

Lecture Notes on Semantics

Matt Fredrikson

Carnegie Mellon University
Lecture 5
September 8, 2025

1 Introduction

This lecture begins Part II: *From Informal to Formal Reasoning*. We need to formalize the various objects of study in the course so far: (1) a programming language, (2) a logic for reasoning about programs, and (3) a proof that the reasoning is correct in the sense that it is consistent with the meaning of programs. Once these aspects have been nailed down, formally, we have a basis for implementing them. We also understand them more thoroughly which means we will be able to use verification tools more effectively and even extend them to cover new computational phenomena.

There are many choices and tradeoffs for such a study, such as the extent of the features in the programming language, the expressive power of the logic, and the pragmatics of using the language and its logic for verification. A natural first idea would be to use WhyML, but it is too complex for us to study in the kind detail we wish to in this course. We conclude that the language should be small, but have the essential features that help us understand WhyML. Within that context, we should also decide whether to focus on functional or imperative aspects of WhyML. We have the opportunity to study functional programming and type theory in a number of other courses in the curriculum including 15-312 *Foundations of Programming Languages* and 15-417 *Constructive Logic*. In this course we are therefore taking a different perspective and study a *small imperative programming language*.

Even in the context of reasoning about imperative programs there are different traditions and approaches. Historically, there is *Hoare logic* that studies triples $P\{\alpha\}Q$ consisting of a precondition P , a program α , and a postcondition Q . This has been generalized to handle heap-allocated objects in *separation logic* and shared-memory concurrency in *concurrent separation logic*. We follow a different trajectory in choosing *dynamic logic*

inspired by traditional *modal logic*. It has also been generalized in multiple ways, including *differential dynamic logic* which supports reasoning about hybrid discrete and continuous evolving systems. Differential dynamic logic is at the core of 15-424 *Logical Foundations of Cyberphysical Systems*.

In outline, we will introduce a small imperative language and a language of formulas and then define the meaning (“semantics”) of both programs and formulas. This will answer the questions *How do programs execute?* and *When are formulas true?* We will ignore aspects of *concrete syntax* and work with *abstract syntax*, being unconcerned with how to parse or type-check programs. You can learn more about those aspects of programming languages in 15-312 (mentioned above) and 15-411 *Compiler Design*.

Learning goals. After this lecture, you should be able to:

- Simulate the dynamics of simple while programs
- Determine if programs are semantically equivalent
- Define the meaning of imperative language constructs
- Reason semantically about arithmetic formulas
- Specify semantics relationally

2 Straight-Line Programs

We now present our small imperative programming language in stages. The development is inherently open-ended in the sense that we will introduce more constructs as our study goes on.

For the sake of simplicity we assume that all variables range of the integers \mathbb{Z} . We have a simple language of *arithmetic expressions* e , with the usual conventions that we do not detail here. We use a, b, c for integer constants and x to stand for variables.

Arithmetic Expressions $e ::= c \mid x \mid e_1 + e_2 \mid e_1 - e_2 \mid e_1 * e_2 \mid \dots$

Since expressions contains variables, their meaning is determined with respect to a *state* that assigns integers to variables. We use ω, μ, ν to range over states and assume that they are defined on all variables. We write $\omega(x) = c$ if ω maps x to c . Then the value of expression e in state ω is written as

$$\omega[e] = c$$

and is easily defined based on the structure of e . For example:

$$\begin{aligned} \omega[c] &= c \\ \omega[x] &= \omega(x) \\ \omega[e_1 + e_2] &= \omega[e_1] + \omega[e_2] \\ \omega[e_1 - e_2] &= \omega[e_1] - \omega[e_2] \\ &\dots \end{aligned}$$

The last two equations may look somewhat odd—we have to keep in mind that ‘+’ and ‘-’ on the left-hand side are pieces of syntax that form expressions while ‘+’ and ‘-’ on the right-hand side are the mathematical operations on integers. Other operations are defined analogously.

Programs are denoted by α and β and we start here with two simple constructs: *assignment* $x \leftarrow e$ and *sequential composition* $\alpha ; \beta$.

$$\text{Programs } \alpha, \beta ::= x \leftarrow e \mid \alpha ; \beta \mid \dots$$

The meaning of a program is a *relation* between the *prestate* and *poststate* of its execution. It is a relation instead of a function because we would like to accommodate nonterminating programs (no possible poststate) and also nondeterministic programs (multiple possible poststates). We write

$$\omega \llbracket \alpha \rrbracket \nu$$

if the meaning of the program α relates prestate ω to poststate ν .

We define the meaning of assignment $x \leftarrow e$ to evaluate e in the current state to c and then update the state to map x to c . In symbols:

$$\omega \llbracket x \leftarrow e \rrbracket \nu \text{ iff } \nu = \omega[x \mapsto c] \text{ where } c = \omega \llbracket e \rrbracket$$

Here we use the notation $\omega[x \mapsto c]$ for the result of updating the state ω by mapping x to c (no matter what it was before).

The meaning of sequential composition $\alpha ; \beta$ is to execute first α and then β from the resulting state. That is:

$$\omega \llbracket \alpha ; \beta \rrbracket \nu \text{ iff there is a } \mu \text{ such that } \omega \llbracket \alpha \rrbracket \mu \text{ and } \mu \llbracket \beta \rrbracket \nu$$

In other words, the relation denoted by $\alpha ; \beta$ is the *composition* of the relations denoted by α and β .

As an example, let's compute

$$(\omega[x \mapsto a]) \llbracket x \leftarrow x + 2 \rrbracket \nu$$

and we find $\nu = (\omega[x \mapsto a])[x \mapsto a + 2] = \omega[x \mapsto a + 2]$ Slightly more complicated is

$$(\omega[x \mapsto a]) \llbracket x \leftarrow x + 1 ; x \leftarrow x + 1 \rrbracket \nu$$

We determine that there is an intermediate state $\mu = \omega[x \mapsto a + 1]$ and a final state $\nu = \omega[x \mapsto a + 2]$.

So, both of these programs define the same relation between $\omega[x \mapsto a]$ and $\omega[x \mapsto a + 2]$ Therefore we can state that these two programs are semantically equivalent

$$\llbracket x \leftarrow x + 2 \rrbracket = \llbracket x \leftarrow x + 1 ; x \leftarrow x + 1 \rrbracket$$

They have the same meaning because they have the same effects on the state. This, by the way, might fail to be true if the language were extended to allow shared memory

concurrency because another process can intervene after the first assignment on the right, while the left atomically increments x by two. Lesson: we always have to be careful about the extent of the language when we reason about it, be it semantically (as here) or logically (as in the next lecture).

As another example we consider this strange way to swap the values between two variables x and y without an auxiliary variable. We would like to prove:

$$(\omega[x \mapsto a, y \mapsto b]) \llbracket x \leftarrow x + y ; y \leftarrow x - y ; x \leftarrow x - y \rrbracket (\omega[x \mapsto b, y \mapsto a])$$

For this we have to calculate the intermediate states. Those are

$$\omega[x \mapsto a + b, y \mapsto b]$$

after the first assignment and

$$\omega[x \mapsto a + b, y \mapsto a]$$

after the second assignment, after which we reach the desired poststate.

3 Conditionals

We now add conditionals $\text{if } P \text{ then } \alpha \text{ else } \beta$ to our language, read as “*if* P *then* α *else* β ”. A characteristic of the dynamic logic approach is that formulas P do double duty: on one hand they serve as conditions in if-then-else programs and (shortly) guards on while loops. On the other hand we also use them to *reason* about programs as shown in the next lectures.

$$\begin{array}{ll} \text{Programs} & \alpha, \beta ::= x \leftarrow e \mid \alpha ; \beta \mid \text{if } P \text{ then } \alpha \text{ else } \beta \mid \dots \\ \text{Formulas} & P, Q ::= e_1 = e_2 \mid e_1 \leq e_2 \mid \top \mid \perp \mid P \wedge Q \mid P \vee Q \mid P \rightarrow Q \mid \neg P \\ & \mid \forall x. P \mid \exists x. P \mid \dots \end{array}$$

In concrete syntax we usually write \top (top) as “true”, \perp (bottom) as “false”, \neg as “not” and we write out the quantifiers as “forall” and “exists”.

In order to define the meaning of the conditional, we first need to define the meaning of the formulas, in mathematical terms. Because variables (and therefore quantifiers) range just over integers, the language of formulas we are concerned with is that of *integer arithmetic*. We define their meaning relative to an assignment ω of values to variables

$$\omega \models P \quad P \text{ is true in state } \omega$$

It is defined on the structure of P .

$\omega \models \top$	always
$\omega \models \perp$	never
$\omega \models e_1 = e_2$	iff $\omega[e_1] = \omega[e_2]$
$\omega \models e_1 \leq e_2$	iff $\omega[e_1] \leq \omega[e_2]$
$\omega \models P \wedge Q$	iff $\omega \models P$ and $\omega \models Q$
$\omega \models P \vee Q$	iff $\omega \models P$ or $\omega \models Q$
$\omega \models \neg P$	iff $\omega \not\models P$
$\omega \models P \rightarrow Q$	iff whenever $\omega \models P$ then also $\omega \models Q$
$\omega \models \forall x. P$	iff $\omega[x \mapsto a] \models P$ for all $a \in \mathbb{Z}$
$\omega \models \exists x. P$	iff $\omega[x \mapsto a] \models P$ for some $a \in \mathbb{Z}$

Because quantified integer arithmetic is undecidable, this definition is not effective in the sense that we cannot use it directly to determine whether a give formula is true. This is a problem if we want to actually execute our programs containing conditionals. So we usually restrict the formulas that can appear in conditionals to be quantifier-free, in which case it is easy to determine whether they are true or false.

With this out of the way, we can now define the meaning of the conditional by cases on the truth of P .

$$\omega[\text{if } P \alpha \beta] \nu \quad \text{iff} \quad \begin{cases} \omega[\alpha] \nu & \text{when } \omega \models P \\ \omega[\beta] \nu & \text{when } \omega \not\models P \end{cases}$$

4 While Loops

The abstract syntax for while loops is $\text{while } P \alpha$ which should somehow be the same as $\text{if } P (\alpha ; \text{while } P \alpha) \text{skip}$, where `skip` is a program that has no effect. Although it is perfectly possible to make this work as a so-called *inductive definition*, it has the issue that $\text{while } P \alpha$ appears on both sides. So we break it down by “guessing” the number of iterations of the loop, using an auxiliary relation $[\text{while } P \alpha]^n$ indexed by an $n \geq 0$. If $n = 0$ we must exit the loop so P should be false, and if $n > 0$ we should go around the loop once, followed by $n - 1$ more iterations.

$$\begin{aligned} \omega[\text{while } P \alpha] \nu &\quad \text{iff} \quad \text{there exists an } n \geq 0 \text{ such that } \omega[\text{while } P \alpha]^n \nu \\ \omega[\text{while } P \alpha]^0 \nu &\quad \text{iff} \quad \omega \not\models P \text{ and } \omega = \nu \\ \omega[\text{while } P \alpha]^{n+1} \nu &\quad \text{iff} \quad \omega \models P \text{ and there exists a } \mu \text{ such that } \omega[\alpha] \mu \text{ and } \mu[\text{while } P \alpha]^n \nu \end{aligned}$$

We can appeal to this definition to compute the meaning of a few simple programs. Actually, we will look at whole families of programs because it doesn’t matter what some of the components are. For example, any program $\text{while false } \alpha$ will behave the same, regardless of α . Instead of looking up the answer immediately, we suggest solving these yourself first with careful reference to the definitions.

$$\begin{aligned} \omega[\text{while true } \alpha] \nu \\ \omega[\text{while false } \alpha] \nu \\ \omega[x \leftarrow x] \nu \end{aligned}$$

We calculate

$$\begin{aligned}\omega[\text{while true } \alpha]_\nu & \text{ never} \\ \omega[\text{while false } \alpha]_\nu & \text{ iff } \nu = \omega \\ \omega[x \leftarrow x]_\nu & \text{ iff } \nu = \omega\end{aligned}$$

We see, for example, that

$$[\text{while false } \alpha] = [x \leftarrow x]$$

where the equality here denotes an equality between two relations. Further examples in the next section.

5 Tests

As the final language construct we consider the *test* or *guard* $?P$. Intuitively, it does nothing if P is true in the current state and “aborts” the computation if P is false. By “abort” we mean that there is no poststate, a semantics shared by a nonterminating while loop.

$$\omega[?P]_\nu \text{ iff } \omega \models P \text{ and } \omega = \nu$$

With this, we can define

$$\begin{aligned}\text{skip} & \triangleq ?\text{true} \quad \text{does nothing} \\ \text{abort} & \triangleq ?\text{false} \quad \text{aborts}\end{aligned}$$

These are *notational definitions* in the sense that the new program on the right expands to the program on the left. If we want to compute the semantics of the new kind of program we would just expand the definition and then compute the semantics of the result.

Tests can be used to model preconditions. A program such as

$$?(n \geq 0) ; \alpha$$

tests the condition $n \geq 0$ and proceeds with α if it is true. Therefore we can assume this condition while reasoning about the effect of α . If the condition is false then the computation aborts, so the final states only reflect the initial states that satisfy the test.

As an example, consider (once again) the following program to compute the $\text{fib}(n)$.

```
?(n ≥ 0) ;
i ← 0 ;
a ← 0 ;
b ← 1 ;
while (i < n)
  ( b ← b + a ;
    a ← b - a ;
    i ← i + 1 )
```

If we start in a state $\omega[n \mapsto c]$ then when we reach the while loop we have

$$\omega_0 = \omega[n \mapsto c, a \mapsto \text{fib}(0), b \mapsto \text{fib}(1), i \mapsto 0]$$

After iteration $k \leq c$, we have

$$\omega_k = \omega[n \mapsto c, a \mapsto \text{fib}(k), b \mapsto \text{fib}(k+1), i \mapsto k]$$

So the final state of the whole loop, which is also the final state of the program, has the form

$$\omega_c = \omega[n \mapsto c, a \mapsto \text{fib}(c), b \mapsto \text{fib}(c+1), i \mapsto c]$$

6 Side Note on Conditionals and Arithmetic

Goldbach's conjecture, proposed in 1742 and still open, states that every even natural number greater than 2 is the sum of two primes. We can actually express this quite easily in arithmetic.

$$\begin{aligned} \text{prime}(p) &\triangleq \neg \exists a. \exists b. a > 1 \wedge b > 1 \wedge a * b = p \\ \text{even}(a) &\triangleq \exists b. 2 * b = a \\ \text{goldbach} &\triangleq \forall n. n > 2 \wedge \text{even}(n) \rightarrow \exists p. \exists q. \text{prime}(p) \wedge \text{prime}(q) \wedge p + q = n \end{aligned}$$

Note that `goldbach` does not depend on any variables. According to our semantics it should therefore be either true or false. This in turns means that for

$$\omega \llbracket \text{if } \text{goldbach} (x \leftarrow 1) (x \leftarrow 0) \rrbracket \nu$$

we have $\nu(x) = 1$ if $\omega \models \text{goldbach}$ and $\nu(x) = 0$ if $\omega \not\models \text{goldbach}$. So you can appreciate the difficulty of trying to execute this program!

References