

Lecture Notes on Dynamic Semantics

15-411: Compiler Design
Frank Pfenning, Rob Simmons

Lecture 13
October 13, 2015

1 Introduction

In the previous lecture we specified the *static semantics* of a small imperative language. In this lecture we proceed to discuss its *dynamic semantics*, that is, how programs execute. The relationship between the dynamic semantics for a language and what a compiler actually implements is usually much less direct than for the static semantics. That's because a compiler doesn't actually run the program. Instead, it translates it from some source language to a target language, and then the program in the target language is actually executed.

In our context, the purpose of writing down the dynamic semantics is primarily to precisely specify how programs are supposed to execute. Previously, we defined how programs were supposed to execute either informally ("use left-to-right evaluation") or by explicitly specifying the way an elaborated AST is compiled into a lower-level language where effects are explicitly ordered. As our language becomes more complex, it will become increasingly important to make our language specification formal in order to avoid the ambiguity inherent in informal descriptions.

We *could* continue to define our dynamic semantics by compiling the language into a simpler form. (In fact, we are doing this to some degree, because we are defining our dynamic semantics in terms of the elaborated AST, which is a simplification of the concrete syntax. There is a trade-off here: the more work we do before defining the dynamic semantics, the harder it is to see the relationship between the original program and the program whose behavior we are defining.) Writing this down the dynamic semantics in terms of the source language, or something close to it, has a number of advantages. The programmer is saved from needing to understand the behavior of the compiler in order to understand the specification of the dynamic semantics. Conversely, the compiler writer is able to implement the

dynamic semantics in any way that conforms to the specification, rather than being tied to a specific implementation strategy they have specified.

Defining both our formal dynamic semantics and static semantics over the same elaborated AST also facilitates mathematically proving properties of the programming language. Much of the theory of programming languages is concerned with just that and therefore requires a dynamic semantics. Furthermore, if we define an operational specification of our language both before and after compilation, we can consider proving that they always compute the same result. These topics, however, are outside the scope of this course.

2 Evaluating Expressions

When trying to specify the operational semantics of a programming language, there are a bewildering array of choices regarding the *style* of presentation. Some choices are natural semantics, structural operational semantics, abstract machines, substructural operational semantics, and many more. We use the mechanism of *abstract machines*, despite some of its shortcomings.

In an *abstract machine semantics*, which is a form of so-called *small-step operational semantics*, we step through the evaluation of an expression e until we have reached a value v . So the basic judgment might be written $e \rightarrow e'$. However, this is much too simplistic. For example, it does not represent the call stack, or the current value of the variables that are recorded in an *environment*, or what to do with the eventual value. We will introduce such semantic artifacts one by one, as they are needed.

Consider the expression $e_1 + e_2$. By the left-to-right evaluation rule, we first have to evaluate e_1 and then e_2 . So when we evaluate e_1 we have to “remember” that we still have to evaluate e_2 then sum up the value. The information on what we still have to do is collected in a so-called *continuation* K . We write the judgment as

$$e \triangleright K$$

which we read as “evaluate expression e and pass the result to the continuation K ”. In the continuation there is a “hole” (written as an underscore character $_$) in which we plug in the value passed to it. So:

$$e_1 + e_2 \triangleright K \quad \longrightarrow \quad e_1 \triangleright (_ + e_2, K)$$

When e_1 has been reduced to a value c_1 , we plug it into the hole and evaluate e_2 next;

$$c_1 \triangleright (_ + e_2, K) \quad \longrightarrow \quad e_2 \triangleright (c_1 + _, K)$$

Finally, when e_2 has been reduced to a value c_2 we perform the actual addition and pass the result to K .

$$c_2 \triangleright (c_1 + _, K) \quad \longrightarrow \quad c \triangleright K \quad (c = c_1 + c_2 \bmod 2^{32})$$

In the last rule we appeal to the mathematical operation of addition modulo 2^{32} on two given integers modulo 2^{32} . Since we describe C0, we assume that constants are 32-bit words in two's complement representation. All other binary modular arithmetic operations \oplus are handled in a similar way, so we summarize them as

$$\begin{array}{lll} e_1 \oplus e_2 \triangleright K & \longrightarrow & e_1 \triangleright (_ \oplus e_2, K) \\ c_1 \triangleright (_ \oplus e_2, K) & \longrightarrow & e_2 \triangleright (c_1 \oplus _, K) \\ c_2 \triangleright (c_1 \oplus _, K) & \longrightarrow & c \triangleright K \quad (c = c_1 \oplus c_2 \bmod 2^{32}) \end{array}$$

For an effectful operation such as division, the last step could also raise an arithmetic exception arith (SIGFPE on the x86 architecture family, numbered 8). How do we represent that? We would like to abort the computation entirely and go to a state where the final outcome is reported as an arithmetic exception. We describe this as follows:

$$\begin{array}{lll} e_1 \odot e_2 \triangleright K & \longrightarrow & e_1 \triangleright (_ \odot e_2, K) \\ c_1 \triangleright (_ \odot e_2, K) & \longrightarrow & e_2 \triangleright (c_1 \odot _, K) \\ c_2 \triangleright (c_1 \odot _, K) & \longrightarrow & c \triangleright K \quad (c = c_1 \odot c_2) \\ c_2 \triangleright (c_1 \odot _, K) & \longrightarrow & \text{exception(arith)} \quad (c_1 \odot c_2 \text{ undefined}) \end{array}$$

Here, some care must be taken to define the value $c_1 \odot c_2$ correctly in the cases of division and modulus, and the conditions under which the result is mathematically "undefined" (like division by zero) and therefore must raise an exception. We have specified this in previous lectures and assignments, so we won't detail the conditions here.

What happens when evaluation finishes normally? In the case of the empty continuation we stop the abstract machine and return $\text{value}(c)$

$$c \triangleright \cdot \longrightarrow \text{value}(c)$$

Boolean expression and short-circuiting expressions work similarly.

$$\begin{array}{lll} e_1 \ \&\& \ e_2 \triangleright K & \longrightarrow & e_1 \triangleright (_ \ \&\& \ e_2, K) \\ \text{false} \triangleright (_ \ \&\& \ e_2, K) & \longrightarrow & \text{false} \triangleright K \\ \text{true} \triangleright (_ \ \&\& \ e_2, K) & \longrightarrow & e_2 \triangleright K \end{array}$$

Notice how e_2 is ignored in case false is returned to the continuation $(_ \ \&\& \ e_2, K)$, which encodes the short-circuiting behavior of the conjunction. Also note that we now have two kinds of values: boolean values false and true as well as integer values c . We could have also defined boolean expressions as evaluating to integers with a rule like $\text{true} \triangleright K \longrightarrow 1 \triangleright K$. By making true a value, we explicitly capture the fact that we expect integers and boolean expressions to not be confused at runtime.

3 Variables

We can continue along this line, but we get stuck for variables. Where do their values come from? We need to add an *environment* η that maps variables to their values. We write

$$\eta ::= \cdot \mid \eta, x \mapsto v$$

and $\eta[x \mapsto v]$ for either adding $x \mapsto v$ to η or overwriting the current value of x by v (if $\eta(x)$ is already defined). The state of the abstract machine now contains the environment η . We separate by a turnstile (\vdash) from the expression to evaluate and its continuation.

$$\eta \vdash e \triangleright K$$

The rules so far just carry this along. For example:

$$\begin{array}{lll} \eta \vdash e_1 \oplus e_2 \triangleright K & \longrightarrow & \eta \vdash e_1 \triangleright (_ \oplus e_2, K) \\ \eta \vdash c_1 \triangleright (_ \oplus e_2, K) & \longrightarrow & \eta \vdash e_2 \triangleright (c_1 \oplus _, K) \\ \eta \vdash c_2 \triangleright (c_1 \oplus _, K) & \longrightarrow & \eta \vdash c \triangleright K \quad (c = c_1 \oplus c_2 \bmod 2^{32}) \end{array}$$

Variables are just looked up in the environment.

$$\eta \vdash x \triangleright K \quad \longrightarrow \quad \eta \vdash \eta(x) \triangleright K$$

Because we are interested in evaluating only expressions that have already passed all static semantic checks of the language, we know that $\eta(x)$ will be defined (all variables must be initialized before they are used).

4 Executing Statements

Executing statements in L3, the fragment of C0 we have considered so far, can either complete normally, return from the current function with a return statement, or raise an exception. The “normal” execution of a statement does not pass a value to its continuation; instead it has an effect on its environment by assigning to its existing variables or declaring new ones. We write this as

$$\eta \vdash s \blacktriangleright K$$

where the continuation K should start with a statement. Statements don’t return values, so every statement that terminates eventually becomes a nop. For example:

$$\begin{array}{lll} \eta \vdash \text{seq}(s_1, s_2) \blacktriangleright K & \longrightarrow & \eta \vdash s_1 \blacktriangleright (s_2, K) \\ \eta \vdash \text{nop} \blacktriangleright (s, K) & \longrightarrow & \eta \vdash s \blacktriangleright K \end{array}$$

The last line codifies that if there is no further statement to execute, we grab the first statement from the continuation. When executing an assignment we first have

to evaluate the assignment, then change the variable value in the environment.

$$\begin{array}{ll} \eta \vdash \text{assign}(x, e) \blacktriangleright K & \longrightarrow \eta \vdash e \triangleright (\text{assign}(x, _), K) \\ \eta \vdash v \blacktriangleright (\text{assign}(x, _), K) & \longrightarrow \eta[x \mapsto v] \vdash \text{nop} \blacktriangleright K \end{array}$$

Conditionals follow the pattern of the short-circuiting conjunction.

$$\begin{array}{ll} \eta \vdash \text{if}(e, s_1, s_2) \blacktriangleright K & \longrightarrow \eta \vdash e \triangleright (\text{if}(_, s_1, s_2), K) \\ \eta \vdash \text{true} \triangleright (\text{if}(_, s_1, s_2)) & \longrightarrow \eta \vdash s_1 \blacktriangleright K \\ \eta \vdash \text{false} \triangleright (\text{if}(_, s_1, s_2)) & \longrightarrow \eta \vdash s_2 \blacktriangleright K \end{array}$$

While loops are a bit more complicated. We take a slight shortcut by using the identity

$$\text{while}(e, s) \equiv \text{if}(e, \text{seq}(s, \text{while}(e, s)), \text{nop})$$

to avoid writing out several rules implementing the right-hand side of this identity directly.

$$\eta \vdash \text{while}(e, s) \blacktriangleright K \quad \longrightarrow \quad \eta \vdash \text{if}(e, \text{seq}(s, \text{while}(e, s)), \text{nop}) \blacktriangleright K$$

Loops bring up the question of non-termination, which is modeled naturally: we just have abstract machine transitions $s_0 \longrightarrow s_1 \longrightarrow s_2 \longrightarrow \dots$ without ever arriving at a final state. The final states are just $\text{value}(c)$ and $\text{exception}(E)$, where E can currently be either arith or the abort exception caused by a failing assert statement.

$$\begin{array}{ll} \eta \vdash \text{assert}(e) \blacktriangleright K & \longrightarrow \eta \vdash e \triangleright (\text{assert}(_), K) \\ \eta \vdash \text{true} \triangleright (\text{assert}(_), K) & \longrightarrow \eta \vdash \text{nop} \blacktriangleright K \\ \eta \vdash \text{false} \triangleright (\text{assert}(_), K) & \longrightarrow \text{exception}(\text{abort}) \end{array}$$

Declarations are also pretty straightforward, since they just add a new variable with an undefined or useless value to the environment.

$$\eta \vdash \text{decl}(x, \tau, s) \blacktriangleright K \quad \longrightarrow \quad \eta[x \mapsto \text{nothing}] \vdash s \blacktriangleright K$$

In a language that permits shadowing of variables, we would have to save the current value of x and restore it after s has finished executing. Or we would think of η as a list where $\eta(x)$ refers to the value of the rightmost occurrence, which would be removed when we leave the scope of declaration. Or we could statically rename the variables during elaboration so that no shadowing can occur during execution.

At this point we have handled all salient statements except return statements that are tied to function calls. We discuss them in the next section.

5 Function Calls

A function call first has to evaluate the function arguments, from left to right. Then we invoke the function, whose body starts to execute in an environment that maps that formal parameters of the function to the argument values. But meanwhile we have to save the current environment of the caller somewhere. Similarly, we also have to save the continuation of the caller, so that when the callee returns we pass it the return value. So a stack frame $\langle \eta, K \rangle$ consists of an environment η and a continuation K .

$$\text{Call stack } S ::= \cdot | S, \langle \eta, K \rangle$$

Now states representing evaluation of expression and execution of statements have the forms

$$\begin{aligned} \text{Evaluation } S ; \eta \vdash e \triangleright K \\ \text{Execution } S ; \eta \vdash s \blacktriangleright K \end{aligned}$$

We only show the special case of evaluation function calls with two and zero arguments. This gets the point across while avoiding nitpicky details of working with multiple-argument functions.

$$\begin{aligned} S ; \eta \vdash f(e_1, e_2) \triangleright K &\longrightarrow S ; \eta \vdash e_1 \triangleright (f(_, e_2), K) \\ S ; \eta \vdash c_1 \triangleright (f(_, e_2), K) &\longrightarrow S ; \eta \vdash e_2 \triangleright (f(c_1, _), K) \\ S ; \eta \vdash c_2 \triangleright (f(c_1, _), K) &\longrightarrow (S, \langle \eta, K \rangle) ; [x_1 \mapsto c_1, x_2 \mapsto c_2] \vdash s \blacktriangleright \cdot \\ &\quad (\text{given that } f \text{ is defined as } f(x_1, x_2)\{s\}) \\ S ; \eta \vdash f() \triangleright K &\longrightarrow (S, \langle \eta, K \rangle) ; \cdot \vdash s \blacktriangleright \cdot \\ &\quad (\text{given that } f \text{ is defined as } f()\{s\}) \end{aligned}$$

In the next-to-last rule we see a new environment with values for x_1 and x_2 and a new stack frame save the caller's environment η and continuation K .

When executing a return statement we simply have to restore the caller's environment and continuation from the stack and pass the return value to the caller's continuation.

$$\begin{aligned} S ; \eta \vdash \text{return}(e) \blacktriangleright K &\longrightarrow S ; \eta \vdash e \triangleright (\text{return}(_), K) \\ S, \langle \eta', K' \rangle ; \eta \vdash v \triangleright (\text{return}(_), K) &\longrightarrow S ; \eta' \vdash v \triangleright K' \end{aligned}$$

In order to support functions returning void, we can use a useless value, nothing, and elaborate "return;" as `return(nothing)`. We also assume that elaboration adds an additional `return(nothing)` statement to the end of every void function. As an alternative, we could explicitly add a rule for function calls that finish without a return statement:

$$S, \langle \eta', K' \rangle ; \eta \vdash \text{nop} \blacktriangleright \cdot \longrightarrow S ; \eta' \vdash \text{nothing} \triangleright K'$$

6 Statics, Dynamics, and Safety

We start the machine initially in a state where we call the main function, and we stop the abstract machine if we reach this continuation.

$$\begin{array}{l} \cdot ; \cdot \vdash \text{main}() \triangleright \cdot \quad \text{(initial state)} \\ \cdot ; \eta \vdash c \triangleright \cdot \quad \longrightarrow \quad \text{value}(c) \quad \text{final state} \end{array}$$

which will eventually step to $\text{value}(c)$, where c is returned by the main function. We have defined four kinds of *machine state* ST :

- $S ; \eta \vdash e \triangleright K$ – Evaluating the expression e with the continuation K
- $S ; \eta \vdash s \blacktriangleright K$ – Evaluating the statement s with the continuation K
- $\text{value}(c)$ – Final state, return a value
- $\text{exception}(E)$ – Final state, report an error

We've also written a bunch of transition rules $ST \longrightarrow ST'$. Can we say anything about whether those rules are reasonable or not?

Obviously we don't expect any final state to be on the left-hand side of a transition. On that basis we could say that any rule like $\text{value}(3) \longrightarrow \text{value}(4)$ is an unreasonable rule. We might also care that the language is deterministic: that every machine state ST has at most one state ST' such that $ST \longrightarrow ST'$.

In the other direction, there are many examples of non-final states that have *no* transition: $S ; \eta \vdash 42 \triangleright (\text{if}(_, s_1, s_2); K)$ is one example, and $\cdot ; \cdot \vdash \text{nop} \blacktriangleright \cdot$ is another. Such states are called *stuck* – it is literally undefined what the program should do in such a state. The central relationship between the static semantics and dynamic semantics is that any program that passes the static semantics should be *free of undefined behavior*, that is, free of stuck states. This can be expressed mathematically:

Theorem 1 (No undefined behavior) *If a program passes all the static semantics, and*

$$\cdot ; \cdot \vdash \text{main}() \longrightarrow ST_1 \longrightarrow \dots \longrightarrow ST_n$$

then either ST_n is a final state or else ST_n is not-stuck because there exists a state ST' such that $ST_n \longrightarrow ST'$.

In a course like 15-312, we would learn how to prove these sorts of theorems, but just stating the theorem is still useful as a specification. If we can find a *counterexample*, a program that passes the static semantics and yet gets stuck in a non-final state according to the dynamic semantics, then we know that we need to change either or static or dynamic semantics.

7 Summary

We use \odot to stand for either a pure operation \oplus , or a potentially effectful operation \otimes as well as shift, comparison, and bitwise operators.

Expressions	e	$::=$	$c \mid e_1 \odot e_2 \mid \text{true} \mid \text{false} \mid e_1 \ \&\& \ e_2 \mid x \mid f(e_1, e_2) \mid f()$
Statements	s	$::=$	$\text{nop} \mid \text{seq}(s_1, s_2) \mid \text{assign}(x, e) \mid \text{decl}(x, \tau, s)$ $\mid \text{if}(e, s_1, s_2) \mid \text{while}(e, s) \mid \text{return}(e) \mid \text{assert}(e)$
Values	v	$::=$	$c \mid \text{true} \mid \text{false} \mid \text{nothing}$
Environments	η	$::=$	$\cdot \mid \eta, x \mapsto c$
Stacks	S	$::=$	$\cdot \mid S, \langle \eta, K \rangle$
Cont. frames	ϕ	$::=$	$_ \odot e \mid c \odot _ \mid _ \ \&\& \ e \mid f(_, e) \mid f(c, _)$ $\mid s \mid \text{assign}(x, _) \mid \text{if}(_, s_1, s_2) \mid \text{return}(_) \mid \text{assert}(_)$
Continuations	K	$::=$	$\cdot \mid \phi, K$
Exceptions	E	$::=$	$\text{arith} \mid \text{abort}$

$S ; \eta \vdash e_1 \odot e_2 \triangleright K$	\longrightarrow	$S ; \eta \vdash e_1 \triangleright (_ \odot e_2 , K)$
$S ; \eta \vdash c_1 \triangleright (_ \odot e_2 , K)$	\longrightarrow	$S ; \eta \vdash e_2 \triangleright (c_1 \odot _ , K)$
$S ; \eta \vdash c_2 \triangleright (c_1 \odot _ , K)$	\longrightarrow	$S ; \eta \vdash c \triangleright K \quad (c = c_1 \odot c_2)$
$S ; \eta \vdash c_2 \triangleright (c_1 \odot _ , K)$	\longrightarrow	exception(arith) $\quad (c_1 \odot c_2 \text{ undefined})$
$S ; \eta \vdash e_1 \ \&\& \ e_2 \triangleright K$	\longrightarrow	$S ; \eta \vdash e_1 \triangleright (_ \ \&\& \ e_2 , K)$
$S ; \eta \vdash \text{false} \triangleright (_ \ \&\& \ e_2 , K)$	\longrightarrow	$S ; \eta \vdash \text{false} \triangleright K$
$S ; \eta \vdash \text{true} \triangleright (_ \ \&\& \ e_2 , K)$	\longrightarrow	$S ; \eta \vdash e_2 \triangleright K$
$S ; \eta \vdash x \triangleright K$	\longrightarrow	$S ; \eta \vdash \eta(x) \triangleright K$
$S ; \eta \vdash \text{nop} \blacktriangleright (s , K)$	\longrightarrow	$S ; \eta \vdash s \blacktriangleright K$
$S ; \eta \vdash \text{assign}(x, e) \blacktriangleright K$	\longrightarrow	$S ; \eta \vdash e \triangleright (\text{assign}(x, _) , K)$
$S ; \eta \vdash c \triangleright (\text{assign}(x, _) , K)$	\longrightarrow	$S ; \eta[x \mapsto c] \vdash \text{nop} \blacktriangleright K$
$S ; \eta \vdash \text{decl}(x, \tau, s) \blacktriangleright K$	\longrightarrow	$S ; \eta[x \mapsto \text{nothing}] \vdash s \blacktriangleright K$
$S ; \eta \vdash \text{assert}(e) \blacktriangleright K$	\longrightarrow	$S ; \eta \vdash e \triangleright (\text{assert}(_) , K)$
$S ; \eta \vdash \text{true} \triangleright (\text{assert}(_) , K)$	\longrightarrow	$S ; \eta \vdash \text{nop} \blacktriangleright K$
$S ; \eta \vdash \text{false} \triangleright (\text{assert}(_) , K)$	\longrightarrow	exception(abort)
$S ; \eta \vdash \text{if}(e, s_1, s_2) \blacktriangleright K$	\longrightarrow	$S ; \eta \vdash e \triangleright (\text{if}(_ , s_1, s_2) , K)$
$S ; \eta \vdash \text{true} \triangleright (\text{if}(_ , s_1, s_2))$	\longrightarrow	$S ; \eta \vdash s_1 \blacktriangleright K$
$S ; \eta \vdash \text{false} \triangleright (\text{if}(_ , s_1, s_2))$	\longrightarrow	$S ; \eta \vdash s_2 \blacktriangleright K$
$S ; \eta \vdash \text{while}(e, s) \blacktriangleright K$	\longrightarrow	$S ; \eta \vdash \text{if}(e, \text{seq}(s, \text{while}(e, s)), \text{nop}) \blacktriangleright K$
$S ; \eta \vdash f(e_1, e_2) \triangleright K$	\longrightarrow	$S ; \eta \vdash e_1 \triangleright (f(_ , e_2) , K)$
$S ; \eta \vdash c_1 \triangleright (f(_ , e_2) , K)$	\longrightarrow	$S ; \eta \vdash e_2 \triangleright (f(c_1, _) , K)$
$S ; \eta \vdash c_2 \triangleright (f(c_1, _) , K)$	\longrightarrow	$(S , \langle \eta , K \rangle) ; [x_1 \mapsto c_1, x_2 \mapsto c_2] \vdash s \blacktriangleright \cdot$ (given that f is defined as $f(x_1, x_2)\{s\}$)
$S ; \eta \vdash f() \triangleright K$	\longrightarrow	$(S , \langle \eta , K \rangle) ; \cdot \vdash s \blacktriangleright \cdot$ (given that f is defined as $f()\{s\}$)
$S ; \eta \vdash \text{return}(e) \blacktriangleright K$	\longrightarrow	$S ; \eta \vdash e \triangleright (\text{return}(_) , K)$
$(S , \langle \eta' , K' \rangle) ; \eta \vdash v \triangleright (\text{return}(_) , K)$	\longrightarrow	$S ; \eta' \vdash v \triangleright K'$
$\cdot ; \eta \vdash c \triangleright \cdot$	\longrightarrow	value(c)