is contained in a popular introductory textbook on formal methods. Unfortunately, it contains a flaw. What is it?

## 4 Total correctness

Sharp eyes might wonder about a problem with the *While* rule presented above. What happens if the program never terminates? In fact, using the *While* rule it is possible to get some strange results. Reexamine the verification for *Factorial* in Figure 3. Suppose we run this program from an initial context in which  $x$  is negative. In this case, the *while* loop will never terminate, since  $z$  starts at 0 and increments until it hits  $x$ —which it never will if  $x$  is negative.

We need to refine our understanding of what we mean by our Hoare triple  $\{P\} c \{Q\}$ . The postcondition  $Q$  need only hold **if the command  $c$  terminates**. If  $c$  never terminates, then we never reach a final context in which  $Q$  is supposed to hold, and so it does not actually matter what  $Q$  is. Program logics that do not guarantee termination are called logics of *partial correctness*. As a more extreme example of the cost/benefit of using a logic of partial correctness, consider the verification in Figure 4 that a simple infinite loop can be specified with postcondition  $\perp$ . Actually a postcondition of  $\perp$  for a loop is not so confusing, since it means that the loop must never have terminated. The more reasonable-looking postcondition of the *Factorial* program is more problematic, since someone looking at the specification might not realize that the precondition of  $\top$  is not enough to guarantee that *Factorial* will terminate.

The solution is to develop a Hoare logic of *total correctness*. We write<sup>4</sup>

$$[P] c [Q]$$

---

<sup>4</sup>We used a different notation to separate partial and total correctness in the lectures; either is fine as long as you are consistent.

$$\begin{array}{c}
\frac{}{\overline{[P] \text{ skip } [P]}} \text{TSkip} \\
\frac{}{\overline{[[x \mapsto e] P] x := e [P]}} \text{TAssignment} \\
\frac{P \rightarrow P' \quad Q' \rightarrow Q \quad [P'] c [Q']}{[P] c [Q]} \text{TConsequence} \\
\frac{[P] c_1 [Q] \quad [Q] c_2 [R]}{[P] c_1; c_2 [R]} \text{TSequence} \\
\frac{[P \wedge [b]] c_1 [Q] \quad [P \wedge \neg[b]] c_2 [Q]}{[P] \text{ if } (b) \{c_1\} \text{ else } \{c_2\} [Q]} \text{TIf}
\end{array}$$

Figure 5: Hoare axioms for total correctness (except for `while`)

to mean that given any starting context  $\rho_\alpha$  that satisfies  $P$ , running the command  $c$  is guaranteed to terminate in some context  $\rho_\omega$  that satisfies  $Q$ . Although we are not yet in a position to prove it, logics of total correctness are stronger than logics of partial correctness, in the sense that

$$\frac{[P] c [Q]}{\{P\} c \{Q\}} \text{Weaken}$$

A common but informal way to make the same point is that

$$\text{Partial Correctness} + \text{Termination} = \text{Total Correctness}$$

Except for the `While` rule, the axioms of a Hoare logic of total correctness are similar to the axioms of a Hoare logic of partial correctness; we give these rules in Figure 5. The difference is that we require that our subcommands terminate (*e.g.*, in `TSequence` we require  $[P] c_1 [Q]$  instead of  $\{P\} c_1 \{Q\}$ ), and produce a guarantee of termination (*e.g.*,  $[P] c_1; c_2 [R]$  instead of  $\{P\} c_1; c_2 \{R\}$ ).

When we get to developing a rule for the `while`  $(b) \{c\}$  command, things are a bit trickier. As we have observed, the previous rule does not guarantee termination, and merely knowing that the subcommand  $c$  terminates does not guarantee that the loop itself terminates. We will have to do something fancier.

A *termination measure*  $t$  is a function from contexts into the natural numbers. Suppose  $F$  is a formula in logic that has  $t$  as a free variable; then we can define an assertion  $t@F$  that evaluates  $F$  after substituting the value of the termination measure (on the current context) for  $t$  as follows:

$$\rho \Vdash t@P \equiv [t \mapsto t(\rho)]P$$

This may be a bit confusing, so let us consider an example. Suppose our example termination measure  $t_a$  is the absolute value of program variable  $a$ :

$$t_a(\rho) \equiv |\rho(a)|$$

Recall from earlier that in our example context  $\rho_a$ , variable  $a$  had value 3. Suppose our example formula is  $F_a \equiv t_a < 5$ . Then we can simplify as follows:

$$\rho_a \Vdash t_a @ F_a = [t_a \mapsto t_a(\rho_a)]F_a = [t_a \mapsto 3](t_a < 5) = 3 < 5$$

Since  $3 < 5$ , we conclude that on  $\rho_a$ ,  $t_a @ F_a$  holds.

Given this idea of termination measures, we are ready to define the Hoare rule of total correctness for the `while` command as follows:

$$\frac{\forall n ( [P \wedge [b] \wedge t@(t = n)] c [P \wedge t@(t < n)] )}{[P] \text{ while } (b) \{c\} [P \wedge \neg[b]]} \text{TWhile}$$

This may look scary but it is not that bad. To conclude that the loop terminates, we require a somewhat stronger precondition above the bar. We need to prove that for all  $n : \mathbb{N}$ , if we start in a context that satisfies  $P$ , and in which  $b$  has evaluated to true, and in which the termination measure  $t$  is equal to  $n$ , then  $c$  will terminate in a context that satisfies  $P$  **and in which the termination measure  $t$  has strictly decreased**. The termination argument is roughly as follows: we start in some context  $\rho$ , in which the termination measure  $t$  is equal to some  $n$ . If the loop is finished ( $b$  evaluates to false), then we are done; if not, then our premise guarantees that executing  $c$  will strictly decrease the termination measure, and so we loop around and try again. Since the naturals are bounded below by 0, we know that we cannot strictly decrease the termination measure forever, so at some point we must terminate.

Using the TWhile rule can be tricky since you need to develop termination measures that decrease over each loop body; this is a major reason that we are willing to use logics of partial correctness, since at least they provide some guarantees without being quite as hard to use. In fact, there are many programs for which no-one has been able to figure out a reasonable termination measure.

Still, sometimes it is a very nice to be able to conclude that a program terminates; in Figure 6 we give an example of using the TWhile rule to prove the total correctness of the Parity function using the termination measure  $t \equiv |y|$ . Notice how we introduce a fresh metavariable  $n$  in the body of the loop; the value of this metavariable is **constant** as we go around the loop so that we can verify that the termination measure has decreased at the end. We highlight the parts of the formulas that verify that the termination measure decreases.

## 5 Soundness of Hoare Logic

We have presented a series of commands, and two Hoare logics (one of partial and one of total correctness) that can be used to verify programs using those commands. Although we have tried to argue why the rules of those logics are reasonable, you may not be convinced that they are. This is a reasonable concern since it is quite easy to write down reasonable-looking Hoare rules that turn out to be unsound; for example, consider the following

$$\frac{[P \wedge [b]] c [P]}{[P] \text{ while } (b) \{c\} [P \wedge \neg[b]]} \text{BadTWhile}$$

$$\begin{array}{l}
[x \geq 0] \\
[[y \mapsto x](x = y \wedge y \geq 0)] \\
y := x; \\
[x = y \wedge y \geq 0] \\
[\exists n (x = y + (2 \times n) \wedge y \geq 0)] \\
\mathbf{while} (2 \leq y) \\
\{ \\
\quad [(\exists n (x = y + (2 \times n)) \wedge y \geq 0) \wedge [2 \leq y] \wedge \boxed{t @ (t = n)}] \\
\quad [\exists n (x = y + (2 \times n)) \wedge y \geq 2 \wedge \boxed{|y| = n}] \\
\quad [[y \mapsto y-2] (\exists n (x = y + (2 \times n)) \wedge y \geq 0 \wedge \boxed{y = n - 2})] \\
\quad y := y-2 \\
\quad [\exists n (x = y + (2 \times n)) \wedge y \geq 0 \wedge \boxed{y = n - 2}] \\
\quad [\exists n (x = y + (2 \times n)) \wedge y \geq 0 \wedge \boxed{|y| < n}] \\
\quad [\exists n (x = y + (2 \times n)) \wedge y \geq 0 \wedge \boxed{t @ (t < n)}] \\
\} \\
[[\exists n (x = y + (2 \times n)) \wedge y \geq 0] \wedge \neg [2 \leq y]] \\
[[\exists n (x = y + (2 \times n)) \wedge y \geq 0] \wedge (y < 2)] \\
[[ (y \Downarrow 0) \wedge \mathbf{even}(x) ] \vee [ (y \Downarrow 1) \wedge \mathbf{odd}(x) ]]
\end{array}$$

Figure 6: Totally correct verification for Parity using  $t = |y|$

This is a seemingly-reasonable attempt to directly lift the `While` rule into the realm of total correctness. Of course, as discussed above, it fails: just because a loop body always terminates does not mean that the loop itself always terminates. Just consider the program `while (true) {skip}`!

In fact, as soon as one starts making a serious attempt to generalize our Hoare logics to more realistic languages (*e.g.*, with memory, input/output, concurrency), one needs to start writing down Hoare rules to reason about those new features. Since those new features are very complicated, it rapidly becomes very hard to be confident in their associated Hoare rules.

What is needed is a semantic model, and associated soundness proof, of Hoare logic. That is, we must provide a formal definition to the Hoare triple (the semantic model) and then prove each of our Hoare rules using our definition (the soundness proof). This is our remaining task today.

**Operational semantics.** A precise definition of the Hoare triple depends on a precise definition for how our language operates—in other words, on our language’s *operational semantics*. To provide that operational semantics, we will define a three-place *step relation* of the form:

$$c \vdash \rho \rightsquigarrow \rho'$$

Here  $c$  is a command and  $\rho$  is a starting context; meanwhile,  $\rho'$  is an ending context—that is, the context that results from running  $c$ . Technically, the relation is defined as the least relation<sup>5</sup> satisfying the seven rules given below.

The `eSkip` rule formalizes the execution behavior of the `skip` command:

$$\frac{}{\text{skip} \vdash \rho \rightsquigarrow \rho} \text{eSkip}$$

That is, given a starting context  $\rho$ , evaluating the `skip` command does nothing and returns the original context  $\rho$ . More interesting is the `eAssign` rule:

$$\frac{}{x := e \vdash \rho \rightsquigarrow [x \mapsto \text{neval}(\rho, e)]\rho} \text{eAssign}$$

In words, executing the command  $x := e$  transforms the starting context  $\rho$  into a new context  $\rho'$  that is equal to  $\rho$  everywhere except at  $x$ , where  $\rho'$  now contains the result of evaluating  $e$ . Hopefully this matches your intuition for  $x := e$ .

We now turn to the commands that contain subcommands. The first is the `eSequence` rule, which specifies how two commands should be run in sequence:

$$\frac{c_1 \vdash \rho_1 \rightsquigarrow \rho_2 \quad c_2 \vdash \rho_2 \rightsquigarrow \rho_3}{c_1; c_2 \vdash \rho_1 \rightsquigarrow \rho_3} \text{eSequence}$$

To determine the result of  $c_1; c_2$ , you take the starting context  $\rho_1$  and run  $c_1$  until it results in the intermediate context  $\rho_2$ . Then, you take  $\rho_2$  and run  $c_2$  until it results in the final context  $\rho_3$ , which is the result of the entire sequence.

<sup>5</sup>*i.e.*, the relation is inductively defined.

We need two rules to define the behavior of `if (b) {c1} else {c2}`: one when  $b$  is true and the other when  $b$  is false. Here they are:

$$\frac{\text{beval}(\rho, b) = \top \quad c_1 \vdash \rho_1 \rightsquigarrow \rho_2}{\text{if } (b) \{c_1\} \text{ else } \{c_2\} \vdash \rho_1 \rightsquigarrow \rho_2} \text{elf1}$$

$$\frac{\text{beval}(\rho, b) = \perp \quad c_2 \vdash \rho_1 \rightsquigarrow \rho_2}{\text{if } (b) \{c_1\} \text{ else } \{c_2\} \vdash \rho_1 \rightsquigarrow \rho_2} \text{elf2}$$

The `elf1` rule says that if the boolean expression  $b$  evaluates to true ( $\top$ ) in the starting context  $\rho_1$ , then the result of the entire conditional command (`if`) is found by running the first command  $c_1$ . The `elf2` rule covers the other case, when  $b$  evaluates to false ( $\perp$ ); in this case we run the second command  $c_2$ .

Similarly, we need two rules to define the behavior of the `while` loop:

$$\frac{\text{beval}(\rho, b) = \perp}{\text{while } (b) \{c\} \vdash \rho \rightsquigarrow \rho} \text{eWhile1}$$

$$\frac{\text{beval}(\rho, b) = \top \quad c \vdash \rho_1 \rightsquigarrow \rho_2 \quad \text{while } (b) \{c\} \vdash \rho_2 \rightsquigarrow \rho_3}{\text{while } (b) \{c\} \vdash \rho_1 \rightsquigarrow \rho_3} \text{eWhile2}$$

The `eWhile1` rule is straightforward enough: if the boolean expression  $b$  evaluates to false, then the loop will end and return the starting context  $\rho$ . The `eWhile2` rule is a little more complicated, and has three premises. The first premise is quite simple:  $b$  must evaluate to true. The second premise says that we must then run the command  $c$  in the initial context  $\rho_1$  to get to an intermediate context  $\rho_2$ ; this is very similar to what we do in the `eSequence`, `elf1`, and `elf2` rules: “recursing” on a subcommand. The third premise is the particularly interesting one: here we are **not** recursing on a subcommand, since we take the intermediate context  $\rho_2$  and determine what final context  $\rho_3$  the original loop `while (b) {c}` will produce. Since the step relation is inductively defined, we will only be able to recurse some finite number of times—eventually we must hit the “base case” of `eWhile1`. **The step relation will only be defined for loops that terminate.** Consider the infinite loop “`while (true) {skip}`”; given any starting context  $\rho$ , there is no context  $\rho'$  such that

$$\text{while } (\text{true}) \{ \text{skip} \} \vdash \rho \rightsquigarrow \rho'$$

On the other hand, for terminating loops, the step relation computes the expected final context by “unrolling” the loop the required finite number of times.

All of this may seem a bit odd: essentially, we are formally saying that our programs have no semantics (that is, do not execute!) if they would run forever in the naïve intuition. Indeed, for more complicated programming languages, this style of step relation—called *big-step*—is inadequate. A big step operational semantics, like the one given above, computes the result of the entire program in one “big step”: that is, the step relation takes a command, before-state, and after-state. This runs into lots of problems when we care about what the

program does along the way, such as input and output (to/from *e.g.*, keyboard, mouse, screen, printer, network). To reason about those kinds of languages, we need a more complicated *small-step* semantics, where each time we apply the step relation we compute the result of some simple command; in a small-step semantics, the step relation takes an input command and context, as we do here, but outputs **both** a next command and a next context<sup>6</sup>. For example, a small-step rule for the **skip** command might look as follows:

$$(\rho, \mathbf{skip}; c) \rightsquigarrow (\rho, c)$$

If we ran the program for another step using the small-step semantics, it would start working on the command  $c$ . In a small-step semantics, infinite loops are easy to model and reason about since we can just keep on stepping forever.

Both partial and total Hoare triples only care about the final context if the program terminates. Therefore, it does not matter if we cannot model nonterminating programs with our semantics, and so for our simple language big-step semantics are adequate (and simpler than small-step).

**Exercise 3.** Prove that the big-step relation defined by the seven rules **eSkip**—**eWhile2** is deterministic. That is, if  $c \vdash \rho_1 \rightsquigarrow \rho_2$  and  $c \vdash \rho_1 \rightsquigarrow \rho'_2$ , then  $\rho'_1 = \rho'_2$ . Hint: do induction on the  $c \vdash \rho \rightsquigarrow \rho'$  relation.

**Exercise 4.** Define a small-step operational semantics for our language. On terminating programs, this language should result in the same final context as the big-step semantics. You may find it convenient to require that all programs have a final **skip** command to signal when the small-step relation can stop.

**Semantics of the Hoare triple.** We turn to the final topic in our study of Hoare logic, which is providing a model for the Hoare triples themselves: that is, what does  $\{P\} c \{Q\}$  mean? Actually, almost all of the heavy lifting has already been done. We begin by drawing on our connection to modal logic.

Our step relation, given a fixed command  $c$ , relates two worlds: the pre- and post-contexts of  $c$ . Thus, we define a relation  $S_c$  on contexts as follows:

$$\rho S_c \rho' \equiv c \vdash \rho \rightsquigarrow \rho'$$

Now that we have defined a relation on worlds, we can define the modal operators  $\Box$  and  $\Diamond$  (specialized to  $S$  and command  $c$ ) in exactly the usual way:

$$\rho \Vdash \Box_c^S P \equiv \forall \rho' (\rho S_c \rho' \rightarrow \rho' \Vdash P)$$

$$\rho \Vdash \Diamond_c^S P \equiv \exists \rho' (\rho S_c \rho' \wedge \rho' \Vdash P)$$

What do these operators mean? If  $\rho \Vdash \Box_c^S P$  holds, then  $P$  will hold in any final context reachable from  $\rho$  by running the command  $c$ —that is, if we run  $c$ ,  $P$

---

<sup>6</sup>Another option might be to define the big-step relation coinductively.

will hold. In contrast, if  $\rho \Vdash \diamond_c^S P$  holds, then **it is possible** to run  $c$  from the context  $\rho$ , and  $P$  will hold on the resulting context. Since our semantics do not allow us to run nonterminating programs, sometimes it is **not** possible to run  $c$ .

Now that we have expressed our modal operators, the definitions of the Hoare triples for partial and total correctness are easy to state:

$$\begin{aligned} \{P\} c \{Q\} &\equiv \forall \rho (\rho \Vdash (P \rightarrow \Box_c^S Q)) \\ [P] c [Q] &\equiv \forall \rho (\rho \Vdash (P \rightarrow \diamond_c^S Q)) \end{aligned}$$

Pleasingly, the definitions differ only by the modality they utilize.  $\{P\} c \{Q\}$  means that given any context  $\rho$  on which  $P$  holds,  $Q$  will hold if we execute  $c$ . In contrast,  $[P] c [Q]$  means that given any context  $\rho$ , if  $P$  holds on  $\rho$ , then we **can** execute  $c$ , and afterwards  $Q$  will hold. Since our semantics is only defined on terminating programs, a guarantee that  $c$  can run implies that  $c$  will terminate.

**Exercise 5.** Prove that total correctness is strictly stronger than partial correctness. That is, show that  $[P] c [Q]$  implies  $\{P\} c \{Q\}$ . Note that you will need the fact that our step relation is deterministic (see Exercise 3).

**Exercise 6. (★)** In the context of a nondeterministic semantics, our definition of Hoare triples for total correctness is not stronger than our definition for Hoare triples for partial correctness. Give a reasonable definition for Hoare triples of total correctness that is stronger in such a context.

The final task is to prove the various Hoare axioms given earlier (Skip, Assignment, Consequence, ..., While, TSkip, TAssignment, TConsequence, ..., TWhile) from the definitions. Most of these are quite simple.

**Exercise 7.** Prove all of the Hoare rules of partial correctness except While. You cannot just combine your proofs for exercises 5 and 8 since the **premises** of the rules of partial correctness allow nonterminating subcommands. You must prove the rules directly (but none of them should require a lots of work).

**Exercise 8.** Prove all of the Hoare rules of total correctness except TWhile. Again, none of them, individually, should require lots of work.

**Exercise 9. (★)** Prove the While rule of partial correctness.

**Exercise 10. (★)** Prove the TWhile rule of total correctness. Hint: do induction on the value of the termination measure.