

6. Automata-Theoretic Models of Higher-Order Computation

based on "Recursive Schemes, Krivine Machines, and Collapsible Pushdown Automata",

Sylvain Salvati & Igor Walukiewicz, 2013.

Goal: Introduce Krivine machines and collapsible pushdown automata, operational models capturing higher-order computation.

Idea: Krivine machines \rightarrow Operational understanding of reduction
Collapsible pushdown automata \rightarrow Pushdown view to functional programming.

6.1 Krivine Machines

Idea: Compute head normal form.

Technically: Use explicit substitutions called environments.

Environments assign closures to free variables,

very much like valuations σ in λ MF or in logic.

Closures are pairs consisting of a λ -term and again an environment.

Definition:

• Closures C and environments ρ are defined by mutual recursion:

$$C ::= (M, \rho)$$

$$\rho ::= \emptyset \mid \rho[x \mapsto C].$$

Here, M is a λ -term, \emptyset stands for the empty environment, and $\rho[x \mapsto C]$ coincides with ρ on all variables except x , which it maps to C .

- We will require that in a closure (M, ρ) , the environment ρ is defined for every free variable of M . Intuitively, the closure denotes the λ -term obtained by substituting in M every free variable x by the λ -term denoted by the closure $\rho(x)$. // Recursion.

For example,

$$(\lambda y. xy, [\rho \mapsto (a z, [z \mapsto b])])$$

denotes

$$\lambda y. (a b) y.$$

- Krivine machines are interpreters for λ -terms. As such, they do not have a syntax. Their state at runtime (configuration) is solely determined by the input.

Definition:

- A configuration of a Krivine machine is a triple (M, ρ, S) ,

where

- M is a λ -term,
- ρ is an environment, and
- $S \in C^*$ is a stack of closures, with the topmost stack element written to the left.

Computations of Krivine machines are defined by means of a transition relation among configurations. It is defined by the following rules:

- (1) $(\lambda x. M, \rho, (N, \rho'), S) \rightarrow (M, \rho[x \mapsto (N, \rho')], S)$
- (2) $(MN, \rho, S) \rightarrow (M, \rho, (N, \rho'), S)$
- (3) $(Y \lambda x. M, \rho, S) \rightarrow (M, \rho[x \mapsto (Y \lambda x. M, \rho)], S)$
- (4) $(x, \rho, S) \rightarrow (M, \rho', S), \text{ where } \rho(x) = (M, \rho').$

Note:

- The machine is deterministic.
- It is actually nonsense to refer to Krivine machines in plural, there is precisely one.

Intuition:

Rule (1): To evaluate $\lambda x. M$, look for the argument at the top of the stack and bind it to x . Then calculate the value of M .

Rule (2): To evaluate an application MN , put N onto the stack together with the current environment. This allows us to evaluate N when needed. Continue with the evaluation of M .

Rule (3): This combines the previous two rules.

Rule (4): Take the value of the varidsk from the environment and evaluate it.

The value is not just a term but also an environment, giving the right meaning of free variables in a term.

Example:

Consider $(\lambda x y z. x y z) a b c$

which (with more brackets) is

$((((\lambda x. \lambda y. \lambda z. (x y) z) a) b) c)$.

Here,

$x, a: \sigma \rightarrow \sigma \rightarrow \sigma$

$y, z, b, c: \sigma$.

The Krivine machine has the following transitions:

$((((\lambda x. \lambda y. \lambda z. (x y) z) a) b) c, \emptyset, \epsilon)$

$\xrightarrow{(2)} (((\lambda x. \lambda y. \lambda z. (x y) z) a) b, \emptyset, (c, \emptyset))$

$\xrightarrow{(2)} (\lambda x. \lambda y. \lambda z. (x y) z) a, \emptyset, (b, \emptyset). (c, \emptyset)$

$\xrightarrow{(2)} (\lambda x. \lambda y. \lambda z. (x y) z, \emptyset, (a, \emptyset). (b, \emptyset). (c, \emptyset))$

$\xrightarrow{(1)} (\lambda y. \lambda z. (x y) z, [x \mapsto (a, \emptyset)], (b, \emptyset). (c, \emptyset))$

$\xrightarrow{(1)} (\lambda z. (x y) z, [x \mapsto (a, \emptyset), y \mapsto (b, \emptyset)], (c, \emptyset))$

$\xrightarrow{(1)} ((x y) z, \underbrace{[x \mapsto (a, \emptyset), y \mapsto (b, \emptyset), z \mapsto (c, \emptyset)]}_{=: \mathcal{P}}, \epsilon)$

$=: \mathcal{P}$

$$\xrightarrow{(2)} (x y, \rho, (z, \rho))$$

$$\xrightarrow{(2)} (x, \rho, (y, \rho), (z, \rho))$$

$$\xrightarrow{(4)} (a, \emptyset, (y, \rho), (z, \rho)).$$

There are no more transitions left.

Indeed, the head normal form is

$$(a b) c.$$

We will be interested in configurations reachable from $(M, \emptyset, \varepsilon)$.

Every such configuration satisfies strong invariants, summarized by the next lemma.

Lemma:

Let M be a λ -term of type σ

and (N, ρ, S) a configuration reachable from $(M, \emptyset, \varepsilon)$.

Then

(1) N is a subterm of M (hence typable and from a finite set).

(2) Environment ρ associates to a free variable x^α a closure (K, ρ') so that K also has type α .

We also say the closure is of type α .

Moreover, K is a subterm of M .

(3) The number of elements in S is determined by the type of N .

There are k elements when the type of N is $\alpha_1 \rightarrow \dots \rightarrow \alpha_k \rightarrow \sigma$.

We explain how Turing machines compute Böhm trees.

Definition:

Consider configuration (M, S, ϵ) .

- If the computation of the Turing machine from (M, S, ϵ) does not terminate, we set

$$KT_{\text{Tree}}(M, S, \epsilon) := \emptyset.$$

- If the computation terminates, then

$$(M, S, \epsilon) \rightarrow^* (b, S', (M_1, S_1) \dots (M_k, S_k)),$$

for some $b \neq \gamma$ and $b \neq \emptyset$.

In this situation,

$$KT_{\text{Tree}}(M, S, \epsilon) := \begin{array}{c} b \\ / \quad \dots \quad \backslash \\ KT_{\text{Tree}}(M_1, S_1, \epsilon) \quad \dots \quad KT_{\text{Tree}}(M_k, S_k, \epsilon). \end{array}$$

By the lemma above, k is the arity of b .

If we are working over tree signatures, $k=0$ or $k=2$.

- We now define

$$KT_{\text{Tree}}(M) := KT_{\text{Tree}}(M, \emptyset, \epsilon),$$

for M closed end of type σ .

The Krivine tree we just defined is exactly the Böhm tree.

Lemma:

Let M be closed end of type σ .

Then $KT_{\text{Tree}}(M) = BT(M)$.

6.2 Collapsible Pushdown Automata

- Goal:
- Define higher-order stacks: stacks of stacks of stacks of...
 - Define collapsible pushdown automata that work with such higher-order stacks.
 - What is special is a collapse operation that, intuitively, restores an environment.

This relates to Rule (4) in Krivine machines.

Intuitively, an order- n stack is a stack of order- $(n-1)$ stacks.

The stack characters are annotated by collapse links that point to a position in the stack.

This can be understood as the context in which the character was created.

Definition:

- Let Σ be a stack alphabet (alphabets are always finite) together with a partition function

$$\lambda: \Sigma \rightarrow [1, n], \quad n \geq 1.$$

// The partition assumption is not standard but often used.

- An order-0 stack (with up to order- n collapse links) is an unannotated stack symbol $a^i \in \Sigma \times \mathbb{N}$.
- An order- k stack (with up to order- n collapse links), $k \geq 1$, is a non-empty sequence $w = [w_1 \dots w_\ell]_k$ (with $\ell > 0$) such that each w_i is an order- $(k-1)$ stack (with up to order- n collapse links).

- By Stacks_n we denote the set of order- n stacks
(with up to order- n collapse links).

The top-of-stack is drawn to the left.

We use the following operations on stacks.

Definition:

Given an order- k stack with up to order- n collapse links,

- top_k returns the topmost element of the topmost order- k' stack.

The definition is by induction on the order:

Let $w = [w_1 \dots w_k]_k$.

Then

$$\text{top}_k(w) := w_1$$

$$\text{top}_{k'}(w) := \text{top}_{k'}(w_1), \text{ where } k' < k.$$

- The operation bot_k^i removes all but the last i elements from the topmost order- k stack.

It does not change the order ($\text{top}_{k'}$ returns an element of order $-(k'-1)$)

and requires $i \in [1, k]$ (with $w = [w_1 \dots w_k]_k$).

We have

$$\text{bot}_k^i(w) := [w_{k-i+1} \dots w_k]_k$$

$$\text{bot}_{k'}^i(w) := [\text{bot}_{k'}^i(w_1) \cdot w_2 \dots w_k]_k, \text{ where } k' < k.$$

- For technical convenience, $\text{top}_{n+1}(w) := w$.

- The idea on a collapse link i on $a \in \Sigma$ with $\lambda(a) = k$ is to collapse a stack w down to $\text{bot}_k^i(w)$,

provided this is defined.

Then $i=0$, the link is considered null.

We omit irrelevant collapse links to improve readability.

- To give an easy definition of (higher-order) push operations, we introduce another auxiliary operation.

Definition:

Let u be a $(k-1)$ -stack and $v = [v_1 \dots v_n]_n$ be an n -stack, with $k \in [1, n]$.

We define $u :_k v$ as the stack obtained by adding u on top of the topmost order- k stack in v . Formally,

$$u :_k v := [u, v_1, \dots, v_n]_n, \quad \text{if } k=n$$

$$u :_k v := [(u :_k v_1) v_2 \dots v_n]_n, \quad \text{if } k < n.$$

Example:

Let $\lambda(a) = 3$ and $\lambda(b) = 2$.

Let $w = \left[\left[[a^2 b^2]_2 \right]_2 \left[[b^2]_2 \right]_2 \left[[b^0]_2 \right]_2 \right]_3$ be an order-3 stack.

Then $\text{top}_2(w) = a^2$.

The destination of this link is

$$\text{bot}_3^2(w) = \left[\left[[b^0]_2 \right]_2 \right]_3.$$

Furthermore,

$$\text{bot}_2^2(w) = \left[\left[[b^2]_1 \right]_2 \left[[b^0]_1 \right]_2 \right]_3.$$

We have

$$\text{top}_2(w) = [a^2 b^2]_2.$$

Now

$$\text{top}_2(w) :_2 \text{bot}_2^2(w) = w.$$

Using the above auxiliary operations,
 collapsible pushdown automata modify order- n stacks as follows.

Definition:

We define the set

$$\text{Ops}_n := \{ \text{push}_k, \dots, \text{push}_n \} \cup \{ \text{push}_a, \text{rew}_a \mid a \in \Sigma \} \\ \cup \{ \text{pop}_k, \dots, \text{pop}_n \} \cup \{ \text{collapse} \}$$

of operations on order- n stacks used by collapsible pushdown automata.

The definition of the operations is as follows, with $w \in \text{Stacks}_n$:

$$\text{push}_k(w) := \text{top}_k(w) :_k w,$$

$$\text{push}_a(w) := a^{k-1} :_k w, \quad \text{Here } \text{top}_{k+1}(w) = [w_1 \dots w_k]_k \\ \text{with } k = \lambda(a) \text{ the link order,}$$

$$\text{pop}_k(w) := v, \quad \text{if } w = u :_k v,$$

$$\text{collapse}(w) := \text{bot}_k^i(w), \quad \text{if } \text{top}_k(w) = a^i \text{ and } \lambda(a) = k,$$

$$\text{rew}_b(w) := b^i :_k v, \quad \text{if } w = a^i :_k v \text{ and } \lambda(a) = \lambda(b).$$

Note:

Our definition of stacks does not permit the empty stack.

This means pop_k is undefined if the resulting v is empty.

Similarly, collapse is undefined if $i=0$.

Thus, the empty stack cannot be reached by the above operations.

Instead, the offending operations will be unavailable.

The same holds for a rewrite operation

that would change the order of a link.

Explanation:

- Operation push_k of order $k > 1$
copies the topmost order- k stack.
- Order-1 push_k pushes a onto the topmost order-1 stack,
annotated with an order- $\lambda(a)$ collapse link.
When executed on a stack w ,
the link destination is $\text{pop}_{\lambda(a)}(w)$.
- pop_k removes the topmost element from the topmost order- k stack.
- The rewrite operation rew_a modifies the topmost character
while maintaining the link.
- Collapse, when executed on a_i with $\lambda(a) = k$,
pops the topmost order- k stack
down to the last i elements.

Example:

Let $\lambda(a) = 3$, $\lambda(b) = 2$, and $w = [[[[a^2 b^2]_2 [b^2]_2]_2 [[b^0]_2]_2]_3$.

From the order-3 link 1 of the topmost a ,

we have $\text{collapse}(w) = u$ with $u = [[[b^0]_2]_2]_3$.

Now $\text{push}_3(u) = [[[[b^0]_2]_2 [[b^0]_2]_2]_3$.

If push_k on this stack results in

$$v = [[[[a^2 b^0]_2]_2 [[b^0]_2]_2]_3.$$

We have

$$\text{pop}_3(v) = u = \text{collapse}(v).$$

Note:

There is a subtlety in the interplay
of collapse links and higher-order pushes.

For a push, links pointing outside of $u = \text{top}_k(w)$

have the same destination in both copies of u .

Links pointing within u point to different sub-stacks.

Besides the order- n stack, collapsible pushdown automata
have a finite control.

Definition:

• An order- n collapsible pushdown automaton over (Σ, Δ)

is a tuple

$$C = (P, R, p_{\text{init}}, a_{\text{init}})$$

with

$\hookrightarrow P$ a finite set of control states
with initial state $p_{\text{init}} \in P$,

$\hookrightarrow R \subseteq P \times \Sigma \times \text{Ops}_n \times P$ a finite set of transitions
(or rules),

$\hookrightarrow a_{\text{init}} \in \Sigma$ an initial stack character.

• A configuration of C is a pair

$$c = (p, w) \in P \times \text{Stacks}_n.$$

The transition relation among configurations

is defined by

$$(p, w) \rightarrow (p', w'), \text{ if there is } (p, a, o, p') \in R$$

with $a = \text{top}_2(w)$ and $w' = \sigma(w)$.

The initial configuration is $(p_{\text{init}}, w_{\text{init}})$, $w_{\text{init}} = [\dots [a_{\text{init}}]_2 \dots]_n$.

To begin from another stack, one can adjust the rules
and build up the stack as required.

- A computation is a (potentially infinite) sequence of configurations, starting in the initial one,

c_0, c_1, \dots with $c_0 = c_{init}$ and $c_i \rightarrow c_{i+1}$ for all $i \in \mathbb{N}$.

Recall that transitions do not empty the stack
nor change the order of links.

Note:

- It is standard to assign priorities to states
and consider infinite computations.
- It is also common to assume an ownership partitioning of the states
and study parity games over collapsible pushdown automata.
- Also the relationship with Böhm trees can be made explicit.