

Koka: Programming with Row-polymorphic Effect Types

Daan Leijen

Microsoft Research,
daan@microsoft.com

We propose a programming model where effects are treated in a disciplined way, and where the potential side-effects of a function are apparent in its type signature. The type and effect of expressions can also be inferred automatically, and we describe a polymorphic type inference system based on Hindley-Milner style inference. A novel feature is that we support polymorphic effects through row-polymorphism using duplicate labels. Moreover, we show that our effects are not just syntactic labels but have a deep semantic connection to the program. For example, if an expression can be typed without an *exn* effect, then it will never throw an unhandled exception. Similar to Haskell's *runST* we show how we can safely encapsulate stateful operations. Through the state effect, we can also safely combine state with let-polymorphism without needing either imperative type variables or a syntactic value restriction. Finally, our system is implemented fully in a new language called Koka¹ and has been used successfully on various small to medium-sized sample programs ranging from a Markdown processor to a tier-splitting chat application. You can try out Koka live at www.rise4fun.com/koka/tutorial.

1. Introduction

We propose a programming model where effects are a part of the type signature of a function. Currently, types only tell us something about the input and output value of a function but say nothing about all *other* behaviors; for example, if the function writes to the console or can throw an exception. In our system, the squaring function:

```
function sqr(x : int) { x * x }
```

will get the type:

```
sqr : int → total int
```

signifying that *sqr* has no side effect at all and behaves as a total function from integers to integers. If we add a *print* statement though:

```
function sqr(x : int) { print(x); x * x }
```

the (inferred) type indicates that *sqr* has an input-output (*io*) effect:

```
sqr : int → io int
```

Note that there was no need to change the original function nor to promote the expression *x*x* into the *io* effect. One of our goals is to make effects convenient for the programmer, so we automatically combine

¹Koka means 'effect' or 'effective' in Japanese.

effects. In particular, this makes it convenient for the programmer to use precise effects without having to insert coercions.

There have been many proposals for effects systems in the past [2,8,22,23,26,31,34,36,40]. However, many such systems suffer from being syntactical in nature (i.e. effects are just labels), or by being quite restricted, for example being monomorphic or applied to a very specific set of effects. Some of the more general systems suffer from having complicated effect types, especially in a polymorphic setting that generally requires sub-effect constraints.

Our main contribution in this paper is the novel combination of existing techniques into the design of a practical ML-like language with strong effect typing. In addition, many known techniques are applied in a novel way: ranging from effect types as rows with duplicate labels to the safety of *runST* in a strict setting. In particular:

- We describe a novel effect system based on row polymorphism which allows *duplicated* effects. This simplifies the effect types and provides natural types to effect elimination forms, like catching exceptions.
- The effect types are not just syntactic labels but they have a deep semantic connection to the program (Section 6). For example, we can prove that if an expression that can be typed without an *exn* effect, then it will never throw an unhandled exception; or if an expression can be typed without a *div* effect, then it always terminates.
- The interaction between polymorphism and mutable state is fraught with danger. We show that by modeling state as an effect we can safely combine mutability with let-polymorphism without needing either imperative type variables, nor a syntactic value restriction. Moreover, we can safely encapsulate local state operations and we prove that such encapsulation is sound where no references or stateful behavior can escape the encapsulation scope.
The interaction between divergence and higher-order mutable state is also tricky. Again, we show how explicit heap effects allow us to safely infer whether stateful operations may diverge.
- We have an extensive experience with the type system within the Koka language. The Koka language fully implements the effect types as described in this paper and we have used it successfully in various small to medium sized code examples ranging from a fully compliant Markdown text processor to a tier-splitting chat application (Section 2.8).

2. Overview

Types tell us about the behavior of functions. For example, if suppose we have the type of a function *foo* in ML with type $int \rightarrow int$. We know that *foo* is well defined on inputs of type *int* and returns values of type *int*. But that is only part of the story, the type tells us nothing about all *other* behaviors: i.e. if it accesses the file system perhaps, or throws exceptions, or never returns a result at all.

Even ‘pure’ functional languages like Haskell do not fare much better at this. Suppose our function has the Haskell type $Int \rightarrow Int$. Even though we know now there is no arbitrary side-effect, we still do not know whether this function terminates or can throw exceptions. Due to laziness, we do not even know if the result itself, when demanded, will raise an exception or diverge; i.e. even a simple transformation like $x * 0$ to 0 is not possible under Haskell’s notion of purity.

In essence, in both ML and Haskell the types are not precise enough to describe many aspects of the static behavior of a program. In the Haskell case, the real type is more like $(Int_{\perp} \rightarrow Int_{\perp})_{\perp}$ while the type signature of the ML program should really include that any kind of side-effect might happen.

Functional programming has been done wrong! We believe it is essential for types to include potential behaviors like divergence, exceptions, or statefulness. Being able to reason about these aspects is crucial in many domains, including safe parallel execution, optimization, query embedding, tier-splitting, etc.

2.1. Effect types

To address the previous problems, we take a fresh look at programming with side-effects in the context of a new language called Koka [18,19].

Like ML, Koka has strict semantics where arguments are evaluated before calling a function. This implies that an expression with type *int* can really be modeled semantically as an integer (and not as a delayed computation that can potentially diverge or raise an exception).

As a consequence, the *only point where side effects can occur is during function application*. We write function types as $(\tau_1, \dots, \tau_n) \rightarrow \varepsilon \tau$ to denote that a function takes arguments of type τ_1 to τ_n , and returns a value of type τ with a potential side effect ε . As apparent from the type, functions need to be fully applied and are not curried. This is to make it immediately apparent where side effects can occur. For example, in ML, an expression like $f\ x\ y$ can have side effects at different points depending on the arity of the function f . In our system this is immediately apparent, as one writes either $f(x, y)$ or $(f(x))(y)$.

2.2. Basic effects

The effects in our system are extensible, but the basic effects defined in Koka are *total*, *exn*, *div*, *ndet*, *alloc* $\langle h \rangle$, *read* $\langle h \rangle$, *write* $\langle h \rangle$, and *io*. Of course *total* is not really an effect but signifies the absence of any effect and is assigned to pure mathematical functions. When a function can throw an exception, it gets the *exn* effect. Potential divergence or non-termination is signified by the *div* effect. Currently, Koka uses a simple termination analysis based on inductive data types to assign this effect to recursive functions. Non-deterministic functions get the *ndet* effect. The effects *alloc* $\langle h \rangle$, *read* $\langle h \rangle$ and *write* $\langle h \rangle$ are used for stateful functions over a heap h . Finally *io* is used for functions that do any input/output operations. Here are some type signatures of common functions in Koka:

```

random   : ()  $\rightarrow$  ndet double
print    : string  $\rightarrow$  io ()
error    :  $\forall \alpha. \textit{string} \rightarrow \textit{exn } \alpha$ 
(:=)     :  $\forall \alpha. (\textit{ref} \langle h, a \rangle, a) \rightarrow \textit{write} \langle h \rangle ()$ 

```

Note that we use angled brackets to denote type application as usual in languages like C# or Scala. We also use angled brackets to denote a *row* of effects. For example, the program:

```
function sqr(x : int) {error("hi"); sqr(x); x * x}
```

will get the type

```
sqr : int  $\rightarrow$   $\langle \textit{exn}, \textit{div} \rangle$  int
```

where we combined the two basic effects *exn* and *div* into a row of effects $\langle \textit{exn}, \textit{div} \rangle$. The combination of the exception and divergence effect corresponds exactly to Haskell's notion of purity, and we call this

effect *pure*. Common type aliases are:

```
alias total = ⟨⟩
alias pure  = ⟨exn, div⟩
alias st⟨h⟩ = ⟨alloc⟨h⟩, read⟨h⟩, write⟨h⟩⟩
alias io    = ⟨st⟨ioheap⟩, pure, ndet⟩
```

This hierarchy is clearly inspired by Haskell’s standard monads and we use this as a starting point for more refined effects which we hope to explore in Koka. For example, blocking, client/server side effects, reversible operations, etc.

2.3. Polymorphic effects

Often, the effect of a function is determined by the effects of functions passed to it. For example, the *map* function which maps a function over all elements of a list will have the type:

$$\text{map} : \forall \alpha \beta \mu. (\text{list} \langle \alpha \rangle, \beta \rightarrow \mu \beta) \rightarrow \mu \text{list} \langle \beta \rangle$$

where the effect of the *map* function itself is completely determined by the effect of its argument. In this case, a simple and obvious type is assigned to *map*, but one can quickly create more complex examples where the type may not be obvious at first. Consider the following program:

```
function foo(f, g) { f(); g(); error("hr") }
```

Clearly, the effect of *foo* is a combination of the effects of *f* and *g*, and the *exn* effect. One possible design choice is to have a \cup operation on effect types, and write the type of *foo* as:

$$\forall \mu_1 \mu_2. ((\rightarrow \mu_1 ()) , (\rightarrow \mu_2 ())) \rightarrow (\mu_1 \cup \mu_2 \cup \text{exn}) ()$$

Unfortunately, this quickly gets us in trouble during type inference: unification can lead to constraints of the form $\mu_1 \cup \mu_2 \sim \mu_3 \cup \mu_4$ which cannot be solved uniquely and must become part of the type language. Another design choice is to introduce subtyping over effects and write the type of *foo* as:

$$\forall \mu_1 \mu_2 \mu_3. (\mu_1 \leq \mu_3, \mu_2 \leq \mu_3, \langle \text{exn} \rangle \leq \mu_3) \Rightarrow ((\rightarrow \mu_1 ()) , (\rightarrow \mu_2 ())) \rightarrow \mu_3 ()$$

This is the choice made in an earlier version of Koka described as a technical report [36]. However, in our experience with that system in practice we felt the constraints often became quite complex and the combination of polymorphism with subtyping can make type inference undecidable.

The approach we advocate in this paper and which is adopted by Koka is the use of row-polymorphism on effects. Row polymorphism is well understood and used for many inference systems for record calculi [7,17,21,30,32,33]. We use the notation $\langle l \mid \mu \rangle$ to extend an effect row μ with an effect constant l . Rows can now have two forms, either a *closed* effect $\langle \text{exn}, \text{div} \rangle$, or an *open* effect ending in an effect variable $\langle \text{exn}, \text{div} \mid \mu \rangle$. Using an open effect, our system infers the following type for *foo*:

$$\text{foo} : \forall \mu. ((\rightarrow \langle \text{exn} \mid \mu \rangle ()) , (\rightarrow \langle \text{exn} \mid \mu \rangle ())) \rightarrow \langle \text{exn} \mid \mu \rangle ()$$

The reader may worry at this point that the row polymorphic type is more restrictive than the earlier type using subtype constraints: indeed, the row polymorphic type requires that each function argument now has the same effect $\langle \text{exn} \mid \mu \rangle$. However, in a calling context $\text{foo}(f, g)$ our system ensures that we always infer a polymorphic open effect for each expression f and g . For example, $f : () \rightarrow \langle \text{exn} \mid \mu_1 \rangle ()$ and $g : () \rightarrow \langle \text{div} \mid \mu_2 \rangle ()$. This allows the types $\langle \text{exn} \mid \mu_1 \rangle$ and $\langle \text{div} \mid \mu_2 \rangle$ to unify into a common type $\langle \text{exn}, \text{div} \mid \mu_3 \rangle$ such that they can be applied to *foo*, resulting in an inferred effect $\langle \text{exn}, \text{div} \mid \mu_3 \rangle$ for $\text{foo}(f, g)$.

2.4. Duplicate effects

Our effect rows differ in an important way from the usual approaches in that effect labels can be duplicated, i.e. $\langle \text{exn}, \text{exn} \rangle \not\equiv \langle \text{exn} \rangle$ (1). This was first described by Leijen [17] where this was used to enable scoped labels in record types. Enabling duplicate labels is crucial for our approach. First of all, it enables principal unification without needing extra constraints and secondly, it enables us to give precise types to effect elimination forms (like catching exceptions).

In particular, during unification we may end up with constraints of the form $\langle \text{exn} | \mu \rangle \sim \langle \text{exn} \rangle$. With regular row-polymorphism, such constraint can have multiple solutions, namely $\mu = \langle \rangle$ or $\mu = \langle \text{exn} \rangle$. This was first observed by Wand [41] in the context of records. Usually, this problem is fixed by either introducing *lacks* constraints [7] or polymorphic presence and absence flags on each label [29] (as used by Lindley and Cheney [21] for an effect system in the context of database queries). With rows allowing duplicate labels, we avoid additional machinery since in our case $\mu = \langle \rangle$ is the only solution to the above constraint (due to (1)).

Moreover, duplicate labels make it easy to give types to effect elimination forms. For example, catching effects removes the *exn* effect:

$$\text{catch} : \forall \alpha \mu. (() \rightarrow \langle \text{exn} | \mu \rangle \alpha, \text{exception} \rightarrow \mu \alpha) \rightarrow \mu \alpha$$

Here we assume that *catch* takes two functions, the action and the exception handler that takes as an argument the thrown *exception*. Here, the *exn* effect of the action is discarded in the final effect μ since all exceptions are handled by the handler. But of course, the handler can itself throw an exception and have an *exn* effect itself. In that case μ will unify with a type of the form $\langle \text{exn} | \mu' \rangle$ giving action the effect $\langle \text{exn} | \text{exn} | \mu' \rangle$ where *exn* occurs duplicated, which gives us exactly the right behavior. Note that with *lacks* constraints we would not be able to type this example because there would be a $\text{exn} \notin \mu$ constraint. We can type this example using flags but the type would arguably be more complex with a polymorphic presence/absence flag on the *exn* label in the result effect, something like:

$$\text{catch} : \forall \mu \alpha \phi. (() \rightarrow \langle \text{exn}_\bullet | \mu \rangle \alpha, \text{exception} \rightarrow \langle \text{exn}_\phi | \mu \rangle \alpha) \rightarrow \langle \text{exn}_\phi | \mu \rangle \alpha$$

There is one situation where an approach with flags is more expressive though: with flags one can state specifically that a certain effect must be absent. This is used for example in the effect system by Lindley and Cheney [21] to enforce that database queries never have the *wild* effect (*io*). In our setting we can only enforce absence of an effect by explicitly listing a closed row of the allowed effects which is less modular. In our current experience this has not yet proven to be a problem in practice though.

2.4.1. Injection

There are some situations where having duplicate effects is quite different from other approaches though. Intuitively, we can do a monadic translation of a Koka program where effect types get translated to a sequence of monad transformers. Under such interpretation, duplicate effects would have real semantic significance. For example, we could provide an *injection* operation for exceptions:

$$\text{inject} : \forall \mu \alpha. (() \rightarrow \langle \text{exn} | \mu \rangle \alpha) \rightarrow (() \rightarrow \langle \text{exn} | \text{exn} | \mu \rangle \alpha)$$

Semantically, it would inject an extra exception layer, such that a *catch* operation would only catch exceptions raised from the outer *exn*, while passing the inner injected exceptions through. Internally, one

can implement this by maintaining level numbers on thrown exceptions – increasing them on an *inject* and decreasing them on a *catch* (and only catching level 0 exceptions).

Now, suppose we have some library code whose exceptions we do not want to handle, but we do want to handle the exceptions in our own code. In that case, we could write something like:

```
catch(function() {
  ... my code ...
  x = inject(... library code ... )()
  ... my code ...
  inject(... more library code ... )()
}, handler)
```

In the example, all exceptions in ‘my code’ are caught, while exceptions raised in the library code are only caught by an outer exception handler. For this article though, we will not further formalize *inject* but only describe the core calculus.

2.5. Heap effects

One of the most useful side-effects is of course mutable state. Here is an example where we give a linear version of the fibonacci function using imperative updates:

```
function fib(n : int) {
  val x = ref(0); val y = ref(1)
  repeat(n) {
    val y0 = !y; y := !x+!y; x := y0
  }
  !x
}
```

Here x and y are bound to freshly allocated references of type $ref\langle h, int \rangle$. The operator $(!)$ dereferences a reference while the operator $(:=)$ is used for assignment to references. Due to the reading and writing of x and y of type $ref\langle h, int \rangle$, the effect inferred for the body of the function is $st\langle h \rangle$ for some heap h :

$$fib : \forall h. int \rightarrow st\langle h \rangle int$$

However, we can of course consider the function fib to be total: for any input, it always returns the same output since the heap h cannot be modified or observed from outside this function. In particular, we can safely remove the effect $st\langle h \rangle$ whenever the function is polymorphic in the heap h and where h is not among the free type variables of argument types or result type. This notion corresponds directly to the use of the higher-ranked $runST$ function in Haskell [28] (which we will call just run):

$$run : \forall \mu \alpha. (\forall h. () \rightarrow \langle st\langle h \rangle \mid \mu \rangle \alpha) \rightarrow \mu \alpha$$

Koka will automatically insert a run wrapper at generalization points if it can be applied, and infers a total type for the above fibonacci function:

$$fib : int \rightarrow total\ int$$

Again, using row polymorphism is quite natural to express in the type of run where the $st\langle h \rangle$ effect can be dismissed.

One complex example from a type inference perspective where we applied Koka, is the Garsia-Wachs algorithm as described by Filliâtre [6]. The given algorithm was originally written in ML and uses updateable references in the leaf nodes of the trees to achieve efficiency comparable to the reference C implementation. However, Filliâtre remarks that these side effects are local and not observable to any caller. We implemented Filliâtre’s algorithm in Koka and our system correctly inferred that the state effect can be discarded and assigned a pure effect to the Garsia-Wachs algorithm [18].

2.6. Heap safety

Combining polymorphism and imperative state is fraught with difficulty and requires great care. In particular, *let*-polymorphism may lead to unsoundness if references can be given a polymorphic type. A classical example from ML is:

```
let r = ref [] in (r := [true], !r + 1)
```

Here, we let bind r to a reference with type $\forall\alpha. \text{ref}\langle \text{list}\langle\alpha\rangle\rangle$. The problem is that this type can instantiate later to both a reference to an integer list and a boolean list. Intuitively, the problem is that the first binding of r generalized over type variables that are actually free in the heap. The ML language considered many solutions to prevent this from happening, ranging from imperative type variables [37] to the current syntactic value restriction, where only value expressions can be generalized.

In our system, no such tricks are necessary. Using the effect types, we restrict generalization to expressions that are total, and we reject the ML example since we will not generalize over the type of r since it has an $\text{alloc}\langle h\rangle$ effect. We prove in Section 5.1 that our approach is semantically sound. In contrast to the value restriction, we can still generalize over any expression that is not stateful regardless of its syntactic form.

The addition of *run* adds further requirements where we must ensure that encapsulated stateful computations truly behave like a pure function and do not ‘leak’ the state. For example, it would be unsound to let a reference escape its encapsulation:

```
run(function(){ ref(1) })
```

or to encapsulate a computation where its effects can still be observed.

We prove in Section 5.2 that well-typed terms never exhibit such behavior. To our knowledge we are the first to prove this formally for a strict language in combination with exceptions and divergence. A similar result is by Launchbury and Sabry [16] where they prove heap safety of the Haskell’s ST monad in the context of a lazy store with lazy evaluation.

2.7. Divergence

Koka uses a simple termination checker (based on [1]) to assign the divergence effect to potentially non-terminating functions. To do this safely, Koka has three kinds of data types, inductive, co-inductive, and arbitrary recursive data types. In particular, we restrict (co)inductive data types such that the type itself cannot occur in a negative position. Any function that matches on an arbitrary recursive data type is assumed to be potentially divergent since one can encode the Y combinator using such data type and write a non-terminating function that is not syntactically recursive.

Recursively defined functions should of course include the divergence effect in general. However, if the termination checker finds that each recursive call decreases the size of an inductive data type (or is

productive for a co- inductive data type), then we do not assign the divergent effect. The current analysis is quite limited and syntactically fragile but seems to work well enough in practice (Section 2.8). For our purpose, we prefer a predictable analysis with clear rules.

However, in combination with higher-order mutable state, we can still define functions that are not syntactically recursive, but fail to terminate. Consider ‘Landin’s knot’:

```
function diverge() {
  val r := ref(id)
  function foo() { (!r)() }
  r := foo
  foo()
}
```

In this function, we first create a reference r initialized with the identify function. Next we define a local function foo which calls the function in r . Then we assign foo itself to r and call foo , which will now never terminate even though there is no syntactic recursion.

How can we infer in general that $diverge$ must include the div effect? It turns out that in essence reading from the heap may result in divergence. A conservative approach would be to assign the div effect to the type of read (!). For simplicity, this is what we will do in the formal development.

But in Koka we use a more sophisticated approach. In order to cause divergence, we actually need to read a function from the heap which accesses the heap itself. Fortunately, our effect system makes this behavior already apparent in the inferred types! – in our example, the effect of foo contains $read\langle h \rangle$, which is being stored in a reference in the same heap of type $ref\langle h, () \rightarrow read\langle h \rangle () \rangle$.

The trick is now that we generate a type constraint $hdiv\langle h, \tau, \varepsilon \rangle$ for every heap read that keeps track of heap type h , the type of the value that was read τ and the current effect ε . The constraint $hdiv\langle h, \tau, \varepsilon \rangle$ expresses that if $h \in \text{ftv}(\tau)$ then the effect ε must include divergence. In particular, this constraint is fine-grained enough that any reading of a non-function type, or non-stateful functions will never cause divergence (and we can dismiss the constraint) The drawback is that if τ is polymorphic at generalization time, we need to keep the constraint around (as we cannot decide at that point whether h will ever be in $\text{ftv}(\tau)$), which in turn means we need to use a system of qualified types [12]. Currently this is not fully implemented yet in Koka, and if at generalization time we cannot guarantee τ will never contain a reference to the heap h , we conservatively assume that the function may diverge.

2.8. Koka in practice

When designing a new type system it is always a question how well it will work in practice: does it infer the types you expect? Do the types become too complicated? Is the termination checker strong enough? etc. We have implemented the effect inference and various extensions in the Koka language which is freely available on the web [18]. The Koka system currently has a JavaScript backend and can generate code that runs in NodeJS or a web page. We have written many small to medium sized samples to see how well the system works in practice.

Markdown One application is a fully compliant Markdown text processor. This program consists of three phases where it first parses block elements, performs block analysis, collects link definitions, numbers sections, and finally renders the inline elements in each block. The program passes the full Markdown test suite. In fact, this article has been written itself as a Madoko program.

Remarkably, almost all functions are inferred to be *total*, and only a handful of driver functions perform side effects, like reading input files. For efficiency though, many internal functions use local state. For example, when rendering all inline elements in a block, we use a local mutable string builder (of type *builder* $\langle h \rangle$) to build the result string in constant time (actual Koka code):

```
function formatInline( ctx : inlineCtx, txt : string ) : string {
  formatAcc( ctx, builder(), txt )
}

function formatAcc( ctx : inlineCtx, acc : builder<h>, txt : string ) : st<h> string {
  if (txt == "") return acc.string
  val (s, next) = matchRules( ctx.grammar, ctx, txt, id )
  formatAcc( ctx, acc.append(s), txt.substr1(next) )
}
```

Note how *formatAcc* is stateful due to the calls to the *append* and *string* methods of the string builder *acc*, but the outer function *formatInline* is inferred to be *total* since Koka can automatically apply the *run* function and encapsulate the state: indeed it is not observable if we use a mutable string builder internally or not. This pattern also occurs for example in the block analysis phase where we use a mutable hashmap to build the dictionary of link definitions.

Safe tier-splitting Most of the HTML5 DOM and the Node API's are available in Koka which allows us to write more substantial programs and evaluate the effect inference system in practice. We use the effect *dom* for functions that may have any side effect through a DOM call.

On the web, many programs are split in a server and client part communicating with each other. It is advantageous to write both the client and server part as one program since that enables us to share one common type definition for the data they exchange. Also, their interaction will be more apparent, and they can share common functionality, like date parsing, to ensure that both parts behave consistently.

Safely splitting a program into a server and client part is difficult though. For example, the client code may call a library function that itself calls a function that can only be run on the server (like writing to a log file), or the other way around. Moreover, if the client and server part both access a shared global variable (or both call a library function that uses an internal global variable) then we cannot split this code anymore.

The Koka effect types tackle both problems and enable fully safe tier splitting. Our main tier splitting function has the following (simplified) type signature:

```
function tiersplit(
  serverPart : () → server ( (α → server ()) → server (β → server ()) ),
  clientPart : (β → client ()) → client (α → client ())
) : io ()
```

where the *server* and *client* effects are defined as:

```
alias server = io
alias client = <dom, div>
```

The *tiersplit* function takes a server and client function and sets up a socket connection. On the server it will call the server part function which can initialize. Now, both the client and server part can be called

| | | |
|-----------------|--|---|
| kinds | $\kappa ::= * e k h$ | values, effect rows, effect constants, heaps |
| | $ (\kappa_1, \dots, \kappa_n) \rightarrow \kappa$ | type constructor |
| types | $\tau^k ::= \alpha^k$ | type variable (using μ for effects, ξ for heaps) |
| | $ c^\kappa$ | type constant |
| | $ c^{\kappa_0} \langle \tau_1^{\kappa_1}, \dots, \tau_n^{\kappa_n} \rangle$ | $\kappa_0 = (\kappa_1, \dots, \kappa_n) \rightarrow \kappa$ |
| schemes | $\sigma ::= \forall \alpha_1 \dots \alpha_n. \tau^*$ | |
| constants | $()$ | $:: *$ unit type |
| | $(_ \rightarrow _)$ | $:: (*, e, *) \rightarrow *$ functions |
| | $\langle \rangle$ | $:: e$ empty effect |
| | $\langle _ _ \rangle$ | $:: (k, e) \rightarrow e$ effect extension |
| | <i>ref</i> | $:: (h, *) \rightarrow *$ references |
| | <i>exn, div</i> | $:: k$ partial, divergent |
| | <i>st</i> | $:: h \rightarrow k$ stateful |
| syntactic sugar | $\tau_1 \rightarrow \tau_2$ | $= \tau_1 \rightarrow \langle \rangle \tau_2$ |
| | $\langle l_1, \dots, l_n \varepsilon \rangle$ | $= \langle l_1 \dots \langle l_n \varepsilon \rangle \dots \rangle$ |
| | $\langle l_1, \dots, l_n \rangle$ | $= \langle l_1, \dots, l_n \langle \rangle \rangle$ |

Figure 1. Syntax of types and kinds. An extra restriction is that effect constants cannot be type variables, i.e. α^k is illegal.

for each fresh connection where *tiersplit* supplies a *send* function that takes a message of type α for client messages, and β for server messages. Both the client and server part return a fresh ‘connection’ function that handles incoming messages from the server or client respectively. Note how this type guarantees that messages sent to the client, and messages handled by the client, are both of type α , while for the server messages they will be β . Furthermore, because the effect types for *server* and *client* are closed, the client and server part are only able to call functions available for the client or server respectively.

Finally, the Koka effect system also prevents accidental sharing of global state by the client and server part. Both the client and server can use state that is contained in their handler. In that case the $st\langle h \rangle$ effect will be inferred, and discarded because h will generalize. However, if either function tries to access a shared variable in an outer scope, then the h will *not* generalize (because the variable will have type $ref\langle h, a \rangle$ and therefore h is not free in the environment), in which case the inferred $st\langle h \rangle$ effect cannot be removed. Again, this will lead to a unification failure and the program will be statically rejected.

3. The type system

In this section we are going to give a formal definition of our polymorphic effect system for a small core-calculus that captures the essence of Koka. We call this λ^k . Figure 1 defines the syntax of types. The well-formedness of types τ is guaranteed by a simple kind system. We put the kind κ of a type τ in superscript, as τ^k . We have the usual kind $*$ and \rightarrow , but also kinds for effect rows (e), effect constants (k), and heaps (h). Often the kind of a type is immediately apparent or not relevant, and most of the time we will not denote the kind to reduce clutter, and just write plain types τ . For clarity, we are using α for regular type variables, μ for effect type variables, and ξ for heap type variables.

$$\begin{array}{c}
\text{(EQ-REFL)} \ \varepsilon \equiv \varepsilon \quad \text{(EQ-TRANS)} \ \frac{\varepsilon_1 \equiv \varepsilon_2 \quad \varepsilon_2 \equiv \varepsilon_3}{\varepsilon_1 \equiv \varepsilon_3} \\
\\
\text{(EQ-HEAD)} \ \frac{\varepsilon_1 \equiv \varepsilon_2}{\langle l | \varepsilon_1 \rangle \equiv \langle l | \varepsilon_2 \rangle} \quad \text{(EQ-SWAP)} \ \frac{l_1 \neq l_2}{\langle l_1 | \langle l_2 | \varepsilon \rangle \rangle \equiv \langle l_2 | \langle l_1 | \varepsilon \rangle \rangle} \\
\\
\text{(UNEQ-LAB)} \ \frac{c \neq c'}{c \langle \tau_1, \dots, \tau_n \rangle \neq c' \langle \tau'_1, \dots, \tau'_n \rangle}
\end{array}$$

Figure 2. Effect equivalence.

Effect types are defined as a row of effect labels l . Such effect row is either empty $\langle \rangle$, a polymorphic effect variable μ , or an extension of an effect row ε with an effect constant l , written as $\langle l | \varepsilon \rangle$. The effect constants can be anything that is interesting to our language. For our purposes we will restrict the constants to exceptions (*exn*), divergence (*div*), and heap operations (*st*). It is no problem to generalize this to the more fine-grained hierarchy of Koka but this simplifies the presentation and proofs. The kind system ensures that an effect is always either a *closed effect* of the form $\langle l_1, \dots, l_n \rangle$, or an *open effect* of the form $\langle l_1, \dots, l_n | \mu \rangle$.

Figure 2 defines an equivalence relation (\equiv) between effect types where we consider effects equivalent regardless of the order of the effect constants. In contrast to many record calculi [7,29,32] effect rows *do* allow duplicate labels where an effect $\langle \text{exn}, \text{exn} \rangle$ is allowed (and not equal to the effect $\langle \text{exn} \rangle$). The definition of effect equality is essentially the same as for scoped labels [17] where we ignore the type components. Note that for rule (EQ-SWAP) we use the rule (UNEQ-LAB) to compare effect constants where the type arguments are not taken into account: intuitively we consider the effect constants as the ‘labels’ of an effect record. Most constants compare directly. The only exception in our system is the state effect where $st \langle h_1 \rangle \neq st \langle h_2 \rangle$ does *not* hold even if $h_1 \neq h_2$.

Using effect equality, we define $l \in \varepsilon$ iff $e \equiv \langle l | \varepsilon' \rangle$ for some ε' .

3.1. Type rules

Figure 3 defines the formal type rules of our effect system. The rules are defined over a small expression calculus:

$$\begin{array}{ll}
e ::= x \mid p \mid e_1 e_2 \mid \lambda x. e & \text{(variables, primitives, applications, functions)} \\
\mid x \leftarrow e_1; e_2 \mid \text{let } x = e_1 \text{ in } e_2 & \text{(sequence and let bindings)} \\
\mid \text{catch } e_1 e_2 \mid \text{run } e & \text{(catch exceptions and isolate state)}
\end{array}$$

$$p ::= () \mid \text{fix} \mid \text{throw} \mid \text{new} \mid (!) \mid (:=) \quad \text{(primitives)}$$

This expression syntax is meant as a surface syntax, but when we discuss the semantics of the calculus, we will refine and extend the syntax further (see Figure 4). We use the bind expression $x \leftarrow e_1; e_2$ for a monomorphic binding of a variable x to an expression e_1 (which is really just syntactic sugar for the application $(\lambda x. e_2) e_1$). We write $e_1; e_2$ as a shorthand for the expression $x \leftarrow e_1; e_2$ where $x \notin \text{fv}(e_2)$. We have added *run* and *catch* as special expressions since this simplifies the presentation where we can give direct type rules for them. Also, we simplified both *catch* and *throw* by limiting the exception type to the unit type $()$.

The type rules are stated under a type environment Γ which maps variables to types. An environment can be extended using a comma. If Γ' is equal to $\Gamma, x : \sigma$ then $\Gamma'(x) = \sigma$ and $\Gamma'(y) = \Gamma(y)$ for any $y \neq x$.

$$\begin{array}{c}
\text{(VAR)} \frac{\Gamma(x) = \sigma}{\Gamma \vdash x : \sigma \mid \varepsilon} \quad \text{(APP)} \frac{\Gamma \vdash e_1 : \tau_2 \rightarrow \varepsilon \tau \mid \varepsilon \quad \Gamma \vdash e_2 : \tau_2 \mid \varepsilon}{\Gamma \vdash e_1 e_2 : \tau \mid \varepsilon} \\
\\
\text{(LAM)} \frac{\Gamma, x : \tau_1 \vdash e : \tau_2 \mid \varepsilon_2}{\Gamma \vdash \lambda x. e : \tau_1 \rightarrow \varepsilon_2 \tau_2 \mid \varepsilon} \quad \text{(LET)} \frac{\Gamma \vdash e_1 : \sigma \mid \langle \rangle \quad \Gamma, x : \sigma \vdash e_2 : \tau \mid \varepsilon}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau \mid \varepsilon} \\
\\
\text{(GEN)} \frac{\Gamma \vdash e : \tau \mid \langle \rangle \quad \bar{\alpha} \notin \text{ftv}(\Gamma)}{\Gamma \vdash e : \forall \bar{\alpha}. \tau \mid \varepsilon} \quad \text{(RUN)} \frac{\Gamma \vdash e : \tau \mid \langle st \langle \xi \rangle \mid \varepsilon \rangle \quad \xi \notin \text{ftv}(\Gamma, \tau, \varepsilon)}{\Gamma \vdash \text{run } e : \tau \mid \varepsilon} \\
\\
\text{(INST)} \frac{\Gamma \vdash e : \forall \bar{\alpha}. \tau \mid \varepsilon}{\Gamma \vdash e : [\bar{\alpha} := \bar{\tau}] \tau \mid \varepsilon} \quad \text{(CATCH)} \frac{\Gamma \vdash e_1 : \tau \mid \langle \text{exn} \mid \varepsilon \rangle \quad \Gamma \vdash e_2 : () \rightarrow \varepsilon \tau \mid \varepsilon}{\Gamma \vdash \text{catch } e_1 e_2 : \tau \mid \varepsilon} \\
\\
\text{(ALLOC)} \quad \Gamma \vdash \text{ref} \quad : \tau \rightarrow \langle st \langle h \rangle \mid \varepsilon \rangle \text{ref} \langle h, \tau \rangle \mid \varepsilon' \\
\text{(READ)} \quad \Gamma \vdash (!) \quad : \text{ref} \langle h, \tau \rangle \rightarrow \langle st \langle h \rangle, \text{div} \mid \varepsilon \rangle \tau \mid \varepsilon' \\
\text{(WRITE)} \quad \Gamma \vdash (:=) \quad : (\text{ref} \langle h, \tau \rangle, \tau) \rightarrow \langle st \langle h \rangle \mid \varepsilon \rangle () \mid \varepsilon' \\
\text{(THROW)} \quad \Gamma \vdash \text{throw} \quad : () \rightarrow \langle \text{exn} \mid \varepsilon \rangle \tau \mid \varepsilon' \\
\text{(UNIT)} \quad \Gamma \vdash () \quad : () \mid \varepsilon \\
\text{(FIX)} \quad \Gamma \vdash \text{fix} \quad : ((\tau_1 \rightarrow \langle \text{div} \mid \varepsilon \rangle \tau_2) \rightarrow (\tau_1 \rightarrow \langle \text{div} \mid \varepsilon \rangle \tau_2)) \rightarrow (\tau_1 \rightarrow \langle \text{div} \mid \varepsilon \rangle \tau_2) \mid \varepsilon'
\end{array}$$

Figure 3. General type rules with effects.

A type rule of the form $\Gamma \vdash e : \sigma \mid \varepsilon$ states that under an environment Γ the expression e has type σ with an effect ε .

Most of the type rules in Figure 3 are quite standard. The rule (VAR) derives the type of a variable. The derived effect is any arbitrary effect ε . We may have expected to derive only the total effect $\langle \rangle$ since the evaluation of a variable has no effect at all (in our strict setting). However, there is no rule that lets us upgrade the final effect and instead we get to pick the final effect right away. Another way to look at this is that since the variable evaluation has no effect, we are free to assume any arbitrary effect.

The (LAM) rule is similar: the evaluation of a lambda expression is a value and has no effect and we can assume any arbitrary effect ε . Interestingly, the effect derived for the body of the lambda expression, ε_2 , shifts from the derivation on to the derived function type $\tau_1 \rightarrow \varepsilon_2 \tau_2$: indeed, calling this function and evaluating the body causes the effect ε_2 . The (APP) is also standard, and derives an effect ε requiring that its premises derive the same effect as the function effect.

Instantiation (INST) is standard and instantiates a type scheme. The generalization rule (GEN) has an interesting twist: it requires the derived effect to be total $\langle \rangle$. It turns out this is required to ensure a sound semantics as we show in Section 4. Indeed, this is essentially the equivalent of the value restriction in ML [20]. Of course, in ML effects are not inferred by the type system and the value restriction syntactically restricts the expression over which one can generalize. In our setting we can directly express that we only generalize over total expressions. The rule (LET) binds expressions with a polymorphic type scheme σ and just like (GEN) requires that the bound expression has no effect. It turns out that is still sound to allow more effects at generalization and let bindings. In particular, we can allow *exn*, *div*. However, for the formal development we will only consider the empty effect for now.

All other rules are just type rules for the primitive constants. Note that all the effects for the primitive

| | |
|--|--------------------------------------|
| $e ::= v \mid e_1 e_2 \mid \text{let } x = e_1 \text{ in } e_2$ | (values, applications, let bindings) |
| $\quad \mid \text{hp } \varphi. e \mid \text{run } e$ | (heap binding and isolation) |
| $v ::= \lambda x. e \mid \text{catch } e$ | (functions and partial catch) |
| $\quad \mid b$ | (basic value (contains no e)) |
| $b ::= x \mid c$ | (variable and constants) |
| $\quad \mid \text{fix} \mid \text{throw} \mid \text{catch}$ | (fixpoint and exceptions) |
| $\quad \mid r \mid \text{ref} \mid (!) \mid (:=) \mid (r :=)$ | (references) |
| $w ::= b \mid \text{throw } v$ | (basic value or exception) |
| $a ::= v \mid \text{throw } v \mid \text{hp } \varphi. v \mid \text{hp } \varphi. \text{throw } v$ | (answers) |
| $\varphi ::= \langle r_1 \mapsto v_1 \rangle \dots \langle r_n \mapsto v_n \rangle$ | (heap bindings) |

Figure 4. Full expression syntax

$$\begin{array}{c}
 \forall \langle r_i \mapsto v_i \rangle \in \varphi. \Gamma, \bar{\varphi}_h \vdash v_i : \tau_i \mid \langle \rangle \\
 \text{(HEAP)} \frac{\Gamma, \bar{\varphi}_h \vdash e : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle}{\Gamma \vdash \text{hp } \varphi. e : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle} \quad \text{(CONST)} \frac{\text{typeof}(c) = \sigma}{\Gamma \vdash c : \sigma \mid \varepsilon}
 \end{array}$$

Figure 5. Extra type rules for heap expressions and constants. We write $\bar{\varphi}_h$ for the conversion of a heap φ to a type environment: if φ equals $\langle r_1 \mapsto v_1, \dots, r_n \mapsto v_n \rangle$ then $\bar{\varphi}_h = r_1 : \text{ref} \langle h, \tau_1 \rangle, \dots, r_n : \text{ref} \langle h, \tau_n \rangle$ for some τ_1 to τ_n .

constants are open and can be freely chosen (just as in the (VAR) rule). This is important as it allows us to always assume more effects than induced by the operation.

Given these type rules, we can construct an efficient type inference algorithm that infers principal types and is sound and complete with respect to the type rules. This is described in detail in a separate technical report [19] (Appendix A).

4. Semantics of effects

In this section we are going to define a precise semantics for our language, and show that well-typed programs cannot go ‘wrong’. In contrast to our earlier soundness and completeness result for the type inference algorithm, the soundness proof of the type system in Hindley-Milner does not carry over easily in our setting: indeed, we are going to model many complex effects which is fraught with danger.

First, we strengthen our expression syntax by separating out value expressions v , as shown in Figure 4. We also define basic values b as values that cannot contain expressions themselves. Moreover, we added a few new expressions, namely heap bindings ($\text{hp } \varphi. e$), a partially applied catch ($\text{catch } e$), a partially applied assignments ($v :=$), and general constants (c). Also, we denote heap variables using r . An expression $\text{hp} \langle r_1 \mapsto v_1 \rangle, \dots, \langle r_n \mapsto v_n \rangle. e$ binds r_1 to r_n in v_1, \dots, v_n and e . By convention, we always require r_1 to r_n to be distinct, and consider heaps φ equal modulo alpha-renaming.

The surface language never exposes the heap binding construct $\text{hp } \varphi. e$ directly to the user but during evaluation the reductions on heap operations create heaps and use them. In order to give a type to such

| | | | |
|--------------|--|-------------------|--|
| (δ) | $c v$ | \longrightarrow | $\delta(c, v)$ if $\delta(c, v)$ is defined |
| (β) | $(\lambda x. e) v$ | \longrightarrow | $[x \mapsto v]e$ |
| (LET) | $\text{let } x = v \text{ in } e$ | \longrightarrow | $[x \mapsto v]e$ |
| (FIX) | $\text{fix } v$ | \longrightarrow | $v(\lambda x. (\text{fix } v)x)$ |
| | | | |
| (THROW) | $X[\text{throw } v]$ | \longrightarrow | $\text{throw } v$ if $X \neq []$ |
| (CATCHT) | $\text{catch } (\text{throw } v) e$ | \longrightarrow | $e v$ |
| (CATCHV) | $\text{catch } v e$ | \longrightarrow | v |
| | | | |
| (ALLOC) | $\text{ref } v$ | \longrightarrow | $\text{hp } \langle r \mapsto v \rangle . r$ |
| (READ) | $\text{hp } \varphi \langle r \mapsto v \rangle . R[!r]$ | \longrightarrow | $\text{hp } \varphi \langle r \mapsto v \rangle . R[v]$ |
| (WRITE) | $\text{hp } \varphi \langle r \mapsto v_1 \rangle . R[r := v_2]$ | \longrightarrow | $\text{hp } \varphi \langle r \mapsto v_2 \rangle . R[()]$ |
| (MERGE) | $\text{hp } \varphi_1 . \text{hp } \varphi_2 . e$ | \longrightarrow | $\text{hp } \varphi_1 \varphi_2 . e$ |
| (LIFT) | $R[\text{hp } \varphi . e]$ | \longrightarrow | $\text{hp } \varphi . R[e]$ if $R \neq []$ |
| | | | |
| (RUNL) | $\text{run } [\text{hp } \varphi .] \lambda x . e$ | \longrightarrow | $\lambda x . \text{run } ([\text{hp } \varphi .] e)$ |
| (RUNC) | $\text{run } [\text{hp } \varphi .] \text{catch } e$ | \longrightarrow | $\text{catch } (\text{run } ([\text{hp } \varphi .] e))$ |
| (RUNH) | $\text{run } [\text{hp } \varphi .] w$ | \longrightarrow | w if $\text{frv}(w) \not\cap \text{dom}(\varphi)$ |

Evaluation contexts:

| | | | |
|------------|---------|---------|-------------------------------------|
| $X ::= []$ | $ X e$ | $ v X$ | $ \text{let } x = X \text{ in } e$ |
| $R ::= []$ | $ R e$ | $ v R$ | $ \text{let } x = R \text{ in } e$ |
| $E ::= []$ | $ E e$ | $ v E$ | $ \text{let } x = E \text{ in } e$ |
| | | | $ \text{catch } E e$ |
| | | | $ \text{hp } \varphi . E$ |
| | | | $ \text{run } E$ |

Figure 6. Reduction rules and evaluation contexts.

expression, we need an extra type rule for heap bindings, given in Figure 5. Note how each heap value is typed under an environment that contains types for all bindings (much like a recursive let binding). Moreover, a heap binding induces the stateful effect $st\langle h \rangle$. In the type rule for constants we assume a function $\text{typeof}(c)$ that returns a closed type scheme for each constant.

4.1. Reductions

We can now consider primitive reductions for the various expressions as shown in Figure 6. The first four reductions are standard for the lambda calculus. To abstract away from a particular set of constants, we assume a function δ which takes a constant and a closed value to a closed value. If $\text{typeof}(c) = \forall \bar{\alpha}. \tau_1 \rightarrow \varepsilon \tau_2$, with $\theta = [\bar{\alpha} \mapsto \bar{\tau}]$ and $\cdot \vdash v : \theta \tau_1 \mid \langle \rangle$, then $\delta(c, v)$ is defined, and $\cdot \vdash \delta(c, v) : \theta \tau_2 \mid \theta \varepsilon$. The reductions β , (LET) and (FIX) are all standard.

The next three rules deal with exceptions. In particular, the rule (THROW) propagates exceptions under a context X . Since X does not include $\text{catch } e_1 e_2$, $\text{hp } \varphi . e$ or $\text{run } e$, this propagates the exception to the nearest exception handler or state block. The (CATCHT) reduction catches exceptions and passes them on to the handler. If the handler raises an exception itself, that exception will then propagate to its nearest enclosing exception handler.

The next five rules model heap reductions. Allocation creates a heap, while (!) and (:=) read and write from the a heap. Through the R context, these always operate on the nearest enclosing heap since R does not contain $\text{hp } \varphi . e$ or $\text{run } e$ expressions. The rules (LIFT) and (MERGE) let us lift heaps out of

expressions to ensure that all references can be bound in the nearest enclosing heap.

The final three rules deal with state isolation through `run`. We write $[\text{hp } \varphi.]$ to denote an optional heap binding (so we really define six rules for state isolation). The first two rules (RUNL) and (RUNC) push a run operation down into a lambda-expression or partial catch expression. The final rule (RUNH) captures the essence of state isolation and reduces to a new value (or exception) discarding the heap φ . The side condition $\text{frv}(w) \not\cap \text{dom}(\varphi)$ is necessary to ensure well-formedness where a reference should not ‘escape’ its binding.

Using the reduction rules, we can now define an evaluation function. Using the evaluation context E defined in Figure 6, we define $E[e] \mapsto E[e']$ iff $e \rightarrow e'$. The evaluation context ensures strict semantics where only the leftmost- outermost reduction is applicable in an expression. We define the relation \mapsto as the reflexive and transitive closure of \mapsto . We can show that \mapsto is a function even though we need a simple diamond theorem since the order in which (LIFT) and (MERGE) reductions happen is not fixed [42].

The final results, or answers a , that expressions evaluate to, are either values v , exceptions `throw v`, heap bound values `hp φ . v` or heap bound exceptions `hp φ . throw v` (as defined in Figure 4).

Our modeling of the heap is slightly unconventional. Usually, a heap is defined either as a parameter to the semantics, or always scoped on the outside. This has the advantage that one doesn’t need operations like (MERGE) and (LIFT). However, for us it is important that the run operation can completely discard the heap which is hard to do in the conventional approach. In particular, in Section 6 we want to state the *purity* Theorem 3 that says that if $\Gamma \vdash e : \tau \mid \varepsilon$ where $st\langle h \rangle \notin \varepsilon$ then we never have $e \mapsto \text{hp } \varphi. v$. Such theorem is difficult to state if we modeled the heap more conventionally.

5. Semantic soundness

We now show that well-typed programs cannot go ‘wrong’. Our proof closely follows the subject reduction proofs of Wright and Felleisen [42]. Our main theorem is:

Theorem 1. (Semantic soundness)

If $\cdot \vdash e : \tau \mid \varepsilon$ then either $e \uparrow$ or $e \mapsto a$ where $\cdot \vdash a : \tau \mid \varepsilon$.

where we use the notation $e \uparrow$ for a never ending reduction. The proof of this theorem consists of showing two main lemmas:

- Show that reduction in the operational semantics preserves well-typing (called subject reduction).
- Show that *faulty* expressions are not typable.

If programs are closed and well-typed, we know from subject reduction that we can only reduce to well-typed terms, which can be either faulty, an answer, or an expression containing a further redex. Since faulty expressions are not typable it must be that evaluation either produces a well-typed answer or diverges. Often, proofs of soundness are carried out using *progress* instead of *faulty* expressions but it turns out that for proving the soundness of state isolation, our current approach works better.

5.1. Subject reduction

The subject reduction theorem states that a well-typed term remains well-typed under reduction:

Lemma 1. (Subject reduction)

If $\Gamma \vdash e_1 : \tau \mid \varepsilon$ and $e_1 \rightarrow e_2$ then $\Gamma \vdash e_2 : \tau \mid \varepsilon$.

To show that subject reduction holds, we need to establish various lemmas. Two particularly important lemmas are the substitution and extension lemmas:

Lemma 2. (Substitution)

If $\Gamma, x : \forall \bar{\alpha}. \tau \vdash e : \tau' \mid \varepsilon$ where $x \notin \text{dom}(\Gamma)$, $\Gamma \vdash v : \tau \mid \langle \rangle$, and $\bar{\alpha} \not\cap \text{ftv}(\Gamma)$, then $\Gamma \vdash [x \mapsto v]e : \tau' \mid \varepsilon$.

Lemma 3. (Extension)

If $\Gamma \vdash e : \tau \mid \varepsilon$ and for all $x \in \text{fv}(e)$ we have $\Gamma(x) = \Gamma'(x)$, then $\Gamma' \vdash e : \tau \mid \varepsilon$.

The proofs of these lemmas from [42] carry over directly to our system. However, to show subject reduction, we require an extra lemma to reason about state effects.

Lemma 4. (Stateful effects)

If $\Gamma \vdash e : \tau \mid \langle st\langle h \rangle \mid \varepsilon \rangle$ and $\Gamma \vdash R[e] : \tau' \mid \varepsilon'$ then $st\langle h \rangle \in \varepsilon'$.

The above lemma essentially states that a stateful effect cannot be discarded in an R context. Later we will generalize this lemma to arbitrary contexts and effects but for subject reduction this lemma is strong enough.

Proof. (Lemma 4) We proceed by induction on the structure of R :

Case $R = []$: By definition $st\langle h \rangle \in \langle st\langle h \rangle \mid \varepsilon \rangle$.

Case $R = R' e_2$: We have $\Gamma \vdash (R'[e]) e_2 : \tau' \mid \varepsilon'$ and by (APP) we have $\Gamma \vdash R'[e] : \tau_2 \rightarrow \varepsilon' \tau' \mid \varepsilon'$. By induction, $st\langle h \rangle \in \varepsilon'$.

Case $R = \nu R'$: Similar to previous case.

Case $R = \text{let } x = R' \text{ in } e_2$: By (LET) we have $\Gamma \vdash R'[e] : \tau_1 \mid \langle \rangle$ but that contradicts our assumption.

Case $R = \text{catch } R' e_2$: By (CATCH) we have $\Gamma \vdash \text{catch } R'[e] e_2 : \tau' \mid \varepsilon'$ where $\Gamma \vdash R'[e] : \tau' \mid \langle \text{exn} \mid \varepsilon' \rangle$. By induction $st\langle h \rangle \in \langle \text{exn} \mid \varepsilon' \rangle$ which implies that $st\langle h \rangle \in \varepsilon'$.

Before proving subject reduction we need one more type rule on our internal language. The (EXTEND) rule states that we can always assume a stateful effect for an expression:

$$\text{(ST-EXTEND)} \frac{\Gamma \vdash e : \tau \mid \varepsilon}{\Gamma \vdash e : \tau \mid \langle st\langle h \rangle \mid \varepsilon \rangle}$$

Now we are ready to prove the subject reduction theorem:

Proof. (Lemma 1) We prove this by induction on the reduction rules of \longrightarrow . We will not repeat all cases here and refer to [42], but instead concentrate on the interesting cases, especially with regard to state isolation.

Case $\text{let } x = \nu \text{ in } e \longrightarrow [x \mapsto \nu]e$: From (LET) we have $\Gamma \vdash \nu : \tau' \mid \langle \rangle$ and $\Gamma, x : \text{gen}(\Gamma, \tau') \vdash e : \tau \mid \varepsilon$. By definition, $\text{gen}(\Gamma, \tau') = \forall \bar{\alpha}. \tau'$ where $\bar{\alpha} \notin \text{ftv}(\Gamma)$ and by Lemma 2 we have $\Gamma \vdash [x \mapsto \nu]e : \tau \mid \varepsilon$.

Case $R[\text{hp } \varphi. e] \longrightarrow \text{hp } \varphi. R[e]$: This case is proven by induction over the structure of R :

case $R = []$: Does not apply due to the side condition \longrightarrow .

case $R = R' e'$: Then $\Gamma \vdash R'[\text{hp } \varphi. e] e' : \tau \mid \varepsilon$ and by (APP) we have $\Gamma \vdash R'[\text{hp } \varphi. e] : \tau_2 \rightarrow \varepsilon \tau \mid \varepsilon$ (1) and $\Gamma \vdash e' : \tau_2 \mid \varepsilon$ (2). By the induction hypothesis and (1), we have $\Gamma \vdash \text{hp } \varphi. R'[e] : \tau_2 \rightarrow \varepsilon \tau \mid \varepsilon$. Then by (HEAP) we know $\Gamma, \bar{\varphi}_h \vdash v_j : \tau_j \mid \langle \rangle$ (3) and $\Gamma, \bar{\varphi}_h \vdash R'[e] : \tau_2 \rightarrow \varepsilon \tau \mid \varepsilon$ (4) where $\varphi = \langle r_1 \mapsto v_1, \dots, r_n \mapsto v_n \rangle$. Since $r_1, \dots, r_n \notin \text{fv}(e')$ we can use (2) and 3 to conclude $\Gamma, \bar{\varphi} \vdash e' : \tau_2 \mid \varepsilon$ (5). Using (APP) with (4) and (5), we have $\Gamma, \bar{\varphi} \vdash R'[e] e' : \tau \mid \varepsilon$ where we can use (HEAP) with (3) to conclude $\Gamma \vdash \text{hp } \varphi. R'[e] e' : \tau \mid \varepsilon$.

case $R = \nu R'$: Similar to the previous case.

case $R = \text{let } x = R' \text{ in } e'$: If this is well-typed, then by rule (LET) we must have $\Gamma \vdash R'[\text{hp } \varphi. e] : \tau' \mid \langle \rangle$. However, due to Lemma 4 and (HEAP), we have $st\langle h \rangle \in \langle \rangle$ which is a contradiction. Note that this case is essential, as it prevents generalization of stateful references. For ML, this is also the tricky proof case that only works if one defines special ‘imperative type variables’ [42] or the value restriction. In our case the effect system ensures safety.

Case $\text{run}([\text{hp } \varphi.] \lambda x. e) \longrightarrow \lambda x. \text{run}([\text{hp } \varphi.] e)$: By rule (RUN) and (HEAP) we have that $\Gamma \vdash \lambda x. e : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$ where $h \notin \text{ftv}(\Gamma, \tau, \varepsilon)$ **(1)**. Applying (LAM) gives $\Gamma, x : \tau_1 \vdash e : \tau_2 \mid \varepsilon_2$ with $\tau = \tau_1 \rightarrow \varepsilon_2 \tau_2$. Using (ST-EXTEND) we can also derive $\Gamma, x : \tau_1 \vdash e : \tau_2 \mid \langle st \langle h \rangle \mid \varepsilon_2 \rangle$. Due to (1) and $h \notin \tau_1$, we can apply (RUN) and (HEAP) again to infer $\Gamma, x : \tau_1 \vdash \text{run}([\text{hp } \varphi.] e) : \tau_2 \mid \varepsilon_2$ and finally (LAM) again to conclude $\Gamma \vdash \lambda x. (\text{run}([\text{hp } \varphi.] e)) : \tau \mid \varepsilon$.

Case $\text{run}([\text{hp } \varphi.] \text{catch } e) \longrightarrow \text{catch}(\text{run}([\text{hp } \varphi.] e))$: Similar to the previous case.

Case $\text{run}([\text{hp } \varphi.] w) \longrightarrow w$ with $\text{frv}(w) \not\cap \text{dom}(\varphi)$ **(1)**: By rule (RUN) and (HEAP) we have that $\Gamma, \bar{\varphi}_h \vdash w : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$ where $h \notin \text{ftv}(\Gamma, \tau, \varepsilon)$ **(2)**. By (1) it must also be that $\Gamma \vdash w : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$ **(3)** (this follows directly if there was no heap binding $\text{hp } \varphi.$). We proceed over the structure of w :

case $w = \text{throw } v$: Then by (3) we have $\Gamma \vdash \text{throw } v : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$, but also $\Gamma \vdash \text{throw } v : \tau \mid \varepsilon$ since we can choose the result effect freely in (THROW).

case $w = r$: By (VAR) and (3), we have $\Gamma \vdash r : \text{ref} \langle h', \tau' \rangle \mid \langle st \langle h \rangle \mid \varepsilon \rangle$, where $h \neq h'$ satisfying (2). Since the result effect is free in (VAR), we can also derive $\Gamma \vdash r : \text{ref} \langle h', \tau' \rangle \mid \varepsilon$

case $w = (r :=)$: As the previous case.

case $w = x$: By (VAR) and (3), we have $\Gamma \vdash x : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$ but in (VAR) the result effect is free, so we can also derive $\Gamma \vdash x : \tau \mid \varepsilon$.

case other: Similarly.

5.2. Faulty expressions

The main purpose of type checking is of course to guarantee that certain bad expressions cannot occur. Apart from the usual errors, like adding a number to a string, we particularly would like to avoid state errors. There are two aspects to this. One of them is notorious where polymorphic types in combination with state can be unsound (which is not the case in our system because of Lemma 1). But in addition, we would like to show that in our system it is not possible to read or write to locations outside the local heap (encapsulated by run), nor is it possible to let local references escape. To make this precise, the *faulty* expressions are defined as:

- Undefined: $c v$ where $\delta(c, v)$ is not defined.
- Escaping read: $\text{run}(\text{hp } \varphi. R[!r])$ where $r \notin \text{dom}(\varphi)$.
- Escaping write: $\text{run}(\text{hp } \varphi. R[r := v])$ where $r \notin \text{dom}(\varphi)$.
- Escaping reference: $\text{run}(\text{hp } \varphi. w)$ where $\text{frv}(w) \cap \text{dom}(\varphi) \neq \emptyset$.
- Not a function: $v e$ where v is not a constant or lambda expression.
- Not a reference: $!v$ or $v := e$ where v is not a reference.
- Not an exception: $\text{throw } v$ where v is not the unit value.

Lemma 5. (*Faulty expressions are untypable*)

If an expression e is faulty, it cannot be typed, i.e. there exists no Γ , τ , and ε such that $\Gamma \vdash e : \tau \mid \varepsilon$.

In the following proof, especially the case for escaping state is interesting. To our knowledge, this is the first proof of the safety of run in a strict setting (and in combination with other effects like exceptions).

Proof. (Lemma 5) Each faulty expression is handled separately. We show here the interesting cases for escaping reads, writes, and references:

Case $\text{run}(\text{hp } \varphi. R[!r])$ with $r \notin \text{dom}(\varphi)$ **(1)**: To be typed in a context Γ we apply (RUN) and (HEAP) and need to show $\Gamma, \bar{\varphi}_h \vdash R[!r] : \tau \mid \langle st \langle h \rangle \mid \varepsilon \rangle$ **(2)**, where $h \notin \text{ftv}(\Gamma, \tau, \varepsilon)$ **(3)**. For $R[!r]$ to be well-typed, we also need $\Gamma, \bar{\varphi}_h \vdash !r : \tau' \mid \langle st \langle h' \rangle \mid \varepsilon' \rangle$ **(4)** where $\Gamma, \bar{\varphi}_h \vdash r : \text{ref} \langle h', \tau' \rangle \mid \langle st \langle h' \rangle \mid \varepsilon' \rangle$ **(5)**. From

Lemma 4, (4), and (2), it must be that $h' = h$ (6). But since $r \notin \text{dom}(\varphi)$ (1), it follows by (5) and (6), that $\Gamma \vdash r : \text{ref}\langle h, \tau' \rangle \mid \langle \text{st}\langle h \rangle \mid \varepsilon' \rangle$. But that means $h \in \text{ftv}(\Gamma)$ contradicting (3).

Case $\text{run}(\text{hp } \varphi. R[(r :=)])$ with $r \notin \text{dom}(\varphi)$: Similar to the previous case.

Case $\text{run}(\text{hp } \varphi. w)$ where $\text{frv}(w) \cap \text{dom}(\varphi) \neq \emptyset$ (1) To be typed in a context Γ we need to show by (HEAP) and (RUN) that $\Gamma, \bar{\varphi}_h \vdash w : \tau \mid \langle \text{st}\langle h \rangle \mid \varepsilon \rangle$ where $h \notin \text{ftv}(\Gamma, \tau, \varepsilon)$ (2). If $w = \text{throw } v$ then by (THROW) the type of v is $()$ and thus v is the unit constant. But $\text{frv}() = \emptyset$ contradicting our assumption. Otherwise, $w = b$ and cannot contain an arbitrary e . Since $\text{frv}(w) \neq \emptyset$ (1), it must be that w is either one of r or $(r :=)$ with $r \in \text{dom}(\varphi)$. To be well-typed, $\Gamma, \bar{\varphi}_h \vdash r : \text{ref}\langle h, \tau' \rangle \mid \varepsilon'$ must hold. However, the possible types for r and $(r :=)$ are $\text{ref}\langle h, \tau' \rangle$ and $\tau' \rightarrow \text{st}\langle h \rangle ()$ and in both cases $h \in \text{ftv}(\tau)$ which contradicts (2).

6. Effectful semantics

Up till now, we have used the effect types to good effect and showed that our system is semantically sound, even though state and polymorphic types are notoriously tricky to combine. Moreover, we showed that local state isolation through run is sound and statically prevents references from escaping.

But the true power of the effect system is really to enable more reasoning about the behavior of a program at a higher level. In particular, the absence of certain effects determines the absence of certain answers. For example, if the exception effect is not inferred, then evaluating the program will never produce an answer of the form $\text{throw } v$ or $\text{hp } \varphi. \text{throw } v!$ It would not be entirely correct to say that such program never throws an exception: indeed, a local catch block can handle such exceptions. We can state the exception property formally as:

Theorem 2. (Exceptions)

If $\Gamma \vdash e : \tau \mid \varepsilon$ where $\text{exn} \notin \varepsilon$ then either $e \uparrow$, $e \mapsto v$ or $e \mapsto \text{hp } \varphi. v$.

Proof. (Theorem 2) By contradiction over the result terms:

Case $e \mapsto \text{throw } v$: By subject reduction (Lemma 1), it must be $\Gamma \vdash \text{throw } v : \tau \mid \varepsilon$. Using the type rule for throw with (APP), it must be the case that $\varepsilon \equiv \langle \text{exn} \mid \varepsilon' \rangle$ contradicting our assumption.

Case $e \mapsto \text{hp } \varphi. \text{throw } v$: Similar to the previous case.

Similarly to the exception case, we can state such theorem over heap effects too. In particular, if the $\text{st}\langle h \rangle$ effect is absent, then evaluation will not produce an answer that contains a heap, i.e. $\text{hp } \varphi. w$. Again, it would not be right to say that the program itself never performs any stateful operations the state can be encapsulated inside a run construct and its stateful behavior is not observable from outside. Formally, we can state this as:

Theorem 3. (State)

If $\Gamma \vdash e : \tau \mid \varepsilon$ where $\text{st}\langle h \rangle \notin \varepsilon$ then either $e \uparrow$, $e \mapsto v$ or $e \mapsto \text{throw } v$.

Proof. (Theorem 3) We prove by contradiction over the result terms:

Case $e \mapsto \text{hp } \varphi. v$: By subject reduction (Lemma 1), it must be $\Gamma \vdash \text{hp } \varphi. v : \tau \mid \varepsilon$. Using (HEAP), it must be the case that $\varepsilon \equiv \langle \text{st}\langle h \rangle \mid \varepsilon' \rangle$ contradicting our assumption.

Case $e \mapsto \text{hp } \varphi. \text{throw } v$: Similar to the previous case.

Finally, our most powerful theorem is about the divergence effect; in particular, if the divergent effect is absent, then evaluation is guaranteed to terminate!

Theorem 4. (Divergence)

If $\Gamma \vdash e : \tau \mid \varepsilon$ where $\text{div} \notin \varepsilon$ then $e \mapsto a$.

The proof of the divergence theorem (Theorem 4) is more complicated as we cannot use subject reduction to show this by contradiction. Instead, we need to do this proof using induction over logical relations [9].

In our case, we say that if $\cdot \vdash e : \tau \mid \varepsilon$, then e is in the set $\mathcal{R}(\tau \mid \varepsilon)$, “the reducible terms of type τ with effect ε ”, if $\text{div} \notin \varepsilon$ and (1) when τ is a non-arrow type, if e halts, and (2) when $\tau = \tau_1 \rightarrow \varepsilon_2 \tau_2$, if e halts and if for all $e_1 \in \mathcal{R}(\tau_1 \mid \varepsilon)$, we have that $e e_1 \in \mathcal{R}(\tau_2 \mid \varepsilon_2)$.

The proof of Theorem 4 is a standard result [9] and is a corollary from the following two main lemmas:

Lemma 6. (*\mathcal{R} is preserved by reduction*) Iff $\cdot \vdash e : \tau \mid \varepsilon$, $e \in \mathcal{R}(\tau \mid \varepsilon)$, and $e \mapsto e'$, then also $e' \in \mathcal{R}(\tau \mid \varepsilon)$.

Lemma 7. (*A well-typed term is in \mathcal{R}*) If $\cdot \vdash e : \tau \mid \varepsilon$ and $\text{div} \notin \varepsilon$, then $e \in \mathcal{R}(\varepsilon \mid \tau)$.

Proof. (Lemma 6) This is shown by induction over the type τ . For atomic types, this holds by definition. For arrow types, $\tau_1 \rightarrow \varepsilon_2 \tau_2$ we must show for a given $e_1 \in \mathcal{R}(\tau_1 \mid \varepsilon)$ that if $e e_1 \in \mathcal{R}(\tau_2 \mid \varepsilon_2)$, then also $e' e_1 \in \mathcal{R}(\tau_2 \mid \varepsilon_2)$ (and the other direction). By (APP), it must be $\varepsilon_2 = \varepsilon$ (1). We can now proceed over the structure of reductions on $e e_1$:

Case (!) e_1 : In this case, since (READ) has a *div* effect, we have by (1) that $\text{div} \in \varepsilon$ contradicting our assumption. Note that if we would have cheated and not include *div* in the type, we would have gotten a reduction to some v which we could not show to be strongly normalizing, and thus if it is an element of $\mathcal{R}(\tau_2 \mid \varepsilon_2)$.

Case $\text{fix } \varepsilon_1$: As the previous case.

Case $(\lambda x. e_2) e_1$: In this case, we can reduce to $e' e_1$, and by the induction hypothesis $e' e_1 \in \mathcal{R}(\tau \mid \varepsilon_2)$ since τ_2 is smaller.

Proof. (Lemma 7) This is proven over the structure of the type derivation. However, as usual, we need to strengthen our induction hypothesis to include the environment. We extend \mathcal{R} over environments to be a set of substitutions:

$$\mathcal{R}(\Gamma) = \{\theta \mid \text{dom}(\Gamma) = \text{dom}(\theta) \wedge \forall (x : \tau \in \Gamma). \theta x \in \mathcal{R}(\tau \mid \langle \rangle)\}$$

where we assume a monomorphic environment for simplicity but we can extend this easily to a (first-order) polymorphic setting. Our strengthened lemma we use for our proof is:

$$\text{if } \Gamma \vdash e : \tau \mid \varepsilon \wedge \theta \in \mathcal{R}(\Gamma) \wedge \text{div} \notin \varepsilon \text{ then } \theta e \in \mathcal{R}(\tau \mid \varepsilon)$$

The induction is standard, and we show only some sample cases:

Case (FIX): Since the result effect is free, we can choose any ε such that $\text{div} \notin \varepsilon$. Indeed, just an occurrence of *fix* is ok – only an application may diverge.

Case (APP): By the induction hypothesis and (APP) we have $\theta e_1 \in \mathcal{R}(\tau_2 \rightarrow \varepsilon \tau \mid \varepsilon)$ and $\theta e_2 \in \mathcal{R}(\tau_2 \mid \varepsilon)$. By definition of $\mathcal{R}(\tau_2 \rightarrow \varepsilon \tau \mid \varepsilon)$, $\theta e_1 \theta e_2 \in \mathcal{R}(\tau_2 \mid \varepsilon)$ and therefore $\theta(e_1 e_2) \in \mathcal{R}(\tau_2 \mid \varepsilon)$. Note that the induction hypothesis ensures that $\text{div} \notin \varepsilon$ and therefore we cannot apply a potentially divergent function (like *fix* or (!)).

Case other: Standard, except that for effect elimination rules, we need to show that *div* is not eliminated.

Contrast the results of this section to many other effect systems [23,31] that just use effect types as syntactic ‘labels’. In our case the theorems of this section are truly significant as they show there is a deep correspondence between the effect types and the dynamic semantics which eventually enables us to reason about our programs much more effectively.

7. Related work

A main contribution of this paper is showing that our notion of mutable state is sound, in particular the combination of mutable state and polymorphic let- bindings is tricky as shown by Tofte [37] for the ML language. Later, variants of the ML value restriction are studied by Leroy [20].

Safe state encapsulation using a lazy state monad was first proven formally by Launchbury and Sabry [16]. Their formalization is quite different though from ours and applies to a lazy store in a monadic setting. In particular, in their formalization there is no separate heap binding, but heaps are always bound at the outer *run*. We tried this, but it proved difficult in our setting; for example, it is hard to state the stateful lemma since answers would never contain an explicit heap. Very similar to our state encapsulation is region inference [38]. Our *run* operation essentially delimits a heap region. Regions are values though, and we can for example not access references in several regions at once.

Independently of our work, Lindley and Cheney [21] also used row polymorphism for effect types. Their approach is based on presence/absence flags [29] to give effect types to database operations in the context of the Links web programming language. The main effects of the database operations are *wild*, *tame*, and *hear*, for arbitrary effects including divergence, pure queries, and asynchronous messages respectively. They report on practical experience exposing effect types to the programmer and discuss various syntax forms to easily denote effect types.

The problems with arbitrary effects have been widely recognized, and there is a large body of work studying how to delimit the scope of effects. There have been many effect typing disciplines proposed. Early work is by Gifford and Lucassen [8,22] which was later extended by Talpin [35] and others [26,34]. These systems are closely related since they describe polymorphic effect systems and use type constraints to give principal types. The system described by Nielson *et al.* [26] also requires the effects to form a complete lattice with meets and joins.

Java contains a simple effect system where each method is labeled with the exceptions it might raise [10]. A system for finding uncaught exceptions was developed for ML by Pessaux *et al.* [27]. A more powerful system for tracking effects was developed by Benton [2] who also studies the semantics of such effect systems [3]. Recent work on effects in Scala [31] shows how even a restricted form of polymorphic effect types can track effects for many programs in practice.

Tolmach [39] describes an effect analysis for ML in terms of effect monads, namely *Total*, *Partial*, *Divergent* and *ST*. This system is not polymorphic though and meant more for internal compiler analysis. In the context proof systems there has been work to show absence of observable side effects for object-oriented programming languages, for example by Naumann [25].

Marino *et al.* recently produced a generic type-and-effect system [23]. This system uses privilege checking to describe analytical effect systems. For example, an effect system could use try-catch statements to grant the *canThrow* privilege inside try blocks. *throw* statements are then only permitted when this privilege is present. Their system is very general and can express many properties but has no semantics on its own. For example, it would be sound for the effect system to have “+” grant the *canThrow* privilege to its arguments, and one has to do an additional proof to show that the effects in these systems actually correspond to an intended meaning.

Wadler and Thiemann showed the close relation between effect systems and monads [40] and showed how any effect system can be translated to a monadic version. For our particular system though a monadic translation is quite involved due to polymorphic effects; essentially we need dependently typed operations and we leave a proper monadic semantics for future work.

Recently, *effect handlers* are proposed to program with effects [4,14,15]. In this work, computational effects are modeled as operations of an algebraic theory. Even though algebraic effects are subsumed

by monads, they can be combined more easily and the specification of handlers offers new ways of programming. This work is quite different than what we described in this paper since we only considered effects that are intrinsic to the language while effect handlers deal specifically with ‘user-defined’ effects. However, we are currently working on integrating user-defined effects in Koka using effect handlers, and investigating how this can work well with the current effect type system.

References

- [1] Andreas Abel (1998): *Foetus – A Termination Checker for Simple Functional Programs*. Unpublished note.
- [2] Nick Benton & Peter Buchlovsky (2007): *Semantics of an effect analysis for exceptions*. In: *TLDI '07: Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*, pp. 15–26.
- [3] Nick Benton, Andrew Kennedy, Lennart Beringer & Martin Hofmann (2007): *Relational semantics for effect-based program transformations with dynamic allocation*. In: *PPDP '07: Proc. of the 9th ACM SIGPLAN int. conf. on Principles and Practice of Declarative Prog.*, pp. 87–96.
- [4] Edwin Brady (2013): *Programming and Reasoning with Algebraic Effects and Dependent Types*. In: *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming, ICFP '13*, ACM, New York, NY, USA, pp. 133–144, doi:[10.1145/2500365.2500581](https://doi.org/10.1145/2500365.2500581).
- [5] Luis Damas & Robin Milner (1982): *Principal type-schemes for functional programs*. In: *9th ACM symp. on Principles of Programming Languages (POPL '82)*, pp. 207–212.
- [6] Jean-Christophe Filliâtre (2008): *A Functional Implementation of the Garsia–wachs Algorithm: (Functional Pearl)*. In: *Proceedings of the 2008 ACM SIGPLAN Workshop on ML, ML '08*, ACM, New York, NY, USA, pp. 91–96, doi:[10.1145/1411304.1411317](https://doi.org/10.1145/1411304.1411317).
- [7] Ben R. Gaster & Mark P. Jones (1996): *A Polymorphic Type System for Extensible Records and Variants*. Technical Report NOTTCS-TR-96-3, University of Nottingham.
- [8] David K. Gifford & John M. Lucassen (1986): *Integrating functional and imperative programming*. In: *LFP '86: Proceedings of the 1986 ACM conference on LISP and functional programming*, pp. 28–38.
- [9] Jean-Yves Girard, Paul Taylor & Yves Lafont (1989): *Proofs and types*. Cambridge University Press.
- [10] James Gosling, Bill Joy & Guy Steele (1996): *The Java Language Specification*. Addison-Wesley.
- [11] J.R. Hindley (1969): *The principal type scheme of an object in combinatory logic*. *Trans. of the American Mathematical Society* 146, pp. 29–60.
- [12] Mark P. Jones (1992): *A theory of qualified types*. In: *4th. European Symposium on Programming (ESOP'92), Lecture Notes in Computer Science 582*, Springer-Verlag, pp. 287–306.
- [13] Mark P. Jones (1995): *A system of constructor classes: overloading and implicit higher-order polymorphism*. *Journal of Functional Programming* 5(1), pp. 1–35.
- [14] Ohad Kammar, Sam Lindley & Nicolas Oury (2013): *Handlers in Action*. In: *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming, ICFP '13*, ACM, New York, NY, USA, pp. 145–158, doi:[10.1145/2500365.2500590](https://doi.org/10.1145/2500365.2500590).
- [15] Ohad Kammar & Gordon D. Plotkin (2012): *Algebraic Foundations for Effect-dependent Optimisations*. In: *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '12*, ACM, New York, NY, USA, pp. 349–360, doi:[10.1145/2103656.2103698](https://doi.org/10.1145/2103656.2103698).
- [16] John Launchbury & Amr Sabry (1997): *Monadic State: Axiomatization and Type Safety*. In: *ICFP'97*, pp. 227–238.
- [17] Daan Leijen (2005): *Extensible records with scoped labels*. In: *In: Proceedings of the 2005 Symposium on Trends in Functional Programming*, pp. 297–312.

- [18] Daan Leijen (2012): *Try Koka online*. <http://rise4fun.com/koka/tutorial> and <http://koka.codeplex.com>.
- [19] Daan Leijen (2013): *Koka: Programming with Row-Polymorphic Effect Types*. Technical Report MSR-TR-2013-79, Microsoft Research.
- [20] Xavier Leroy (1993): *Polymorphism by name for references and continuations*. In: *POPL '93: Proc. of the 20th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pp. 220–231.
- [21] Sam Lindley & James Cheney (2012): *Row-based effect types for database integration*. In: *TLDI'12*, pp. 91–102.
- [22] J. M. Lucassen & D. K. Gifford (1988): *Polymorphic effect systems*. In: *POPL '88*, pp. 47–57.
- [23] Daniel Marino & Todd Millstein (2009): *A generic type-and-effect system*. In: *TLDI '09: Proceedings of the 4th international workshop on Types in language design and implementation*, pp. 39–50.
- [24] Robin Milner (1978): *A theory of type polymorphism in programming*. *Journal of Computer and System Sciences* 17, pp. 248–375.
- [25] David A. Naumann (2007): *Observational purity and encapsulation*. *Theor. Comput. Sci.* 376(3), pp. 205–224.
- [26] Hanne Riis Nielson, Flemming Nielson & Torben Amtoft (1997): *Polymorphic Subtyping for Effect Analysis: The Static Semantics*. In: *Selected papers from the 5th LOMAPS Workshop on Analysis and Verification of Multiple-Agent Languages*, pp. 141–171.
- [27] François Pessaux & Xavier Leroy (1999): *Type-based analysis of uncaught exceptions*. In: *POPL '99*, pp. 276–290.
- [28] Simon L Peyton Jones & John Launchbury (1995): *State in Haskell*. *Lisp and Symbolic Comp.* 8(4), pp. 293–341.
- [29] Didier Rémy (1994): In Carl A. Gunter & John C. Mitchell, editors: *Theoretical Aspects of Object-oriented Programming*, chapter Type Inference for Records in Natural Extension of ML, MIT Press, Cambridge, MA, USA, pp. 67–95. Available at <http://dl.acm.org/citation.cfm?id=186677.186689>.
- [30] Didier Remy (1994): *Programming Objects with ML-ART, an Extension to ML with Abstract and Record Types*. In: *TACS '94: Proc. Int. Conf. on Theoretical Aspects of Computer Software*, pp. 321–346.
- [31] Lukas Rytz, Martin Odersky & Philipp Haller (2012): *Lightweight Polymorphic Effects*. In: *Proceedings of the 26th European Conference on Object-Oriented Programming, ECOOP'12*, Springer-Verlag, Berlin, Heidelberg, pp. 258–282, doi:10.1007/978-3-642-31057-7_13.
- [32] Martin Sulzmann (1997): *Designing record systems*. Technical Report YALEU/DCS/RR-1128, Yale University.
- [33] Martin Sulzmann (1998): *Type systems for records revisited*. Unpublished.
- [34] Jean-Pierre Talpin & Pierre Jouvelot (1994): *The type and effect discipline*. *Inf. Comput.* 111(2), pp. 245–296.
- [35] J.P. Talpin (1993): *Theoretical and practical aspects of type and effect inference*. Ph.D. thesis, Ecole des Mines de Paris and University Paris VI, Paris, France.
- [36] Ross Tate & Daan Leijen (2010): *Convenient Explicit Effects using Type Inference with Subeffects*. Technical Report MSR-TR-2010-80, Microsoft Research.
- [37] Mads Tofte (1990): *Type inference for polymorphic references*. *Inf. Comput.* 89(1), pp. 1–34.
- [38] Mads Tofte & Lars Birkedal (1998): *A region inference algorithm*. *ACM Trans. Program. Lang. Syst.* 20(4), pp. 724–767.
- [39] Andrew P. Tolmach (1998): *Optimizing ML Using a Hierarchy of Monadic Types*. In: *TIC '98*, pp. 97–115.
- [40] Philip Wadler & Peter Thiemann (2003): *The marriage of effects and monads*. *ACM Trans. Comput. Logic* 4(1), pp. 1–32.
- [41] Mitchell Wand (1987): *Complete Type Inference for Simple Objects*. In: *Proceedings of the 2nd. IEEE Symposium on Logic in Computer Science*, pp. 37–44. Corrigendum in LICS'88, page 132.

- [42] Andrew K. Wright & Matthias Felleisen (1994): *A syntactic approach to type soundness*. *Inf. Comput.* 115(1), pp. 38–94.

$$\begin{array}{c}
(\text{VAR})_s \quad \frac{\Gamma(x) = \forall \bar{\alpha}. \tau}{\Gamma \vdash_s x : [\bar{\alpha} \mapsto \bar{\tau}] \tau \mid \varepsilon} \\
\\
\Gamma \vdash_s e_1 : \tau_1 \mid \langle \rangle \quad \bar{\alpha} \notin \text{ftv}(\Gamma) \\
(\text{LET})_s \quad \frac{\Gamma, x : \forall \bar{\alpha}. \tau_1 \vdash_s e_2 : \tau_2 \mid \varepsilon}{\Gamma \vdash_s \text{let } x = e_1 \text{ in } e_2 : \tau_2 \mid \varepsilon}
\end{array}$$

Figure 7. Changed rules for the syntax directed system; Rule (INST) and (GEN) are removed, and all other rules are equivalent to the declarative system (Figure 3)

Appendix

A. Type inference

As a first step toward type inference, we first present in a syntax directed version of our declarative type rules in Figure 7. For this system, the syntax tree completely determines the derivation tree. Effectively, we removed the (INST) and (GEN) rules, and always apply instantiation in the (VAR) rule, and always generalize at let-bindings. This technique is entirely standard [11,12,24] and we can show that the syntax directed system is sound and complete with respect to the declarative rules:

Theorem 5. (*Soundness of the syntax directed rules*)

When $\Gamma \vdash_s e : \tau \mid \varepsilon$ then we also have $\Gamma \vdash e : \tau \mid \varepsilon$.

Theorem 6. (*Completeness of the syntax directed rules*)

When $\Gamma \vdash e : \sigma \mid \varepsilon$ then we also have $\Gamma \vdash_s e : \tau \mid \varepsilon$ where σ can be instantiated to τ .

Both proofs are by straightforward induction using standard techniques as described for example by Jones [12].

A.1. The type inference algorithm

Starting from the syntax directed rules, we can now give a the type inference algorithm for our effect system which is shown in Figure 8. Following Jones [12] we present the algorithm as natural inference rules of the form $\theta \Gamma \vdash e : \tau \mid \varepsilon$ where θ is a substitution, Γ the environment, and e , τ , and ε , the expression, its type, and its effect respectively. The rules can be read as an attribute grammar where θ , τ , and ε are synthesised, and Γ and e inherited. An advantage is that this highlights the correspondence between the syntax directed rules and the inference algorithm.

The algorithm uses unification written as $\tau_1 \sim \tau_2 : \theta$ which unifies τ_1 and τ_2 with a most general substitution θ such that $\theta \tau_1 = \theta \tau_2$. %The unification algorithm is standard and effects are unified using standard row unification allowing for duplicate label as described by Leijen [17]. The gen function generalizes a type with respect to an environment and is defined as:

$$\text{gen}(\Gamma, \tau) = \forall(\text{ftv}(\tau) - \text{ftv}(\Gamma)). \tau$$

We can prove that the inference algorithm is sound and complete with respect to the syntax directed rules (and by Theorem 5 and 6 also sound and complete to the declarative rules):

Theorem 7. (*Soundness*)

If $\theta \Gamma \vdash_i e : \tau \mid \varepsilon$ then there exists a θ' such that $\theta \Gamma \vdash_s e : \tau' \mid \varepsilon'$ where $\theta' \tau = \tau'$ and $\theta' \varepsilon = \varepsilon'$.

$$\begin{array}{c}
(\text{VAR})_i \quad \frac{\Gamma(x) = \forall \bar{\alpha}. \tau}{\emptyset \Gamma \vdash_i x : [\bar{\alpha} \mapsto \bar{\beta}] \tau \mid \mu} \\
\\
(\text{LAM})_i \quad \frac{\theta \Gamma, x : \alpha \vdash_i e : \tau_2 \mid \varepsilon_2}{\theta \Gamma \vdash_i \lambda x. e : \theta \alpha \rightarrow \varepsilon_2 \tau_2 \mid \mu} \\
\\
(\text{APP})_i \quad \frac{\theta_1 \Gamma \vdash_i e_1 : \tau_1 \mid \varepsilon_1 \quad \theta_2(\theta_1 \Gamma) \vdash_i e_2 : \tau_2 \mid \varepsilon_2 \quad \theta_2 \tau_1 \sim (\tau_2 \rightarrow \varepsilon_2 \alpha) : \theta_3 \quad \theta_3 \theta_2 \varepsilon_1 \sim \theta_3 \varepsilon_2 : \theta_4}{\theta_4 \theta_3 \theta_2 \theta_1 \Gamma \vdash_i e_1 e_2 : \theta_4 \theta_3 \alpha \mid \theta_4 \theta_3 \varepsilon_2} \\
\\
(\text{LET})_i \quad \frac{\theta_1 \Gamma \vdash_i e_1 : \tau_1 \mid \varepsilon_1 \quad \varepsilon_1 \sim \langle \rangle : \theta_2 \quad \sigma = \text{gen}(\theta_2 \theta_1 \Gamma, \theta_2 \tau_1) \quad \theta_3(\theta_2 \theta_1 \Gamma, x : \sigma) \vdash_i e_2 : \tau \mid \varepsilon}{\theta_3 \theta_2 \theta_1 \Gamma \vdash_i \text{let } x = e_1 \text{ in } e_2 : \tau \mid \varepsilon} \\
\\
(\text{RUN})_i \quad \frac{\theta_1 \Gamma \vdash_i e : \tau \mid \varepsilon \quad \varepsilon \sim \langle \text{st}(\xi) \mid \mu \rangle : \theta_2 \quad \theta_2 \xi \in \text{TypeVar} \quad \theta_2 \xi \notin \text{ftv}(\theta_2 \theta_1 \Gamma, \theta_2 \tau, \theta_2 \mu)}{\theta_2 \theta_1 \Gamma \vdash_i \text{run } e : \theta_2 \tau \mid \theta_2 \mu} \\
\\
(\text{CATCH})_i \quad \frac{\theta_1 \Gamma \vdash_i e_1 : \tau_1 \mid \varepsilon_1 \quad \theta_2(\theta_1 \Gamma) \vdash_i e_2 : \tau_2 \mid \varepsilon_2 \quad \theta_2 \varepsilon_1 \sim \langle \text{exn} \mid \varepsilon_2 \rangle : \theta_3 \quad \theta_3 \tau_2 \sim () \rightarrow \theta_3 \varepsilon_2 \quad \theta_3 \theta_2 \tau_1 : \theta_4}{\theta_4 \theta_3 \theta_2 \theta_1 \Gamma \vdash \text{catch } e_1 e_2 : \theta_4 \theta_3 \tau_2 \mid \theta_4 \theta_3 \varepsilon_2}
\end{array}$$

Figure 8. Type and effect inference algorithm. Any type variables α , μ , ξ , and $\bar{\alpha}$ are considered fresh.

Theorem 8. (*Completeness*)

If $\theta_1 \Gamma \vdash_s e : \tau_1 \mid \varepsilon_1$ then $\theta_2 \Gamma \vdash_i e : \tau_2 \mid \varepsilon_2$ and there exists a substitution θ such that $\theta_1 \approx \theta \theta_2$, $\tau_1 = \theta \tau_2$ and $\varepsilon_1 = \theta \varepsilon_2$.

Since the inference algorithm is basically just algorithm W [5] together with extra unifications for effect types, the proofs of soundness and completeness are entirely standard. The main extended lemma is for the soundness, completeness, and termination of the unification algorithm which now also unifies effect types.

The unification algorithm is shown in Figure 9. The algorithm is an almost literal adaption of the unification algorithm for records with scoped labels as described by Leijen [17], and the proofs of soundness, completeness, and termination carry over directly.

The first four rules are the standard Robinson unification rules with a slight modification to return only kind-preserving substitutions [7, 13]. The rule (UNI-EFF) unifies effect rows. The operation $\text{tl}(\varepsilon)$ is defined as:

$$\begin{aligned}
\text{tl}(\langle l_1, \dots, l_n \mid \mu \rangle) &= \mu \\
\text{tl}(\langle l_1, \dots, l_n \rangle) &= \langle \rangle
\end{aligned}$$

As described in detail in [17], the check $\text{tl}(\varepsilon_1) \notin \text{dom}(\theta)_1$ is subtle but necessary to guarantee termination of row unification. The final three rules unify an effect with a specific head. In particular, $\varepsilon \simeq l \mid \varepsilon' : \theta$

$$\begin{array}{l}
\text{(UNI-VAR)} \quad \alpha \sim \alpha : [] \\
\text{(UNI-VARL)} \quad \frac{\alpha \notin \text{ftv}(\tau)}{\alpha^k \sim \tau^k : [\alpha \mapsto \tau]} \\
\text{(UNI-VARR)} \quad \frac{\alpha \notin \text{ftv}(\tau)}{\tau^k \sim \alpha^k : [\alpha \mapsto \tau]} \\
\forall i \in 1..n. \quad \theta_{i-1} \dots \theta_1 \tau_i \sim \theta_{i-1} \dots \theta_1 t_i : \theta_i \\
\text{(UNI-CON)} \quad \frac{\kappa = (\kappa_1, \dots, \kappa_n) \rightarrow \kappa'}{c^\kappa \langle \tau_1^{\kappa_1}, \dots, \tau_n^{\kappa_n} \rangle \sim c^\kappa \langle t_1^{\kappa_1}, \dots, t_n^{\kappa_n} \rangle : \theta_n \dots \theta_1} \\
\text{(UNI-EFF)} \quad \frac{\varepsilon_2 \simeq l \mid \varepsilon_3 : \theta_1 \quad \text{tl}(\varepsilon_1) \notin \text{dom}(\theta_1) \quad \theta_1 \varepsilon_1 \sim \theta_1 \varepsilon_3 : \theta_2}{\langle l \mid \varepsilon_1 \rangle \sim \varepsilon_2 : \theta_2 \theta_1} \\
\text{(EFF-HEAD)} \quad \frac{l \equiv l' \quad l \sim l' : \theta}{\langle l' \mid \varepsilon \rangle \simeq l \mid \varepsilon : \theta} \\
\text{(EFF-SWAP)} \quad \frac{l \neq l' \quad \varepsilon \simeq l \mid \varepsilon' : \theta}{\langle l' \mid \varepsilon \rangle \simeq l \mid \langle l \mid \varepsilon' \rangle : \theta} \\
\text{(EFF-TAIL)} \quad \frac{\text{fresh } \mu'}{\mu \simeq l \mid \mu' : [\mu \mapsto \langle l \mid \mu' \rangle]}
\end{array}$$

Figure 9. Unification: $\tau \sim \tau' : \theta$ unifies two types and returns a substitution θ . It uses effect unification $\varepsilon \simeq l \mid \varepsilon' : \theta$ which takes an effect ε and effect primitive l as input, and returns effect tail ε' and a substitution θ .

states that for a given effect row ε , we match it with a given effect constant l , and return an effect tail ε' and substitution θ such that $\theta \varepsilon = \langle \theta l \mid \theta \varepsilon' \rangle$. Each rule basically corresponds to the equivalence rules on effect rows (Figure 2).