STATE OF EFF

Matija Pretnar University of Ljubljana, Slovenia



I was trying to reconstruct the history of Eff, and as far as I recall, the first version appeared in 2010.

Mathematics and Computation
Posts Talks Publication Software About
state vi
← How eff han Programming with effects I: Theory →
Progr (lamba lookup ()
operation S: Vield
(Lambda s: Yield () S point lideas about the algebraic
nature of (lambda s: (s, y)) ; a precise connection combined.
If you have <u>Mercurial</u> installed (type hg at command prompt to find out) you can get en mo
Otherwise, you may also download the latest source as a <u>.zip</u> or <u>.tar.gz</u> , or <u>visit the repository with your browser</u> for other versions. Eff is released under the simplified BSD License.
To compile eff you need <u>Ocaml</u> 3.11 or newer (there is an incompatibility with 3.10 in the Lexer module), <u>ocamlbuild</u> , and <u>Menhir</u> (which are both likely to be bundled with Ocaml). Put the source in a suitable directory and compile it with make to create the Ocaml bytecode executable eff.byte. When you run it you get an interactive shell without line editing capabilities. If you never make any typos that should be fine, otherwise use one of the line editing wrappers, such as <u>riwrap</u> or <u>ledit</u> . A handy shortcut eff runs eff.byte wrapped in riwrap.
Suntay

It waas announced on Andrej's blog, 13 years ago almost to a day. Two interesting notes: it's name was written in lower case, and it used a Python-like syntax.



On our train ride back from the Domains workshop in Swansea, Andrej and I realised that Eff fits much better into an ML family.

Programming with Algebraic Effects and Handlers type 'a ref = effect **operation** lookup: unit -> 'a **operation** update: 'a -> unit end let state r x = handler r#lookup () k -> (**fun** s -> k s s) r#update s' k -> (fun s -> k () s') **val** y -> (**fun** s -> (y, s)) reprint submitted to Florvier November 12, 2013

Eff 2.0 is what appeared in the "Programming with Algebraic Effects and Handlers" paper. At this point, Eff featured dynamic generation of effect instances, inspired by how references are created in OCaml.



After that point, the versioning becomes much more confusing. Version 3.0 appeared some time after that, but I do not know what the main difference from 2.0 was. So let's stick to chronological history from now onwards.



In 2013, Eff got a subtyping-based effect system.



Its inference algorithm was constraint based and used a variant of Francois Pottier's garbage collection.



In 2015, I visited Tom Schrijvers in Leuven, and we started looking at efficient evaluation of handlers.

```
effect Put: int -> unit
effect Get: unit -> int
let rec loop n =
 if n = 0 then () else
    perform (Put (perform (Get ()) + 1));
    loop(n-1)
let state_handler = handler
   effect (Put s') k -> (fun _ -> k () s')
   effect (Get ()) k -> (fun s -> k s s)
  _ -> (fun s -> s)
let main n =
  (with state_handler handle loop n) 0
let main n =
 let rec state_handler_loop m s =
    if m = 0 then s
             else state_handler_loop (m - 1) (s + 1)
  in
  state_handler_loop n 0
```

Tom's idea was to do source-level optimisations of Eff code and then to compile it down to OCaml using some sort of a monadic embedding (OCaml didn't have effects at that point yet).



Our initial results were promising, reaching performance of hand-written OCaml code.



But we ended up having too little examples for a thorough evaulation since the effect system was very fragile and every tweak either made our compilation too conservative and thus too slow, or too agressive and ill-typed.



The first step in resolving the mess was getting rid of instances (since one could use the generativity of modules (which Eff still doesn't have)), though even that did not help.



The next step was to change all the implicit subtyping coercions in Eff into explicit ones, and this was done with the help of Tom Schrijvers and his students, mostly Amr Hany Saleh and Georgios Karachalias.

```
let apply_if p f x =
    if p x then
        f x
    else
        x
```

For example, take the following Eff function.

```
\begin{array}{l} \texttt{fun} \ p \mapsto \texttt{return} \ (\texttt{fun} \ f \mapsto \texttt{return} \ (\texttt{fun} \ x \mapsto (\\ \texttt{do} \ b \leftarrow p \ x;\\ \texttt{if} \ b \ \texttt{then} \ f \ x \ \texttt{else} \ \texttt{return} \ x\\ ))) \end{array}
```

In fine-grain call-by-value, this gets elaborated to the following function with explicit binds and returns.



In the new explicitly annotated version, each variable is assigned a type, and types have to match, which is achieved through coercions. For example, the type a_4 of x doesn't need to match the argument type a_1 of f, but there must be an explicit coercion $\omega_1 : a_1 \le a_4$ witnessing the subtyping. Now, all effect information can be read directly off the syntax.



In parallel, a fork of Eff called EEFF appeared.



This was done by my PhD student Žiga Lukšič and was based on the work of extending the effect system with equations that the handlers need to satisfy. In EEFF, one can specify equations that have to hold, but the compiler does not do any checks along the lines of QuickCheck or SMT solving, just prints out the obligations to the user. If anyone is interested in helping continue this work, let me know!



With the explicit subtyping sorted out, we decided to properly implement it, and this was done mostly in 2020 with the help of my student Filip Koprivec.

```
let apply (exp1, exp2) =
  match exp1.ty.term with
  | Type.Arrow (ty1, drty2) ->
    assert (Type.equal_ty exp2.ty ty1);
    { term = Apply (exp1, exp2); ty = drty2 }
    | _ -> assert false
```

One useful technique we used were smart-constructors of typed terms, which raised an assertion fault as soon as some types did not match. This caught countless bugs in source-level optimizations, which shuffle terms around a lot.



With Eff cleaned up, it was time in 2021 to return to our initial work on optimizations.



This worked smoothly, and the current pipeline is as follows. After desugaring, Eff is first elaborated into an explicitly typed core language. Next, it is translated into an OCaml-like language that features no native effects, just their monadic reification. On both languages, mostly on the core one, we perform source-level optimizations that inline handlers, extract pure computations, ...



Our current work focuses on simplifying subtyping coercions.



Recall that core language features explicit coercions in terms.

```
let apply_if w1 w2 w3 w4 p f x =
    p (x |> w1) >>= fun b ->
    if b then
       (f (x |> w2)) |> coer_comp w3
    else
       return (x |> w4)
```

When translating to OCaml, these coercions gain computational meaning (eg. they embed pure values into the monad), and any polymorphic coercion parameters get translated into additional function arguments. Very simple functions have a handful of such additional arguments, while a quick-sort implementation for example, already features a couple hundred of them, which is unacceptable.



Unfortunately, we cannot perform Francois Pottiers garbage collection because that relies on computational irrelevance of coercions and completely removes them. In our case, coercions form part of terms and have to be replaced rather than removed. Instead we rely on heuristics such as collapsing of cycles, collapsing coercions between parameters with singualr bounds, ... We cannot get rid of all the coercions in general, but we can for example remove all of them in Eff's standard library or the quick-sort example.



The correctness of this simplifications is established in a paper that Filip and I are submitting soon to LMCS.



What are the next steps for Eff?



The answer lies in Millet, a fine-grain call-by-value based ML-like language.



A few years ago, I worked with Danel Ahman on asynchronous effects. To test our ideas, I've developed a small prototype language called Æff. I got it by taking Eff, removing some stuff and adding some new one. Of course, during writing of Æff I corrected some mistakes present in Eff. How to port those changes back to get an improved version of Eff? What I started is Millet, a template language featuring all the boring but useful things in a language (recursive types, records, interactive loop, parser, lexer, desugarer, simple type-checker, interpreter, ...) which one can fork and add only the interesting bits (handlers, asynchronous operations, ...). Eff and Æff appeared before, and now I am slowly porting them to the point where they can be considered such forks. Now, if Millet gets an additional feature (right now, a couple of students of mine are developing a module system, a LSP, and a Wasm backend), one can "simply" merge those changes into forks. If you are researching a new feature and want to try it out, I encourage you to take a look at Millet.