

# Abstract Machine Semantics

CPSC 509: Programming Language Principles

Ronald Garcia\*

12 November 2011

## Introduction

In these notes we will be talking about *control* structures in programming languages. To do so, we'll introduce yet another style of *operational semantics*, one that makes program control more apparent. We have already discussed *big-step semantics*, which take the form  $t \Downarrow v$ , and we've discussed two kinds of *small-step semantics*: *structural operational semantics* and *reduction semantics*. We'll now talk about a third kind of small-step semantics, called *abstract machine semantics*. Abstract machine semantics operationalize the intuition that we discussed regarding how to think about running a reduction semantic: given a program, you have to somehow find a way to decompose the program into an evaluation context and redex, then reduce the redex, and then plug the result back into the context, and then repeat these steps until a final outcome results. Abstract machines typically inject more systematic method into this process. As such, this semantic approach is a useful guide toward implementing compilers, which must translate your high-level language into low-level instructions that run on some machine with a bunch of registers, some stack memory, and some heap memory.<sup>1</sup>

## A Simple Language

Our running example for this introduction is, as usual, a language of Boolean expressions. The language is bare-bones, but has enough content to capture most of the concepts that concern us at this time. We'll start with the language reduction semantics, and then will return to present it as an abstract machine.

### Syntax

$$\begin{aligned} t &\in \text{TERM}, & v &\in \text{VAL}, & r &\in \text{REDEX}, & E &\in \text{ECTXT} \\ t &::= \text{true} \mid \text{false} \mid \text{if } t \text{ then } t \text{ else } t \mid t \text{ and } t \mid t \text{ or } t \mid t \text{ xor } t \mid \text{not } t \\ v &::= \text{true} \mid \text{false} \\ r &::= \text{if } v \text{ then } t \text{ else } t \mid v \text{ and } t \mid v \text{ or } t \mid v \text{ xor } v \mid \text{not } v \\ E &::= \square \mid E[\text{if } \square \text{ then } t \text{ else } t] \mid E[\square \text{ and } t] \mid E[\square \text{ or } t] \mid E[\square \text{ xor } t] \mid E[v \text{ xor } \square] \mid E[\text{not } \square] \end{aligned}$$

---

\*© Ronald Garcia. Not to be copied, used, or revised without explicit written permission from the copyright owner.

<sup>1</sup>Nowadays, most computers have just plain ole' memory, which you can use to implement your own stack and heap, but these two concepts are so ingrained in the brains of programmers that it's helpful to at least at first maintain the fiction that they are two truly distinct components of a computer.

**Notions of Reduction**  $\rightsquigarrow \subseteq \text{REDEX} \times \text{TERM}$ :

$\text{if true then } t_2 \text{ else } t_3 \rightsquigarrow t_2$   
 $\text{if false then } t_2 \text{ else } t_3 \rightsquigarrow t_3$   
 $\text{true and } t_2 \rightsquigarrow t_2$   
 $\text{false and } t_2 \rightsquigarrow \text{false}$   
 $\text{true or } t_2 \rightsquigarrow \text{true}$   
 $\text{false or } t_2 \rightsquigarrow t_2$   
 $v \text{ xor } v \rightsquigarrow \text{false}$   
 $\text{false xor true} \rightsquigarrow \text{true}$   
 $\text{true xor false} \rightsquigarrow \text{true}$   
 $\text{not true} \rightsquigarrow \text{false}$   
 $\text{not false} \rightsquigarrow \text{true}$

**One-Step Reduction**

$$\frac{r \rightsquigarrow t}{E[r] \longrightarrow E[t]}$$

As is common practice, We omit the definition of multi-step reduction. One-step reduction is often omitted too, but it plays a role in our story about abstract machines.

## Motivation for Abstract Machine Semantics

The reduction semantics style takes the *structural rules* of structural operational semantics, like the rule for **if**:

$$\frac{t_1 \longrightarrow t'_1}{\text{if } t_1 \text{ then } t_2 \text{ else } t_3 \longrightarrow \text{if } t'_1 \text{ then } t_2 \text{ else } t_3}$$

and expresses them succinctly as a grammar of *evaluation contexts*.

$$E ::= \dots \mid E[\text{if } \square \text{ then } t \text{ else } t] \dots$$

We can then view the process of evaluation as a succession of breaking a program down into an evaluation context and redex and applying a notion of reduction:

$$t_0 = E_0[r_0] \longrightarrow E_0[t'_0] = t_1 = E_1[r_1] \longrightarrow E_1[t'_1] = t_2 = E_2[r_2] \longrightarrow \dots$$

After each reduction step  $E_i[r_i] \longrightarrow E_i[t'_i] = t_{i+1}$ , the next step begins with asserting the existence of some different from-scratch decomposition of the program  $t_{i+1} = E_{i+1}[r_{i+1}]$  to find the next redex to reduce. The sufficiency of this multi-step evaluation process typically relies on the truth of a particular theorem.

**Theorem 1** (Unique Decomposition). *For any program  $t$ , either  $t$  is a value  $v$ , or there exists a unique context  $E$  and unique redex  $r$  such that  $t = E[r]$ .*

Proving this theorem is somewhat involved: when set up appropriately, it's tedious but routine. The existence clause of the theorem guarantees that evaluation is either finished (when  $t$  is a value), can make progress ( $t = E[r]$  for some context  $E$  and redex  $r$  that can be reduced according to some notion of reduction (if the semantics is defined for all well-formed programs).

The uniqueness clause of this theorem guarantees that the programming language is deterministic, because if there is an  $E$  and  $r$  such that  $t = E[r]$ , then both  $E$  and  $r$  are unique: there is only one possible step at this point in the program.

Reduction semantics are nice and abstract, but for some purposes, they can be seen as too abstract. Just knowing that unique decomposition holds isn't enough: to evaluate a program, we want not only to know that a decomposition exists: we want to *find* that unique decomposition.

A language implementation modeled after reduction semantics would have to do the work of decomposing and then making a reduction step. Naturally a sufficient proof of unique decomposition contains within it an algorithm to find unique context-redex pairs, but the most straightforward proof technique implies a rather inefficient (but informative) algorithm. In fact, simply solving the decomposition problem is not quite satisfactory. Both reduction semantics and structural operational semantics are woefully inefficient in that each step of program reduction rebuilds the entire program via plugging, and then decomposes it again (which manifests in the reduction example above as  $E_0[t'_1] = t_1 = E_1[r_1]$ ). Most of the time  $E_0$  and  $E_1$  share a great deal of the same structure. A real language implementation need not be so inefficient: it can work directly with  $E_0$  and  $t'_1$  to find  $E_1$  and  $r_1$ .

## Abstract Machine Semantics

We now show how a reduction semantics can be elaborated into a more realistic model of program execution, one that does not keep decomposing and rebuilding a program after each reduction step. An *abstract machine semantics* presents a high-level model of how a computer might actually go about running a program. Historically, abstract machines were the first kind of operational semantics, introduced by Peter Landin in the highly influential 1964 paper *The Mechanical Evaluation of Expressions* [Landin, 1964]. However, abstract machines were considered too low-level by many researchers, and this led to the development of structural operational semantics (by Plotkin) and reduction semantics (by Felleisen). However, it was observed that a reduction semantics is just an abstract representation of an abstract machine.

In this section we expand the reduction semantics above to form an abstract machine. This machine implements the deterministic search for a context-redex pair that can be reduced. Then, after a reduction, the machine proceeds directly from the current decomposition  $E_n[t'_n]$  to find the next redex, without completely plugging the term and starting over from scratch. This approach is much more efficient in practice.

The abstract machine is defined by a set of *configurations*, which for our simple language are comprised of evaluation context-term pairs. The machine operates in four different *modes*, which are reflected in a grammar of configurations:

$$\begin{aligned} C &\in \text{CFG} \\ C &::= \langle E, t \rangle_{\text{focus}} \mid \langle E, r \rangle_{\text{reduce}} \mid \langle E, v \rangle_{\text{return}} \mid v \end{aligned}$$

The *focus* mode represents when the abstract machine is searching downward into a term for a redex. The *reduce* mode represents when the abstract machine has found a redex and is ready to reduce it. The *return* mode represents when the machine has found or produced a value, and is returning that value to the current context in order to find the next piece of work to do. The final mode, a standalone value  $v$ , denotes that reduction has concluded with a final value.

The heart of the abstract machine semantics is the single step relation

$$\longrightarrow \subseteq \text{CFG} \times \text{CFG}$$

which captures how the machine configurations evolve over time. Since our language happens to be deterministic, we can consider  $\longrightarrow$  to be a partial function, but in general that is not the case, so we simply consider it to be a binary relation.

We break the step relation down based on which mode it is in. The *focus* steps determine how to transition by analyzing the top-level structure of the machine's term position.

$$\begin{aligned} \langle E, \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \rangle_{\text{focus}} &\longrightarrow \langle E[\text{if } \square \text{ then } t_2 \text{ else } t_3], t_1 \rangle_{\text{focus}} \\ \langle E, t_1 \text{ and } t_2 \rangle_{\text{focus}} &\longrightarrow \langle E[\square \text{ and } t_2], t_1 \rangle_{\text{focus}} \\ \langle E, t_1 \text{ or } t_2 \rangle_{\text{focus}} &\longrightarrow \langle E[\square \text{ or } t_2], t_1 \rangle_{\text{focus}} \\ \langle E, t_1 \text{ xor } t_2 \rangle_{\text{focus}} &\longrightarrow \langle E[\square \text{ xor } t_2], t_1 \rangle_{\text{focus}} \\ \langle E, \text{not } t \rangle_{\text{focus}} &\longrightarrow \langle E[\text{not } t], t \rangle_{\text{focus}} \\ \langle E, v \rangle_{\text{focus}} &\longrightarrow \langle E, v \rangle_{\text{return}} \end{aligned}$$

In essence, focus steps facilitate the search for a redex by analyzing the structure of a term  $t$  step-by-step, leaving a trail of breadcrumbs behind it as extensions to the evaluation context  $E$ . Upon discovering a value, it transitions to a *return* configuration that has the same context and term, which we now know to be a value.

The *return* steps determine how to transition by analyzing the top-level structure of the machine's context, which corresponds to the innermost frame, if any.

$$\begin{array}{lcl}
\langle \square, v \rangle_{\text{return}} & \longrightarrow & v \\
\langle E[\text{if } \square \text{ then } t \text{ else } t], v \rangle_{\text{return}} & \longrightarrow & \langle E, \text{if } v \text{ then } t \text{ else } t \rangle_{\text{reduce}} \\
\langle E[\square \text{ and } t], v \rangle_{\text{return}} & \longrightarrow & \langle E, v \text{ and } t \rangle_{\text{reduce}} \\
\langle E[\square \text{ or } t], v \rangle_{\text{return}} & \longrightarrow & \langle E, v \text{ and } t \rangle_{\text{reduce}} \\
\langle E[\square \text{ xor } t], v \rangle_{\text{return}} & \longrightarrow & \langle E[v \text{ xor } \square], t \rangle_{\text{focus}} \\
\langle E[v_1 \text{ xor } \square], v \rangle_{\text{return}} & \longrightarrow & \langle E, v_1 \text{ xor } v_2 \rangle_{\text{reduce}} \\
\langle E[\text{not } \square], v \rangle_{\text{return}} & \longrightarrow & \langle E, \text{not } v \rangle_{\text{reduce}}
\end{array}$$

Here is where using inside-out contexts helps us: they are naturally suited for easily examining or updating the innermost *frame* of the evaluation context. If reduction has arrived at a value, and the evaluation context is exhausted, then there is no more work to do: reduction concludes by producing that value. In the case of **and**, **or**, and **not**, a value in the first (or only) position justifies a reduction step, thanks to short-circuit evaluation, so the machine transitions to a *reduce* configuration with the relevant redex in place. In the case of **xor**, both arguments must be values in order to reduce it, but if the evaluation context is  $E[\square \text{ xor } t]$ , then the machine cannot yet be sure that the second argument is a value. As a result, the machine replaces the innermost frame with a new frame that records the discovered knowledge that the first argument is a value, and proceeds to explore the second argument using a focus configuration. However, if the return configuration's evaluation context is  $E[v \text{ xor } \square]$ , then both arguments to **xor** are values and reduction can now proceed. The two return steps for **xor** demonstrate how an evaluation context records information about the current program as the machine traverses it in search of a redex. It's worth observing that the return configuration itself records, albeit temporarily, the discovery of a value  $v$ , and preserves that information as it inspects the context.

The *reduce* steps determine how to transition by analyzing the relevant structure of the redex. They proceed by performing the appropriate notion of reduction from the reduction semantics.

$$\begin{array}{lcl}
\langle E, \text{if true then } t_1 \text{ else } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_1 \rangle_{\text{focus}} \\
\langle E, \text{if false then } t_1 \text{ else } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
\langle E, \text{true and } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
\langle E, \text{false and } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{focus}} \\
\langle E, \text{true or } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{focus}} \\
\langle E, \text{false or } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
\langle E, v \text{ xor } v \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{focus}} \\
\langle E, \text{false xor true} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{focus}} \\
\langle E, \text{true xor false} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{focus}} \\
\langle E, \text{not false} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{focus}} \\
\langle E, \text{not true} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{focus}}
\end{array}$$

Every focus and return machine transition has simply contributed to the search for a redex: plugging the context into the term of a configuration yields the same total program at the beginning and end of any focus and return transition. Not so for reduce, this is where the "real work" gets done.

Each reduce configuration transitions to a focus transition containing the same evaluation context and the outcome of the notion of reduction in term position. The rationale for this design is that the only new information in the outcome transition is the output of reduction, so focusing on that term suffices to find the *next* relevant redex. This design is precisely what avoids the behaviour of reduction semantics that causes it to discard all of the information about the evaluation context and begin decomposition anew.

We can still do a little better, however. Some of our reduction steps tell us enough about reduction such that we can retain even more information. For instance, consider again the transition:

$$\langle E, \text{true or } t_2 \rangle_{\text{reduce}} \longrightarrow \langle E, \text{true} \rangle_{\text{focus}}$$

Just looking at this transition, we can tell that the next step will be to transition to  $\langle E, \text{true} \rangle_{\text{return}}$ , because we know that the term position holds a value. As such, we can do better than transitioning to focus by fast-forwarding to the next inevitable step. Olivier Danvy calls these inevitable steps *corridor transitions* because the transitions are forced down a single possible path to some known configuration after some number of steps (here 1). As such, we are wise to update the currently “lossy” transitions to immediately transition to return configurations:

$$\begin{array}{lcl}
 \langle E, \text{if true then } t_1 \text{ else } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_1 \rangle_{\text{focus}} \\
 \langle E, \text{if false then } t_1 \text{ else } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
 \langle E, \text{true and } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
 \langle E, \text{false and } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{return}} \\
 \langle E, \text{true or } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{return}} \\
 \langle E, \text{false or } t_2 \rangle_{\text{reduce}} & \longrightarrow & \langle E, t_2 \rangle_{\text{focus}} \\
 \langle E, v \text{ xor } v \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{return}} \\
 \langle E, \text{false xor true} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{return}} \\
 \langle E, \text{true xor false} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{return}} \\
 \langle E, \text{not false} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{true} \rangle_{\text{return}} \\
 \langle E, \text{not true} \rangle_{\text{reduce}} & \longrightarrow & \langle E, \text{false} \rangle_{\text{return}}
 \end{array}$$

In some other languages, corridor transitions will lead the machine to immediately extend the evaluation context, or transition immediately back to reduction. The outcome depends on the properties of the language and its reduction semantics. The natural default, though, is to focus on the outcome of the notion of reduction.

Now, given the definition of  $\longrightarrow$ , we can specify the corresponding evaluator for this language. The initial state of the machine has an empty context, because we have not yet analyzed the program, and the entire program is in focus position, ( $\langle \square, t \rangle_{\text{focus}}$ ). The end state of the machine is represented by a final value  $v$ , which must have arisen from the machine is returning a final value to an empty evaluation context (i.e.,  $\langle \square, v \rangle_{\text{return}}$ ). Given these, the definition of the evaluator follows:

$$eval(t) = v \text{ iff } \langle \square, t \rangle_{\text{focus}} \longrightarrow^* v.$$

Other programming languages may have additional possible outcomes (e.g. error states), and the evaluator would account for them.

Notice how the evaluation context literally plays the role of the *runtime stack* that you probably learned about in an early programming class. The basic structure of any computation is that when the machine is in state  $\langle E, t \rangle_{\text{focus}}$ , it is about to analyze  $t$  all the way to some value  $v$ , which will be returned to the current stack as  $\langle E, v \rangle_{\text{return}}$ . In big-step semantics, we can see the part of the same phenomenon by following the big-step reduction tree upwards. An arithmetic summation rule of the form

$$\frac{t_1 \Downarrow n_1 \quad t_2 \Downarrow n_2}{t_1 + t_2 \Downarrow n_3} \text{ where } n_3 = n_1 + n_2$$

first evaluates  $t_1$  all the way down to  $n_1$ , “remembering” that the next thing to do is evaluate  $t_2$ , before summing the results. This is made much more explicit in the abstract machine semantics, since the evaluation context is explicitly carrying this memory. And this is why abstract machine semantics are a clear way of expressing programming language control-constructs. Language features that implement nontraditional control patterns can be expressed by manipulating the evaluation context. We’ll see examples of this in the near future.

## References

P. J. Landin. The mechanical evaluation of expressions. *Comput. J.*, 6(4):308–320, 1964. doi: 10.1093/comjnl/6.4.308. URL <https://doi.org/10.1093/comjnl/6.4.308>.