
Homework (DM)

Compilation and Program Analysis (CAP)

Algebraic Effects and Effect Handlers

Instructions:

1. Every single answer must be informally explained AND formally proved.
2. Using LaTeX is NOT mandatory at all.
3. Vous avez le droit de rédiger en Français.

In this Homework, you will study algebraic effects and handlers. The language we will consider differs in nature from **WHILE**: it has a rather functional nature, as opposed to imperative. Algebraic effects are a relatively modern creation, essentially popularized by Matija Pretnar's PhD thesis in 2010. They have since then percolated into real world languages, and are notably part of OCaml 5.

1 Preliminaries: lambda-calculus

We first introduce λ , a simple purely functional programming language whose syntax is depicted on Figure 1. λ is quite standard, except perhaps that values and computations are explicitly split into two categories. Values are "inert", they do not reduce, and include $()$, the only value of the unit type, booleans, natural numbers, variables, and (anonymous) functions. Computations are proper programs: they include return statements, sequences, conditionals, and function application. Figure 2 defines λ 's (small step) operational semantics. The last four rules describe the reduction of the computations, while the first one lifts this reduction under evaluation contexts.

Values:	Computations:
$v \in \mathcal{V} ::= () \mid \# \mid \text{ff} \mid n$ \mathbf{x} $\lambda \mathbf{x}.c$	$c \in \mathcal{C} ::= \text{return } v$ $\text{let } \mathbf{x} = c_1 \text{ in } c_2$ $\text{if } v \{ c_1 \} \text{ else } \{ c_2 \}$ $v_1 v_2$
Constants Variables Functions	Return Sequence Conditional Application

Figure 1: Grammar for λ **Small step relation for λ : $c \rightsquigarrow c'$**

We define the evaluation contexts C as:

$$C ::= [] \mid \text{let } \mathbf{x} = C \text{ in } c$$

$$\begin{array}{c}
 \frac{c \rightsquigarrow c'}{C[c] \rightsquigarrow C[c']} \\
 \\
 \frac{}{\text{let } \mathbf{x} = \text{return } v \text{ in } c_2 \rightsquigarrow c_2[v/\mathbf{x}]} \qquad \frac{}{\text{if } \# \{ c_1 \} \text{ else } \{ c_2 \} \rightsquigarrow c_1} \\
 \\
 \frac{}{\text{if } \text{ff} \{ c_1 \} \text{ else } \{ c_2 \} \rightsquigarrow c_2} \qquad \frac{}{(\lambda \mathbf{x}.c) v \rightsquigarrow c[v/\mathbf{x}]}
 \end{array}$$

Figure 2: Operational semantics for λ **2 A calculus with effects: λ_{eff}**

The core idea of algebraic effects is to think of your programs as purely functional programs, but with the ability to ask questions to the environment through operations: impure behaviour arises from a set of operations.

For instance, suppose your program interact with a memory cell containing a single bit of information. You would traditionally think of your program as a function taking a boolean as input, your state, and returning a new boolean, the updated state. Rather, we will think of our programs as a series of interaction with this cell through an interface of two operations: `get` to read the value of the cell, and `put` to update the value of the cell.

Consider now a program `flip` negating the content of the cell. It should be thought of as a sequence of two actions: `get` the content of the cell, then `put` the negation of the value read, and `terminate`, returning the unit value `()`. This program can be represented as something very close to an AST, as depicted on the left hand-side of Figure 3.

You may notice that this tree explains how the computation proceeds no matter which value is read, but does not explain how to know what value should be read in a concrete run, or even what is a cell for that matter. Worst, if for some reason our `flip` program were to read the cell twice in a row before updating the value, we would result in the right hand-side of Figure 3. Here, some leaves feel weird, suggesting we could read first `true` and then `false` from our cell without any update in between.

Indeed, we have described a program that interacts with an environment through an interface of operations, but we have not described the environment! This is where effect handlers come into




Figure 3: Flipping the content of a cell: tree representation

Values:

$$v \in \mathcal{V} ::= \dots$$

| h handler

Handlers:

$$h \in \mathcal{H} ::= \text{handler } \left\{ \begin{array}{l} \text{return } \mathbf{x} \mapsto c_r, \\ \text{op}_1(\mathbf{x}, \mathbf{k}) \mapsto c_1; \dots; \text{op}_n(\mathbf{x}, \mathbf{k}) \mapsto c_n \end{array} \right\}$$

Return clause
Operational clauses

Computations:

$$c \in \mathcal{C} ::= \dots$$

| with h handle c Handling

Figure 4: Grammar for λ_{eff} **Small step relation for λ_{eff} :** $c \rightsquigarrow c'$ In the following rules, we set $h = \text{handler } \{ \text{return } \mathbf{x} \mapsto c_r, \text{op}_1(\mathbf{x}, \mathbf{k}) \mapsto c_1; \dots; \text{op}_n(\mathbf{x}, \mathbf{k}) \mapsto c_n \}$

$$\frac{}{\text{with } h \text{ handle return } v \rightsquigarrow c_r[v/\mathbf{x}]}$$

$$\frac{}{\text{with } h \text{ handle } C[\text{perform op}_i(v)] \rightsquigarrow c_i[v/\mathbf{x}, \lambda \mathbf{y}. \text{with } h \text{ handle } C[\text{return } \mathbf{y}]/\mathbf{k}]}$$

Figure 5: Operational semantics for λ_{eff}

play: they provide implementations to the operations. They can be thought of as an extension of exception handling: the control flow is passed to a harness catching the operations, and deciding what to do in its lieu. For instance, the handler may implement our get and put effects as mentioned above, by explicitly passing as argument the cell.

A syntax for λ_{eff} . Let's now package these ideas into a calculus. Figure 4 extends λ 's syntax with two new kind of computations. One can perform an operation op over an argument v with the computation $\text{perform op}(v)$. Furthermore, computations can be wrapped into an effect handler. This handler specifies how the computation must proceed when one of two situations arises. If the computation returns, the Return clause triggers. When an operation is performed, the corresponding Operational clause happens. In this case, the corresponding computation c_i depends on the argument of the operation (bound to \mathbf{x}), but also on the continuation of the

computation (bound to \mathbf{k}): contrary to exception handlers, we usually want to get back to our main computation after falling in our harness! Figure 5 explicit these new semantic rules.

Conventions: to simplify the writing of programs, we introduce the following notations:

- λ_{eff} variables are in bold: \mathbf{x} .
- When the return close of a handler is omitted, it is assumed to be the identity, i.e., $\lambda \mathbf{x}.\text{return } \mathbf{x}$.
- We write op instead of $op()$ when an operation expects no argument, such as `get`.
- We write $c_1;;c_2$ for `let $\mathbf{x} = c_1$ in c_2` if \mathbf{x} does not appear in c_2 .
- We will allow ourselves to use data-types such as pairs or lists in the remaining even if they have not been formally defined in the language.

We can finally rewrite `flip` in λ_{eff} :

$$\text{flip} \triangleq \text{let } \mathbf{x} = \text{perform } \text{get}() \text{ in perform } \text{put}(\mathbf{x})$$

. If we tried to evaluate `flip` according to our operational semantics, we would be immediately stuck. We hence implement a handler that passes around the boolean cell.

$$h_{\text{cell}} \triangleq \text{handler} \left\{ \begin{array}{l} \text{return } \mathbf{x} \mapsto \lambda \mathbf{b}.\text{return } \mathbf{x} \\ \text{get}(\mathbf{k}) \mapsto \lambda \mathbf{b}.\langle \mathbf{k} \ \mathbf{b} \ \mathbf{b} \rangle \\ \text{put}(\mathbf{x}, \mathbf{k}) \mapsto \lambda \mathbf{b}.\langle \mathbf{k} \ () \ \mathbf{x} \rangle \end{array} \right\}$$

In the implementation of `get`, the current value of the cell is passed to the continuation, and the value of the cell is left unchanged. For `put`, we simply pass `()`, but update the value of the cell. We can finally get a full computation, which expects an initial cell to run c_i :

$$\text{(with } h_{\text{cell}} \text{ handle flip) } c_i$$

Question #1

Write a λ_{eff} computation c_{and} that takes a natural number n as argument, and returns whether n is odd if the cell is set to true, and whether n is even if the cell is set to false.

Question #2

Assuming λ_{eff} is extended with lists, propose an alternate handler for `get` and `put` denoted h_{log} that logs the history of all booleans that are read via `get` and returns this history at the end of the computation.

Question #3

Consider

$$c_{\text{ex}} \triangleq \text{let } \mathbf{x} = \text{get} \text{ in let } \mathbf{y} = \text{get} \text{ in put } \mathbf{x};; \text{return orb } \mathbf{x} \ \mathbf{y}$$

where `orb` is the boolean disjunction. Show the traces of execution of `with h_{cell} handle c_{ex}` and `with h_{log} handle c_{ex}` applied to an initial cell set to true.

Question #4

Consider now a new effect: failure. To do so, we introduce a new operation `fail` whose handler discards the continuation:

$$h_{fail} \triangleq \text{handler } \{\text{fail}(\mathbf{k}) \mapsto \text{return } ()\}$$

We can now consider two ways to nest handlers: with h_{fail} handle with h_{log} handle c or conversely with h_{log} handle with h_{fail} handle c .

Are those equivalent? Justify your answer either by an informal argument, or a counter-example.

Question #5

We now consider the case of non-determinism through an operation `toss()` which picks randomly a boolean. Propose an interface and handler h_{toss} implementing the non-determinism by collecting all possible results.

2.1 Algebraic Simplification

We have informally qualified some handlers as correct or incorrect. This intuition has a formal ground, hinted at the name of algebraic effects. Operations form an algebra, i.e., their signature comes with a set of equations that axiomatize their behavior.

For instance, let us consider the case of `get/put`. We may axiomatize the interaction between each pair of operations, i.e.:

$$\text{put } v_1;; \text{put } v_2;; c \equiv \text{put } v_2;; c \quad (1)$$

$$\text{put } v;; \text{let } x = \text{get in } c \equiv \text{put } v;; c[v/x] \quad (2)$$

$$\text{let } x = \text{get in put } x; c \equiv \text{let } x = \text{get in } c \quad (3)$$

$$\text{let } x = \text{get in let } y = \text{get in } c \equiv \text{let } x = \text{get in } c[x/y] \quad (4)$$

Spelled out, these equations state that (1) two consecutive puts reduce to the latest one, (2) a put determines uniquely the result of a get, (3) putting in memory the same cell as the one we have read is useless, and (4) two consecutive gets can be condensed to a single one.

Some care must be taken in this definition: when introducing a piece of state such as with h_{cell} , the equality $c_1 \equiv c_2$ should be understood as (with h_{cell} handle c_1) s_i and (with h_{cell} handle c_2) s_i reduce to the same value for any initial state s_i .

"Correct" handlers for `get/put` are hence defined in the rest of this section as the ones that satisfy these four equations.

Question #6

Propose an alternate handler h_{wrong} for `get` and `put` that would be blatantly semantically incorrect.

Question #7

Check that h_{cell} is correct. Is h_{log} correct?

Question #8

Prove that this axiomatization is not minimal by deriving (4) from the three other equations.

Question #9

These equations can be an opportunity for optimization! Write a program transformation that exploits the cell algebra. This transformation should be defined by mutual induction on the syntax of values and computations, i.e.:

$$\begin{aligned} \llbracket () \rrbracket_v &= () && \dots \\ \llbracket v_1 \ v_2 \rrbracket_c &= \llbracket v_1 \rrbracket_v \ \llbracket v_2 \rrbracket_v && \dots \end{aligned}$$

Justify its correctness informally—in particular, which hypothesis must you impose on the handlers used in your program?

Question #10 (Difficult)

Open question

Does your transformation optimize the following program:

```
let x = get in let y = 2 * x in let z = get in return y + z.
```

Discuss how you could improve the transformation to better optimize programs.

Question #11

Recall the non-determinism operation from Question #4: extend it with equations to suggest an algebra of non-determinism.

3 Compilation to CPS

We now wish to compile our language to another without support for algebraic effects; for instance Javascript, to run our programs in the browser! For this purpose, we must translate algebraic effects into simpler construct. Here, we will target simple functions, using the so-called Continuation Passing Style (CPS). In continuation passing style, the control flow is made explicit via an argument, usually denoted **k** and called the **k**ontinuation, which is the future of the execution. For instance, let us consider the program $\text{flip} \triangleq \text{let } x = \text{perform } \text{get}() \text{ in perform } \text{put}(x)$ defined earlier.

$$\text{flip}_{cps} \triangleq \lambda k. \lambda k_{\text{get}}. \lambda k_{\text{put}}. (k_{\text{get}} (\lambda x. k_{\text{put}} (x, k)))$$

We can then consider the following program:

$$p_{cps} \triangleq \lambda k. (\text{flip}_{cps} (\lambda x. \lambda b. \text{return } (k \ x))) (\lambda k. \lambda b. (k \ b) \ b) (\lambda (x, k). \lambda b. (k \ ()) \ x)$$

In the rest of this section, we will aim to obtain this program from a suitable input.

Question #12

To which handler is this program equivalent? Justify informally.

3.1 Code without closure

To make our task initially simpler, we consider a sublanguage of λ_{eff}^- without λ -expressions and function calls, named λ_{eff}^- . The function $\llbracket e \rrbracket_k^H$ compiles the λ_{eff}^- expression e to Continuation Passing Style given a continuation k and a (meta) handler H . It is shown in [Figure 7](#). The idea is, as we progress through handlers, to map each label op used to designate an algebraic effect with a continuation that indicates the code to execute when this effect is performed. H is thus a map from labels op to their continuations k_{op} . Furthermore, k is a regular function symbol which indicates the current continuation.

$$H ::= \{ \text{op} \mapsto k \} \quad \text{Effect continuation} \\ | H \cup H' \quad \text{Union of handlers}$$

Figure 6: Handler Continuations

$$\begin{aligned} \llbracket \mathbf{x} \rrbracket_k^H &= \mathbf{x} \\ \llbracket v \rrbracket_k^H &= v \quad \text{where } v \in \{(), \#, \text{ff}\} \cup \mathbb{Z} \\ \llbracket \text{return } v \rrbracket_k^H &= k \llbracket v \rrbracket_k^H \\ \llbracket \text{if } v \{ c_1 \} \text{ else } \{ c_2 \} \rrbracket_k^H &= \text{if } v \{ \llbracket c_1 \rrbracket_k^H \} \text{ else } \{ \llbracket c_2 \rrbracket_k^H \} \\ \llbracket \text{let } \mathbf{x} = c_1 \text{ in } c_2 \rrbracket_k^H &= ??? \\ \llbracket \text{perform } \text{op}(v) \rrbracket_k^H &= k_{\text{op}} v \quad \text{where } \{ \text{op} \mapsto k_{\text{op}} \} \in H \\ \left[\left[\text{with } \left\{ \begin{array}{l} \text{return } \mathbf{x} \mapsto c_r, \\ \text{op}_i(\mathbf{x}, \mathbf{k}) \mapsto c_i, \end{array} \right. \text{ for } 0 \leq i < n \right\} \text{ handle } c \right] \right]_k^H &= \text{let } \mathbf{k}_r = \lambda \mathbf{x}. \llbracket c_r \rrbracket_k^H \text{ in} \\ &\quad \text{let } \mathbf{k}_{\text{op}_i} = \lambda \mathbf{x}. \lambda \mathbf{k}. \llbracket c_i \rrbracket_k^H \text{ in} \quad \text{for } 0 \leq i < n \\ &\quad \llbracket c \rrbracket_{\mathbf{k}_r}^{H \cup \{ \text{op}_i \mapsto \mathbf{k}_{\text{op}_i} \mid 0 \leq i < n \}} \end{aligned}$$

Figure 7: Compilation of effects to CPS.

Question #13

Translate c_{ex} to CPS, and execute it with an initial cell set to true.

Question #14

In the translation of λ_{eff}^- , do the meta handlers H appear at runtime? Explain their use.

Question #15

Give a rule for $\llbracket \text{let } x = c_1 \text{ in } c_2 \rrbracket_k^H$ without using let. Where should k be used? Justify the correctness informally.

Question #16

State (without proving) the correctness criterion for the compilation function $\llbracket e \rrbracket_k^H$.

Hint: this doesn't talk about soundness of the source language!

Question #17

With λ_{eff}^- we have neglected lambda-expressions. What is the problem with extending our compilation function to lambda-expressions? Give an example to illustrate the issue.

3.2 Implementing lambdas with dynamic handlers

As a first approach to handle lambda-expressions, we now introduce a new class of values: dynamic handlers! Similarly to before, handlers are maps from operations to continuations. But additionally, they can be bound to variables and composed. For instance,

$$\lambda x. (\lambda H. \text{let } \mathbf{H}' = \mathbf{H} \cup \{incr \mapsto c_{incr}\} \text{ in } \mathbf{H}'(get) + \mathbf{H}'(incr))$$

The extended grammars and rules are given below. The idea is that all lambda-expressions now take two additional arguments: the current continuation, and the handler continuations.

$$\begin{array}{l} v ::= \dots \mid H \text{ Dynamic Handler Continuations} \\ H ::= \dots \mid \mathbf{H} \text{ Handler Variables} \end{array}$$

Figure 8: Dynamic Handler Continuations

$$\begin{array}{l} \llbracket \lambda x. c \rrbracket_k^H = \lambda x. \lambda k'. \lambda H'. \llbracket c \rrbracket_{k'}^{H'} \\ \llbracket v_1 v_2 \rrbracket_k^H = ??? \\ \llbracket \text{perform } op(v) \rrbracket_k^H = ??? \end{array}$$

Figure 9: Compilation rules, extended for lambda-expressions.

Question #18

Give the rule for application. Justify it informally.

Question #19

We have extended the notion of handler to be used dynamically.

Provide the new evaluation rules and extend the rule for perform. What changes compared to the previous rule?

Question #20

Translate `let f = $\lambda().\text{flip}$ in with h_{cell} handle $f()$` with this new technique.

Question #21

Prove your correctness theorem only on the “with h handle c ” case.

Question #22

Consider the following term.

```
let f = λx.(with {rand() ↦ c_rand} handle perform set(x);;perform rand()) in
with h_cell handle f 12
```

Translate it, and execute it. How many time is the operation of handlers union executed ?

Question #23

Propose a code simplification to reduce the number of handlers union.

Question #24

Consider a program without any effects. What performance penalties does our compilation to CPS incur ?

3.3 Optimising lambdas with an effect analysis

We now want to analyze a closure to know which effect it can raise, and use this information to optimise the compilation of lambdas.

Let us consider a judgement $\text{Effects}(c) = \text{op}_0, \dots, \text{op}_{n-1}$ which indicates that the execution of c can only perform effects amongst $\text{op}_0, \dots, \text{op}_{n-1}$. We also consider $\text{Effects}(v) = \frac{\text{op}_0, \dots, \text{op}_{n-1}}{\rightarrow}$ which indicates that v is a function which, when called, can only perform effects amongst $\text{op}_0, \dots, \text{op}_{n-1}$ (no matter what argument the function is called on).

For instance $\text{Effects}(\text{flip}) = \{\text{get}; \text{put}\}$

Question #25

Propose new compilation rules using this judgement to avoid translating effect-less functions to CPS. Justify their correctness.

Question #26 (Difficult)

Propose compilation rules that avoid the use of dynamic handler completely. Justify their correctness.

Question #27 (Difficult)

Open question

What are the techniques that could be used to implement $\text{Effects}(c) = \dots$. Propose ideas.