

Handlers.Js: Compiling Effect Handlers to JavaScript

Daniel Hillerström

The University of Edinburgh, UK

November 16, 2021

(joint work with Sam Lindley, KC Sivaramakrishnan, Robert Atkey, and Jeremy Yallop)

Compiling Effect Handlers to JavaScript

Compiling Effect Handlers to JavaScript **Control Hostile Environments**

Disclaimer

This work is preliminary!
(originally presented at ProWeb'18)

The “take away” slide

Implementation	Extensions	Stack	Signature preserving	Translation
CPS	None	Explicit	No	Global
Abstract machine	None	Explicit	No	Global
Generators/iterators	Generators/iterators	Implicit*	No	Global
Stack inspection	Exception handlers	Explicit (lazy)	Yes [†]	Global

* Trampolining requires an explicit stack representation

[†] Modulo effect typing

A coroutine library in one slide

A signature with a single operation $\text{Yield} \stackrel{\text{def}}{=} \{\text{Yield} : \text{Unit} \rightarrow \text{Unit}\}$

type $\text{Co} \stackrel{\text{def}}{=} \text{Co}(\text{List Co} \rightarrow \text{Unit!Yield})$

$\text{coop} : (\text{Unit} \rightarrow \text{Unit!Yield}) \rightarrow \text{List Co} \rightarrow \text{Unit}$

$\text{coop } p \stackrel{\text{def}}{=} \text{handle } p \langle \rangle \text{ with}$

$\{ \text{return } x \mapsto \lambda ps. \text{case } ps \{ [] \mapsto \langle \rangle; \text{Co } r :: ps' \mapsto r \text{ } ps' \}$
 $\text{Yield } \langle \rangle k \mapsto \lambda ps. \text{let } k' = \text{Co } (\lambda ps. \text{resume } r \text{ with } ps) \text{ in}$
 $\text{let } (\text{Co } r :: ps') = ps ++ [k'] \text{ in}$
 $r \text{ } ps' \}$

$\text{coop_with} : (\text{Unit} \rightarrow \text{Unit!Yield}) \rightarrow \text{Co}$

$\text{coop_with } p \stackrel{\text{def}}{=} \text{Co}(\lambda ps. \text{coop } p \text{ } ps)$

$\text{cooperate} : \text{List } (\text{Unit} \rightarrow \text{Unit!Yield}) \rightarrow \text{Unit}$

$\text{cooperate } rs \stackrel{\text{def}}{=} \text{coop } (\lambda \langle \rangle. \langle \rangle) (\text{List.map } \text{coop_with } rs)$

Kinds of effect handlers

Affine, e.g. exception handlers

$$\{\mathbf{return} \ x \mapsto \text{Some } x \\ \text{Fail } \langle \rangle \ r \mapsto \text{None}\}$$

Linear, e.g. concurrency

$$\lambda q. \{\mathbf{return} \ x \mapsto \mathbf{if} \ \text{Queue.is_empty } q \ \mathbf{then} \ \langle \rangle \ \mathbf{else} \ (\text{Queue.pop } q) \ \langle \rangle \\ \text{Yield } \langle \rangle \ r \mapsto \text{Queue.push } r; (\text{Queue.pop } q) \ \langle \rangle\}$$

Multishot, e.g. nondeterminism

$$\{\mathbf{return} \ x \mapsto [x] \\ \text{Fork } \langle \rangle \ r \mapsto r \ \text{true} \ ++ \ r \ \text{false}\}$$

Tail-resumptive, e.g. dynamic binding

$$\lambda y. \{\mathbf{return} \ x \mapsto x \\ \text{Ask } \langle \rangle \ r \mapsto r \ y\}$$

Three ways of handling effects

Deep capture and resumption

handle $\mathcal{E}[\text{op } V]$ **with** $H \rightsquigarrow N[\mathbf{cont}_{\langle H; \mathcal{E} \rangle} / r, V/x]$, where $\{\text{op } p \ r \mapsto N\} \in H$
resume V **with** $W \rightsquigarrow \mathbf{handle} \ \mathcal{E}[W]$ **with** H , where $V = \mathbf{cont}_{\langle H; \mathcal{E} \rangle}$

Shallow capture and resumption

handle $\mathcal{E}[\text{op } V]$ **with** $H \rightsquigarrow N[\mathbf{cont}_{\langle \mathcal{E} \rangle} / r, V/x]$, where $\{\text{op } p \ r \mapsto N\} \in H$
resume V **with** $W \rightsquigarrow \mathcal{E}[W]$, where $V = \mathbf{cont}_{\langle \mathcal{E} \rangle}$

'Sheep' allocation, capture, and resumption


new $V \rightsquigarrow \mathbf{cont}_{\langle V [] \rangle}$, where $V = \lambda \langle \rangle. M$
handle $\mathcal{E}[\text{op } V]$ **with** $H \rightsquigarrow N[\mathbf{cont}_{\langle \mathcal{E} \rangle} / r, V/x]$, where $\{\text{op } p \ r \mapsto N\} \in H$
resume V **with** $\langle H; W \rangle \rightsquigarrow \mathbf{handle} \ \mathcal{E}[W]$ **with** H , where $V = \mathbf{cont}_{\langle \mathcal{E} \rangle}$

JavaScript: the challenges


Some challenges posed by the JavaScript web runtime

- Event-driven computation: necessitates yields every-so-often
- No tail call elimination
- No access to the control state
- Difficult to predict performance


OCaml: effect handlers via segmented stacks (Sivaramakrishnan et al. 2021)




OCaml: effect handlers via segmented stacks (Sivaramakrishnan et al. 2021)



OCaml: effect handlers via segmented stacks (Sivaramakrishnan et al. 2021)



OCaml: effect handlers via segmented stacks (Sivaramakrishnan et al. 2021)




Effect handlers via generalised continuations (Hillerström et al. 2020)

Generalised continuations are a *stackless* encoding of segmented stacks.

Key idea: Represent segmented stacks as list of lists.

type GCont $\stackrel{\text{def}}{=} \text{List} \langle \text{List Frame} \times \text{Hdef} \rangle$

Pictorially



Continuation passing style

Pros

- Well-established compilation technique
- Well-understood use of λ -calculus
- Well-optimised programming idiom in JavaScript engines

```
[[let x = M in N]] = function(k) {  
    return [[M]](function(x) { return [[N]](k); });  
}
```

Cons

- Difficult to debug (esp. with a trampoline)
- Everything gets heap allocated
- CPS transform for effect handlers is complicated (Hillerström et al. 2017)

Continuation passing style (cont'd)

Naïve idea: Compile effect handlers by argument passing, e.g.

$$\begin{aligned} \llbracket \mathbf{perform} \ell V \rrbracket &\stackrel{\text{def}}{=} \lambda k. \lambda h. h \langle \ell, \langle \llbracket V \rrbracket, \lambda x. k \ x \ h \rangle \rangle \\ \llbracket \mathbf{let} \ x = M \ \mathbf{in} \ N \rrbracket &\stackrel{\text{def}}{=} \lambda k. \llbracket M \rrbracket (\lambda x. \llbracket N \rrbracket k) \\ \llbracket \mathbf{return} \ V \rrbracket &\stackrel{\text{def}}{=} \lambda k. k \ \llbracket V \rrbracket \end{aligned}$$

Continuation passing style (cont'd)

Naïve idea: Compile effect handlers by argument passing, e.g.

$$\begin{aligned}\llbracket \mathbf{perform} \ell V \rrbracket &\stackrel{\text{def}}{=} \lambda k. \lambda h. h \langle \ell, \langle \llbracket V \rrbracket, \lambda x. k \ x \ h \rangle \rangle \\ \llbracket \mathbf{let} \ x = M \ \mathbf{in} \ N \rrbracket &\stackrel{\text{def}}{=} \lambda k. \llbracket M \rrbracket (\lambda x. \llbracket N \rrbracket k) \\ \llbracket \mathbf{return} \ V \rrbracket &\stackrel{\text{def}}{=} \lambda k. k \llbracket V \rrbracket\end{aligned}$$

The image is not *properly tail-recursive*

$$\begin{aligned}\top \llbracket \mathbf{return} \ \langle \rangle \rrbracket &= (\lambda k. k \ \langle \rangle) (\lambda x. \lambda h. x) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow ((\lambda x. \lambda h. x) \ \langle \rangle) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow (\lambda h. \langle \rangle) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow \langle \rangle\end{aligned}$$

(the image also contains *static administrative redexes*)

Continuation passing style (cont'd)

Naïve idea: Compile effect handlers by argument passing, e.g.

$$\begin{aligned}\llbracket \mathbf{perform} \ell V \rrbracket &\stackrel{\text{def}}{=} \lambda k. \lambda h. h \langle \ell, \langle \llbracket V \rrbracket, \lambda x. k \ x \ h \rangle \rangle \\ \llbracket \mathbf{let} \ x = M \ \mathbf{in} \ N \rrbracket &\stackrel{\text{def}}{=} \lambda k. \llbracket M \rrbracket (\lambda x. \llbracket N \rrbracket k) \\ \llbracket \mathbf{return} \ V \rrbracket &\stackrel{\text{def}}{=} \lambda k. k \llbracket V \rrbracket\end{aligned}$$

The image is not *properly tail-recursive*

$$\begin{aligned}\top \llbracket \mathbf{return} \ \langle \rangle \rrbracket &= (\lambda k. k \ \langle \rangle) (\lambda x. \lambda h. x) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow ((\lambda x. \lambda h. x) \ \langle \rangle) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow (\lambda h. \langle \rangle) (\lambda \langle z, _ \rangle. \mathbf{absurd} \ z) \\ &\rightsquigarrow \langle \rangle\end{aligned}$$

(the image also contains *static administrative redexes*)

Continuation passing style (cont'd)

Solution: Uncurry the translation, e.g.

$$\begin{aligned}\llbracket \mathbf{perform} \ell V \rrbracket &\stackrel{\text{def}}{=} \lambda(k :: h :: ks).h \langle \ell, \langle \llbracket V \rrbracket, \lambda x. \lambda ks'. k x (h :: ks') \rangle \rangle ks \\ \llbracket \mathbf{let} x = M \mathbf{in} N \rrbracket &\stackrel{\text{def}}{=} \lambda(k :: ks). \llbracket M \rrbracket ((\lambda x. \lambda ks'. \llbracket N \rrbracket (k :: ks')) :: ks) \\ \llbracket \mathbf{return} V \rrbracket &\stackrel{\text{def}}{=} \lambda(k :: ks). k \llbracket V \rrbracket ks\end{aligned}$$

Now the image is properly tail-recursive

$$\begin{aligned}\top \llbracket \mathbf{return} \langle \rangle \rrbracket &= (\lambda(k :: ks). k \langle \rangle ks) ((\lambda x. \lambda ks. x) :: (\lambda \langle z, _ \rangle. \lambda ks. \mathbf{absurd} z) :: []) \\ &\rightsquigarrow (\lambda x. \lambda ks. x) \langle \rangle ((\lambda \langle z, _ \rangle. \lambda ks. \mathbf{absurd} z) :: []) \\ &\rightsquigarrow \langle \rangle\end{aligned}$$

(the image still contains static administrative redexes)

Continuation passing style (cont'd)

Long story short: Efficient one-pass CPS transform for effect handlers requires partial evaluation, e.g.

$$\begin{aligned} \llbracket \mathbf{perform} \ell V \rrbracket &\stackrel{\text{def}}{=} \bar{\lambda}\bar{\theta}, \bar{\xi}, \chi^{\text{ret}}, \chi^{\text{ops}} \bar{\cdot} \bar{\cdot} \kappa. \downarrow \chi^{\text{ops}} \ @ \ \langle \downarrow \xi, \ell, \langle \llbracket V \rrbracket, \langle \downarrow \theta, \langle \downarrow \xi, \downarrow \chi^{\text{ret}}, \downarrow \chi^{\text{ops}} \rangle \rangle \bar{\cdot} \bar{\cdot} \llbracket \rrbracket \rangle \ @ \ \downarrow \kappa \\ \llbracket \mathbf{let} x = M \mathbf{in} N \rrbracket &\stackrel{\text{def}}{=} \bar{\lambda}\bar{\theta}, \chi \bar{\cdot} \bar{\cdot} \kappa. \llbracket M \rrbracket \ @ \ (\bar{\lambda} \uparrow ((\underline{\lambda} x \mathbf{ks.let} (k \bar{\cdot} \bar{\cdot} ks') = ks \mathbf{in} \llbracket N \rrbracket \ @ \ (\uparrow k \bar{\cdot} \bar{\cdot} \uparrow ks')) \bar{\cdot} \bar{\cdot} \downarrow \theta), \chi \bar{\cdot} \bar{\cdot} \kappa) \\ \llbracket \mathbf{return} V \rrbracket &\stackrel{\text{def}}{=} \bar{\lambda} \kappa. \mathbf{app} (\downarrow \kappa) \llbracket V \rrbracket \end{aligned}$$

Intuition: Explicit passing of the control state (represented as a generalised continuation)

$$\llbracket f \langle \rangle \rrbracket = f \left(\begin{array}{|c|c|c|c|} \hline m_2 & h_2 & m_1 & h_1 \\ \hline \end{array} \right)$$

$$\langle C \mid E \mid K \rangle$$

The CEK machine consists of three components

- Control, the expression being evaluated
- Environment, binding free variables
- Kontinuation, the continuation of C

$$\langle C \mid E \mid K \rangle$$

The CEK machine consists of three components

- Control, the expression being evaluated
- Environment, binding free variables
- Kontinuation, the continuation of C

Classic continuation structure (Felleisen and Friedman 1987)

K : List Frame

$$\langle C \mid E \mid K \rangle$$

The CEK machine consists of three components

- Control, the expression being evaluated
- Environment, binding free variables
- Kontinuation, the continuation of C

Obtain a runtime for handlers by plugging in generalised continuations (Hillerström 2021)

$$K : \text{List} \langle \text{List Frame} \times \text{Hdef} \rangle$$

CEK: abstract machine (cont'd)

Pros

- Fairly easy to debug
- Hot machine loop
- Well-understood technique

Cons

- Allocation intensive
- Everything is boxed
- Basic operations are not JIT optimised

```
let control = initial_program;
let environment = initial_env;
let kontinuation = k_identity; // generalised continuation
let config = NORMAL;

while (true) {
  switch (config) {
    case NORMAL:
      switch (match(control)) {
        // Match AST
        // ...
        continue;
      }
    case ADM_APPLY_CONT:
      // Apply 'kontinuation' to the value in 'control'
      // ...
      continue;
  }
}
```


Generators and iterators: native JavaScript continuations

Generators and iterators provide a form of delimited control (James and Sabry 2011).

The main idea

- Transform every function into a generator
- Transform each handler into a generator that iterates its given computation

$$\begin{aligned} \llbracket \mathbf{perform} \ell V \rrbracket &= \mathbf{yield} \{ \text{'tag':} \ell, \text{'payload':} \llbracket V \rrbracket \} \\ \llbracket \mathbf{let} x = M \mathbf{in} N \rrbracket &= \mathbf{var} x = \llbracket M \rrbracket; \llbracket N \rrbracket \\ \llbracket \mathbf{return} V \rrbracket &= \llbracket V \rrbracket \end{aligned}$$

Pros

- Native JavaScript
- Fairly easy to debug (unless trampolined)

Cons

- Everything lives in a generator (**function***)
- Generators/iterators use simulated continuations under-the-hood

Generators and iterators: native JavaScript continuations (cont'd)

Potential for efficient compilation of tail-resumptive handlers

```
[[Hreader := λv. {return x ↦ x; Ask ⟨⟩ r ↦ r v}]]  
= function* Hreader(v, m) {  
  var g = m();  
  var res = g.next(); // run until first yield  
  while (!res.done) {  
    switch (res.value.tag) {  
      case "Ask":  
        res = g.next(v); // invoke the continuation with 'v'  
        break;  
      default: // effect forwarding  
        var y = yield res.value;  
        res = g.next(y);  
    }  
  }  
}
```

Generalised stack inspection

Generalised stack inspection provides simulate continuations using exceptions (Pettyjohn et al. 2005).

- Enclose every binding in an exception handler
- Throw an exception to assemble the continuation

Pros

- Compatibility with (first-order) legacy code
- Continuation capture is on-demand

Cons


- Difficult to debug
- Deoptimisations often kick in

```
[[perform  $\ell$  V]] = throw new PerformOperation( $\ell$ , [[V]]);
[[let x = M in N]] = function let_x(...) {
    var x;
    try {
        x = [[M]](...);
    } catch (e) {
        if (e instanceof PerformOperation) {
            e.augment([[N]]);
            throw e;
        } else {
            throw e;
        }
    }
    return [[N]](x, ...);
}
[[return V]] = function(...) { return [[V]]; }
```

Generalised stack inspection (cont'd)

Initially there is only the call stack

Call stack




Continuation


Generalised stack inspection (cont'd)

Throwing an exception causes the continuation to materialise

Call stack




Continuation




Generalised stack inspection (cont'd)

Instantiate the abstract handler once we pass over a concrete handler

Call stack




Continuation




Generalised stack inspection (cont'd)

Continue unwinding the call stack

Call stack




Continuation




Generalised stack inspection (cont'd)

Continue unwinding the call stack

Call stack



Continuation




Generalised stack inspection (cont'd)

Notice that the continuation was built in reverse

Call stack

Continuation




Generalised stack inspection (cont'd)


The continuation is reversed prior to invocation

Call stack


Continuation



Preliminary experiments



Preliminary experiments



Effect handlers benchmark suite

Join the effort to build a standardised benchmark suite

<https://github.com/effect-handlers/effect-handlers-bench>

- Selective linear CPS
 - Direct code whenever possible
 - Optimised for one-shot continuations
 - Reuse allocated frame slots
- Scheme Gambit virtual machine (Thivierge and Feeley 2012)
 - Use JavaScript as an assembly language
 - Primitive values
 - Maintains a “shadow” stack using a JavaScript array

Summary

In summary

Implementation	Extensions	Stack	Signature preserving	Translation
CPS	None	Explicit	No	Global
Abstract machine	None	Explicit	No	Global
Generators/iterators	Generators/iterators	Implicit*	No	Global
Stack inspection	Exception handlers	Explicit (lazy)	Yes [†]	Global

* Trampolining requires an explicit stack representation

[†] Modulo effect typing

References I

- Felleisen, Matthias and Daniel P. Friedman (1987). “Control Operators, the SECD-machine, and the λ -Calculus”. In: *Formal Description of Programming Concepts III*, pp. 193–217.
- Ganz, Steven E., Daniel P. Friedman, and Mitchell Wand (1999). “Trampolined Style”. In: *ICFP*. ACM, pp. 18–27.
- Pettyjohn, Greg et al. (2005). “Continuations from generalized stack inspection”. In: *ICFP*. ACM, pp. 216–227.
- James, Roshan P and Amr Sabry (2011). “Yield: Mainstream delimited continuations”. In: *TPDC*.
- Thivierge, Eric and Marc Feeley (2012). “Efficient compilation of tail calls and continuations to JavaScript”. In: *Scheme@ICFP*. ACM, pp. 47–57.
- Hillerström, Daniel et al. (2017). “Continuation Passing Style for Effect Handlers”. In: *FSCD*. Vol. 84. LIPIcs, 18:1–18:19.
- Hillerström, Daniel, Sam Lindley, and Robert Atkey (2020). “Effect handlers via generalised continuations”. In: *J. Funct. Program.* 30, e5.
- Hillerström, Daniel (2021). “Foundations for Programming and Implementing Effect Handlers”. PhD thesis. School of Informatics, The University of Edinburgh, UK.

Sivaramakrishnan, KC et al. (2021). "Retrofitting Effect Handlers onto OCaml". In: *CoRR* abs/2104.00250.