

Abstract Machines

Gabriele Keller, Liam O'Connor, Johannes Åman Pohjola*

October 16, 2025

1 Abstract Machines

Abstract machines are a theoretical model of a calculator, typically consisting of a set of states, including initial and final states, and a set of machine operations, which manipulate the state of the machine. They are an important concept in theoretical computer science, for example to specify the operational semantics of a programming language (as we did we MinHs and TinyImp), or computational and complexity theory. Probably one of the best known examples of an abstract machine is the Turing machine, which was designed by Alan Turing in 1936 as a means to tackle the Entscheidungsproblem (Decision Problem).

Abstract Machines are closely related to Virtual Machines, which are basically Abstract Machines with an implementation. Such virtual machines can be used to achieve portability of high-level programming languages (like the Java Virtual Machine, or the Common Language Framework for the .NET framework) or complete operating systems.

In this course, we use abstract machines as a means to study the operational semantics of programming languages. We started with two simple machines, which have a very small set of operations and states. We will gradually add more details to the machine and bring it closer to a machine we could actually use as a basis for an implementation. We do this in order to make languages easier to reason about - we use the higher level specification (for example, big step semantics) for reasoning about higher level properties such as safety, and use the lower level specification for reasoning about machine characteristics such as performance.

Technically, you have already seen two examples of abstract machines - in the structural operational semantics of a language (the “small step” semantics), and in the evaluation (“big step”) semantics. While these technically constitute abstract machines, they are perhaps *too* abstract for our needs.

We seek to specify the evaluation of a language in a more “low-level” way. Should we start from the evaluation semantics, or the small-step semantics? Notice that the evaluation semantics do not even specify the *order* in which terms should be evaluated. Seeing as nondeterminism doesn't exist in real computers, we must specify such an order. Hence we could say that the small-step semantics are “lower-level” than the big-step. Therefore, it makes sense for us to start with the small-step semantics. We call the small-step semantics of MinHS the *M Machine*.

2 The *C Machine*

One thing still left rather abstract in the *M Machine* is control flow, specifically, the notion of a *runtime stack*. When we want to evaluate a term, say `Plus (Plus (Num 2) (Num 3)) (Num 4)`, the *M Machine* rules tell us that we must first evaluate the inner `Plus (Num 2) (Num 3)` and then return it to the original expression. This is akin to *pushing* the greater expression context onto a stack,

*Just minor edits

evaluating the inner expression and then *popping* the context off the stack again, returning the evaluated expression to it.

We will introduce a new machine, the *C Machine*, that makes this stack explicit. In order to do this, we will need to formalise it, but we will get an informal intuition for it first.

2.1 *C Machine States*

In the *M Machine*, the state of the machine merely consists of the current expression to be evaluated - the notion of the stack is hidden in the deduction tree of the inference rules. In our *C Machine*, however, our state consists of three parts:

- The current expression to be evaluated
- A *stack* of expression *frames*
- A single binary flag that denotes whether the machine is currently *evaluating* an expression, or *returning* a value after evaluating.

To start with, we will introduce some new symbols to represent the stack. Formally:

$$\frac{}{\circ \text{ Stack}} \quad \frac{s \text{ Stack} \quad x \text{ Frame}}{x \triangleright s \text{ Stack}}$$

So, \circ represents the empty stack, and $x \triangleright s$ is a stack with a frame x on the top. But what is a frame? We define a frame as an expression with a hole in it, denoted by \square . For example:

$$\text{Plus } \square \text{ (Num 3)}$$

Is a **Plus** expression awaiting evaluation of its first argument. In this way, our frames are *suspended computations*, which are awaiting further information in order to be evaluated.

2.1.1 A sketch of a *C Machine*

Now that we have the stack, we can have a shot at representing the machine states. Suppose we want to evaluate our earlier example with the *C Machine*:

$$\text{Plus (Plus (Num 2) (Num 3)) Num 4}$$

To begin, we need to come up with an *initial state* for our expression. So, when we start, we have the empty stack (\circ), and the machine starts with the flag set to *evaluate* the expression (denoted by \succ).

$$\circ \succ \text{Plus (Plus (Num 2) (Num 3)) (Num 4)}$$

The rules in the *M Machine* state that in order to evaluate a **Plus** expression, first the first subexpression should be evaluated, then the second. Hence, in this case, our *C Machine* will mirror the behaviour of the *M Machine*; that is, it should evaluate **Plus (Num 2) (Num 3)** first. To achieve this, a stack frame is pushed for **Plus**, with a \square in place of the first subexpression, and the machine is set to evaluate the expression we just replaced:

$$\text{Plus } \square \text{ (Num 4)} \triangleright \circ \succ \text{Plus (Num 2) (Num 3)}$$

Now the machine has to evaluate another **plus**, so another stack frame is pushed:

$$\text{Plus } \square \text{ (Num 3)} \triangleright \text{Plus } \square \text{ (Num 4)} \triangleright \circ \succ \text{Num 2}$$

Now the machine simply has to evaluate a numeric literal. Seeing as no further evaluation need take place (the value associated with the expression can be inferred without further work), the machine switches to *return* (denoted by \prec) the value back into the awaiting stack frame:

$$\text{Plus } \square \text{ (Num 3)} \triangleright \text{Plus } \square \text{ (Num 4)} \triangleright \circ \prec 2$$

Having evaluated the first argument, the machine again suspends computation of the `plus` expression in order to evaluate the second subexpression, which proceeds similarly:

$$\begin{array}{l} \text{Plus } 2 \square \triangleright \text{Plus } \square (\text{Num } 4) \triangleright \circ \succ \text{Num } 3 \\ \text{Plus } 2 \square \triangleright \text{Plus } \square (\text{Num } 4) \triangleright \circ \prec 3 \end{array}$$

As the machine is *returning* the last value necessary for the `plus` frame, it *pops* the frame off the stack, performs a primitive addition operation, and returns the result 5:

$$\text{Plus } \square \text{Num } 4 \triangleright \circ \prec 5$$

The rest of the evaluation proceeds predictably:

$$\begin{array}{l} \text{Plus } 5 \square \triangleright \circ \succ \text{Num } 4 \\ \text{Plus } 5 \square \triangleright \circ \prec 4 \\ \quad \circ \prec 9 \end{array}$$

2.2 Formalising the *C Machine*

Now that we have an informal idea of what our *C Machine* looks like, we can begin to formalise the machine. An abstract machine in general consists of:

- A set of states Σ .
- A set of initial states $I \subseteq \Sigma$.
- A set of final states $F \subseteq \Sigma$.
- A *state transition function*, $\mapsto : \Sigma \rightarrow \Sigma$.

You've seen these before—they're also called *small-step semantics*.

We can specify formally each of these components. The set of states Σ is specified as follows:

$$\frac{s \text{ Stack } \quad e \text{ Expr}}{s \succ e \in \Sigma} \quad \frac{s \text{ Stack } \quad v \text{ Value}}{s \prec v \in \Sigma}$$

That is, Σ is comprised of all evaluating states and all returning states.

The set of initial states I is defined as the set of all evaluating states with an empty stack:

$$\frac{e \text{ Expr}}{\circ \succ e \in I}$$

The final states F are defined as all returning states with an empty stack:

$$\frac{v \text{ Value}}{\circ \prec v \in F}$$

Now we must define the state transition relation for our *C Machine*, \mapsto_c .

2.2.1 Literals

To begin, we will start on the easy part - Evaluating numeric and boolean literals:

$$\frac{}{s \succ \text{Num } v \mapsto_c s \prec v} \quad \frac{}{s \succ \text{Bool } v \mapsto_c s \prec v}$$

So, the machine simply returns the corresponding values unchanged - no further computation is necessary. For function values, we must introduce some new notation, but the principle is the same as for numbers and booleans:

$$\frac{}{s \succ \text{Fun } (f.x \dots) \mapsto_c s \prec \langle\langle f.x \dots \rangle\rangle}$$

2.2.2 Primitive Operations

Now we can specify more complicated rules, such as that for **Plus**. When faced with an unevaluated **Plus** expression, the machine first evaluates the first subexpression, and pushes the rest on the stack:

$$\frac{}{s \succ \mathbf{Plus} \ e_1 \ e_2 \mapsto_c \mathbf{Plus} \ \square \ e_2 \triangleright s \succ e_1}$$

Once that subexpression is evaluated, the machine will begin evaluating the second subexpression:

$$\frac{}{\mathbf{Plus} \ \square \ e_2 \triangleright s \prec v \mapsto_c \mathbf{Plus} \ v \ \square \triangleright s \succ e_2}$$

Finally, when both subexpressions are evaluated, the machine returns the resultant sum, computed via a primitive operation, **+**:

$$\frac{}{\mathbf{Plus} \ v_1 \ \square \triangleright s \prec v_2 \mapsto_c s \prec (v_1+v_2)}$$

The definitions for **Times**, **Greater**, **Equal**, **Not**, **And** etc. are very similar, just making use of different primitive machine operations.

2.2.3 Conditionals

Now, what about **if**? Recall that in the *M Machine*, the machine first evaluates the condition to some result, and, depending on the result, would evaluate just one branch of the conditional.

Similarly, when faced with an unevaluated **If** expression, the *C Machine* evaluates the condition first:

$$\frac{}{s \succ \mathbf{If} \ c \ t \ e \mapsto_c \mathbf{If} \ \square \ t \ e \triangleright s \succ c}$$

If the result is **True**, the first branch is evaluated, otherwise the second:

$$\frac{}{\mathbf{If} \ \square \ t \ e \triangleright s \prec \mathbf{True} \mapsto_c s \succ t} \quad \frac{}{\mathbf{If} \ \square \ t \ e \triangleright s \prec \mathbf{False} \mapsto_c s \succ e}$$

2.2.4 Function Application

Finally, we must deal with function application. Recall that in the *M Machine*, first the function expression was evaluated, then the argument to the function, then finally the body of the function, substituting in the value of the argument. We employ a similar strategy here.

First evaluate the function value:

$$\frac{}{s \succ \mathbf{Apply} \ f \ a \mapsto_c \mathbf{Apply} \ \square \ a \triangleright s \succ f}$$

Then, evaluate the argument:

$$\frac{}{\mathbf{Apply} \ \square \ a \triangleright s \prec \langle\langle f.x. e \rangle\rangle \mapsto_c \mathbf{Apply} \ \langle\langle f.x. e \rangle\rangle \ \square \triangleright s \succ a}$$

Then finally evaluate the function body, substituting the value for the argument, and the function name in case the function is recursive:

$$\frac{}{\mathbf{Apply} \ \langle\langle f.x. e \rangle\rangle \ \square \triangleright s \prec v_a \mapsto_c s \succ e[f := (\mathbf{Fun} \ f.x. e), x = v_a]}$$

2.3 Example

Here we have a simple function that determines if the provided argument is even, applied to the argument 3. To make things shorter, we rename the **Num** abstract syntax expression to simply **N**, **Minus** to **Sub**, and **Apply** to **Ap**.

		o	>					
\mapsto_c	Ap	\square	$(N\ 3)$	\triangleright	o	\triangleright	Ap (Fun $f.x.(If\ (LEq\ x\ (N\ 0))\ (Eq\ x\ (N\ 0))\ (Ap\ f\ (Sub\ x\ (N\ 2))))\ (N\ 3)$)	
\mapsto_c	Ap	\square	$(N\ 3)$	\triangleright	o	\triangleright	Fun $f.x.(If\ (LEq\ x\ (N\ 0))\ (Eq\ x\ (N\ 0))\ (Ap\ f\ (Sub\ x\ (N\ 2))))$	
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	$\langle\langle f.x.(If\ (LEq\ x\ (N\ 0))\ (Eq\ x\ (N\ 0))\ (Ap\ f\ (Sub\ x\ (N\ 2))))\rangle\rangle$	
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	N 3	
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	3	
\mapsto_c					o	\triangleright	If (LEq (N 3) (N 0)) (Eq (N 3) (N 0)) (Ap (Fun \cdots) (Sub (N 3) (N 2)))	
\mapsto_c		If	\square	\cdots	\triangleright	o	\triangleright	LEq (N 3) (N 0)
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	N 3
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	3
\mapsto_c	LEq	3	\square	\triangleright	If	\cdots	\triangleright	N 0
\mapsto_c	LEq	3	\square	\triangleright	If	\cdots	\triangleright	0
\mapsto_c	If	\square	\cdots	(Ap \cdots)	\triangleright	o	\triangleright	False
\mapsto_c						o	\triangleright	(Ap (Fun \cdots) (Sub (N 3) (N 2)))
\mapsto_c	Ap	\square	(Sub \cdots)	\triangleright	o	\triangleright	(Fun \cdots)	
\mapsto_c	Ap	\square	(Sub \cdots)	\triangleright	o	\triangleright	$\langle\langle\cdots\rangle\rangle$	
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	Sub (N 3) (N 2)	
\mapsto_c	Sub	\square	$(N\ 2)$	\triangleright	Ap	\cdots	\triangleright	N 3
\mapsto_c	Sub	\square	$(N\ 2)$	\triangleright	Ap	\cdots	\triangleright	3
\mapsto_c	Sub	3	\square	\triangleright	Ap	\cdots	\triangleright	N 2
\mapsto_c	Sub	3	\square	\triangleright	Ap	\cdots	\triangleright	2
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	1	
\mapsto_c					o	\triangleright	If (LEq (N 1) (N 0)) (Eq (N 1) (N 0)) (Ap (Fun \cdots) (Sub (N 1) (N 2)))	
\mapsto_c						o	\triangleright	LEq (N 1) (N 0)
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	N 1
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	1
\mapsto_c	LEq	1	\square	\triangleright	If	\cdots	\triangleright	N 0
\mapsto_c	LEq	1	\square	\triangleright	If	\cdots	\triangleright	0
\mapsto_c	If	\square	\cdots	(Ap \cdots)	\triangleright	o	\triangleright	False
\mapsto_c						o	\triangleright	(Ap (Fun \cdots) (Sub (N 1) (N 2)))
\mapsto_c	Ap	\square	(Sub \cdots)	\triangleright	o	\triangleright	(Fun \cdots)	
\mapsto_c	Ap	\square	(Sub \cdots)	\triangleright	o	\triangleright	$\langle\langle\cdots\rangle\rangle$	
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	Sub (N 1) (N 2)	
\mapsto_c	Sub	\square	$(N\ 2)$	\triangleright	Ap	\cdots	\triangleright	N 1
\mapsto_c	Sub	\square	$(N\ 2)$	\triangleright	Ap	\cdots	\triangleright	1
\mapsto_c	Sub	1	\square	\triangleright	Ap	\cdots	\triangleright	N 2
\mapsto_c	Sub	1	\square	\triangleright	Ap	\cdots	\triangleright	2
\mapsto_c	Ap	$\langle\langle\cdots\rangle\rangle$	\square	\triangleright	o	\triangleright	-1	
\mapsto_c					o	\triangleright	If (LEq (N -1) (N 0)) (Eq (N -1) (N 0)) (Ap (Fun \cdots) (Sub (N -1) (N 2)))	
\mapsto_c						o	\triangleright	LEq (N -1) (N 0)
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	N -1
\mapsto_c	LEq	\square	$(N\ 0)$	\triangleright	If	\cdots	\triangleright	-1
\mapsto_c	LEq	-1	\square	\triangleright	If	\cdots	\triangleright	N 0
\mapsto_c	LEq	-1	\square	\triangleright	If	\cdots	\triangleright	0
\mapsto_c	If	\square	(Eq (N -1) (N 0))	\cdots	\triangleright	o	\triangleright	True
\mapsto_c						o	\triangleright	Eq (N -1) (N 0)
\mapsto_c	Eq	\square	$(N\ 0)$	\triangleright	o	\triangleright	N -1	
\mapsto_c	Eq	\square	$(N\ 0)$	\triangleright	o	\triangleright	-1	
\mapsto_c	Eq	-1	\square	\triangleright	o	\triangleright	N 0	
\mapsto_c	Eq	-1	\square	\triangleright	o	\triangleright	0	
\mapsto_c					o	\triangleright	False	

Wow, that long just to compute if three is even! No wonder we prefer evaluation semantics! Computers, however, certainly would prefer the *C Machine* - note that every state transition for the *C Machine* is an axiom. This means we can implement it as a single tight **while** loop that moves from state to state until it reaches a state in F .

Note: In an exam situation, you may be asked to present a derivation like the above. It is not necessary to write out every single step, just those steps you believe to be most important.

3 The *E Machine*

Now that we've made control flow more explicit, it becomes easier to see how we would implement the language efficiently on a real computer.

Let's take a look at our primitive machine operations so far:

- **The Numeric Operators** - +, * etc.
- **Comparison Operators** - ==, < etc.
- **Logical Operators** - &&, ||, !
- **Substitution** - $e[x := y]$

The great thing about most of these is that they are often native machine instructions in most computers, so they can be implemented very efficiently.

The one operation that *cannot* be implemented very efficiently is substitution. Substitution is of complexity $O(n)$ in the size of the expression—it would be very difficult to implement it as an efficient machine instruction!

So, we are going to extend our machine once more, to include *environments* in the machine state. We will call our new machine the *E Machine*. This way, we make the machine look up variables in the environment rather than rely on substitution.

We extend our states as follows:

$$\frac{s \text{ Stack} \quad \Gamma \text{ Env} \quad e \text{ Expr}}{s \mid \Gamma \succ e \in \Sigma} \qquad \frac{s \text{ Stack} \quad \Gamma \text{ Env} \quad e \text{ Expr}}{s \mid \Gamma \prec e \in \Sigma}$$

Note this is exactly the same as the *C Machine*, except with environments added. How do we denote environments?

$$\bullet \text{ Env} \qquad \frac{\Gamma \text{ Env} \quad x \text{ Ident} \quad v \text{ Value}}{x = v; \Gamma \text{ Env}}$$

So, \bullet is the empty environment, and we add new bindings to the environment with the notation $x = v$.

Now we need to add some rules for dealing with variables. In the *C Machine*, a variable occurring by itself was a stuck state - the only way for such a situation to arise is if the variable is free, which makes it an invalid expression. In the *M Machine*, variables are to be expected - *if* they occur in the environment:

$$\frac{}{s \mid x = v; \Gamma \succ x \mapsto_E s \mid x = v; \Gamma \prec v}$$

Now, what happens when we call a function? Naturally, we'd want to introduce new bindings to our environment, for the argument and the recursive name. When the function returns, however, we want to *remove* these bindings, as they are no longer in scope. The way we achieve this is somewhat unusual. We extend our stack to be able to include *environments* as well as frames:

$$\frac{}{\circ \text{ Stack}} \qquad \frac{s \text{ Stack} \quad x \text{ Frame}}{x \triangleright s \text{ Stack}} \qquad \frac{s \text{ Stack} \quad \Gamma \text{ Env}}{\Gamma \triangleright s \text{ Stack}}$$

Then, when we apply a function, we add the bindings, and push the *old* environment to the stack:

$$\text{Apply } \langle\langle f.x.e \rangle\rangle \quad \square \triangleright s \mid \Gamma \prec v \mapsto_E \Gamma \triangleright s \mid x = v; f = \langle\langle f.x.e \rangle\rangle; \Gamma \succ e$$

When the function returns, we *restore* the environment from the stack, which has the effect of removing the unwanted bindings:

$$\frac{}{\Gamma \triangleright s \mid \Delta \prec v \mapsto_E s \mid \Gamma \prec v}$$

			○ ●	↘	App (App (Fun $f.x.(Fun\ g.y.x)$) 3) 4	
↳ _E		App □ 4▷○	●	↘	App (Fun $f.x.(Fun\ g.y.x)$) 3	
↳ _E		App □ 3▷App □ 4▷○	●	↘	Fun $f.x.(Fun\ g.y.x)$	
				⋮		
↳ _E	App	⟨⟨●, $f.x.(Fun\ g.y.x)$ ⟩⟩ □▷App □ 4▷○	●	↘	3	
↳ _E		●▷App □ 4▷○	$x = 3; f = \langle\langle\cdot\cdot\rangle\rangle;$	●	↘	Fun $g.y.x$
↳ _E		●▷App □ 4▷○	$x = 3; f = \langle\langle\cdot\cdot\rangle\rangle;$	●	↘	⟨⟨ $x = 3; f; \bullet, g.y.x$ ⟩⟩
↳ _E			App □ 4▷○	●	↘	⟨⟨ $x = 3; f; \bullet, g.y.x$ ⟩⟩
				⋮		
↳ _E		App	⟨⟨ $x = 3; f; \bullet, g.y.x$ ⟩⟩ □▷○	●	↘	4
↳ _E		●▷○	$y = 4; g = \langle\langle\cdot\cdot\rangle\rangle; x = 3; f = \langle\langle\cdot\cdot\rangle\rangle;$	●	↘	x
↳ _E		●▷○	$y = 4; g = \langle\langle\cdot\cdot\rangle\rangle; x = 3; f = \langle\langle\cdot\cdot\rangle\rangle;$	●	↘	3
↳ _E			○ ●	↘	3	
				⊙		