Catch me if you can Looking for type-safe, hierarchical, lightweight, polymorphic and efficient error management in OCaml

David Teller, Arnaud Spiwack, Till Varoquaux

▶ To cite this version:

HAL Id: inria-00432575

https://hal.inria.fr/inria-00432575

Submitted on 16 Nov 2009

Catch me if you can

Looking for type-safe, hierarchical, lightweight, polymorphic and efficient error management in OCaml

David Teller[1], Arnaud Spiwack[2], and Till Varoquaux[3]

[1] LIFO, Université d'Orléans David.Teller@univ-orleans.fr
[2] LIX, École Polytechnique Arnaud.Spiwack@lix.polytechnique.fr
[3] till@pps.jussieu.fr

Abstract. This is the year 2008 and ML-style exceptions are everywhere. Most modern languages, whether academic or industrial, feature some variant of this mechanism. Languages such as Java even feature static coverage-checking for such exceptions, something not available for ML languages, at least not without resorting to external tools.
In this document, we demonstrate a design principle and a tiny library for managing errors in a functional manner, with static coverage-checking, automatically-inferred, structurally typed and hierarchical exceptional cases, with a reasonable run-time penalty. Our work is based on OCaml and features monads, polymorphic variants, compile-time code rewriting and trace elements of black magic.

1   Introduction

Despite our best intentions and precautions, even correct programs may fail. The disk may be full, the password provided by the user may be wrong or the expression whose result the user wants plotted may end up in a division by zero. Indeed, management of dynamic errors and other exceptional circumstances inside programs is a problem nearly as old as programming. Error management techniques should be powerful enough to cover all possible situations and flexible enough to let the programmer deal with whichever cases are his responsibility while letting other modules handle other cases, should be sufficiently noninvasive so as to let the programmer concentrate on the main path of execution while providing guarantees that exceptional circumstances will not remain unmanaged, all without compromising performance or violating the paradigm.

Nowadays, most programming languages feature a mechanism based on (or similar to) the notion of *exceptions*, as pioneered by PL/I [10], usually with the semantics later introduced in ML [12]. A few languages, such as Haskell, define this mechanism as libraries [19], while most make it a language primitive, either because the language is not powerful enough, for the sake of performance, to add sensible debugging information, or as a manner of sharing a common mechanism for programmer errors and manageable errors.

As a support for our discussion on the management of errors and exceptional circumstances, let us introduce the following type for the representation of arithmetic expressions, written in OCaml:

```
type expr = Value of float
          | Div   of expr * expr
          | Add   of expr * expr
```

The implementation of an evaluator for this type is a trivial task:

```
let rec eval = function
  Value f      → f
| Div   (x, y) → (eval x) /. (eval y)
| Add   (x, y) → (eval x) +. (eval y)
(*val eval: expr → float*)
```

However, as such, the interpreter fails to take into account the possibility of division by zero. In order to manage this *exceptional circumstance* (or *error*), we promptly need to rewrite the code into something more complex:

Listing 1.1. Ad-hoc error management

```
type (α, β) result = Ok of α | Error of β

let rec eval = function
  Value f      → Ok f
| Div   (x, y) → (match eval x with
    Error e → Error e
  | Ok x'   → match eval y with
      Error e              → Error e
    | Ok y' when y' = 0. → Error "Divison by 0"
    | Ok y'                → Ok (x' /. y'))
| Add   (x, y) → (match eval x with
    Error e              → Error e
  | Ok x'                → match eval y with
      Error e              → Error e
    | Ok y'                → Ok (x' +. y'))
(*val eval: expr → (float, string) result*)
```

While this function succeeds in managing exceptional cases, the code is clumsy and possibly slow. An alternative is to use the built-in mechanism of exceptions – which we will refer to as "native exceptions"– as follows:

Listing 1.2. Error management with native exceptions

```
exception Error of string

let rec eval = function
  Value f      → f
| Div   (x, y) → let (x', y') = (eval x, eval y) in
    if y' = 0. then raise (Error "division by 0")
    else                 x' /. y'
| Add   (x, y) → eval x +. eval y
(*val eval: expr → float *)
```

This definition of `eval` is easier to write and read, closer to the mathematical definition of arithmetic operations and faster. While native exceptions appear

to be a great win over explicitly returning an error value, the implementation of this mechanism is commonly both less flexible and less safe than ad-hoc error management. In particular, in ML languages, the loss of flexibility appears as the impossibility of defining exceptions which would use any polymorphic type parameters[4]. As for the loss of safety, it is a consequence of `eval` containing no type information could let us determine that the function may fail and what kind of information may accompany the failure. Worse than that: the compiler itself does not have such information and cannot provide guarantees that every exceptional case will eventually be managed. Arguably, the possibility for a native exception to completely escape is comparable to the possibility for a pattern-matching to go wrong, which in turn is comparable to null-pointer exceptions in most modern industrial languages – while technically not a type error, this remains a source of unsafety which we will refer to as "incomplete coverage" in the rest of this document.

While it is possible to complete the type system to guarantee complete coverage, perhaps as a variation on *types and effects* [17], either as part of the mandatory compilation process, as in Java, or as external tools [14], this guarantee does nothing to improve the lack of flexibility. This is quite unfortunate, as numerous situations thus require the manual definition of many superfluous exceptions with identical semantics but different types, or lead to the overuse of magic constants to specify sub-exceptions, or require impure or unsafe hacks to implement simple features. SML provides a way to regain some polymorphism in exceptions defined *locally* using *generative exceptions.* To improve flexibility one needs to change deeply the type system [1]

Another possible approach which may be used to obtain both the readability of exceptions, guarantee of complete coverage and parametric polymorphism, is to implement error management as monads [19], a path followed by Haskell. However, this approach often results in either not-quite-readable and possibly ambiguous type combinations consisting in large hierarchies of algebraic combinators, in the necessity of writing custom error monads or monad transformers, which need to be manually rewritten as often as the list of exceptional cases changes, or in the use of dynamic types. In addition, these monads are perceived as having a large computational cost, due to constant thunking and dethunking of continuations and to the lack of compiler-optimized stack unrolling.

In this document, we attempt to obtain the best of both worlds: polymorphism, type-safety, coverage-check, with the added benefits of automatic inference of error cases and the definition of classes and subclasses of exceptional cases, all of this without the need to modify the programming language. As this is an OCaml-based work, we also take into account the impact of this program-

---

[4] For comparison, the Java type-checker rejects subtypes of `Exception` with parametric polymorphism, the C# *parser* rejects catch clauses with parametric polymorphism but allows general catch clauses followed by unsafe downcasts, while Scala accepts defining, throwing and catching exceptions with parametric polymorphism, but the semantics of the language ignores these type parameters both during compilation and during execution.

ming style in terms of both performance, syntactic sugar, possible compile-time optimizations. Despite the number of claims appearing in the previous sentences, our work is actually based on very simple concepts and does not have the ambition of introducing brand new programming methodologies, nor to revolutionize ML programming. Rather, our more limited objective is to present an interesting design principle and a tiny library for error management, in the spirit of Functional Pearls or of a tutorial, and based on ML-style exceptions, monads, phantom types, polymorphic variants and code rewriting. Some of our results are positive, some negative and, somewhere along the way, we revisit techniques results discarded by previous works on hierarchical exceptions [11] and demonstrate that, when redefined with better language support, they may be used to provide safer (and possibly faster) results.

In a first section we demonstrate briefly an error monad, well-known in the world of Haskell but perhaps less common in the ML family. We then proceed to complete this monad with the use of lightweight types to achieve automatic inference of error cases, before pushing farther these lightweight types to permit the representation of classes and subclasses of exceptional cases, while keeping guarantees of coverage. Once this is done, we study the performance of this monad and progressively move the code from the library to the compiler. Finally, we conclude by a discussion on usability, potential improvements, and comparison with other related works.

The source code for the library is available as a downloadable package [18]. Elements of this work will be included in the rehauled OCaml standard libraries, OCaml Batteries Included.

2   The error monad

As we discussed already, Listing 1.1 shows a rather clumsy manner of managing manually whichever errors may happen during the evaluation of an arithmetic expression. However, after a cursory examination of this extract, we may notice that much of the clumsiness may be factored away by adding an operation to check whether the result of an expression is Ok x, proceed with x if so and abort the operation otherwise. Indeed in the world of monads [19], this is the *binding* operation. In OCaml, this function is typically hidden behind syntactic sugar [4] perform and ⟵, which allows us to rewrite

Listing 1.3. Towards monadic error management

```
let bind m k = match m with
   Ok x     → k x
 | Error _ → m

let rec eval = function
   Value f        → Ok f
 | Div   (x, y) → perform x' ⟵ eval x ; y' ⟵ eval y ;
   if y' = 0. then Error "Division by 0"
   else             Ok (x' /. y')
```

```
  | Add   (x, y) → perform x' ⟵ eval x ; y' ⟵ eval y ;
    Ok (x' +. y')
(*val eval: expr → (float, string) result*)
```

For the sake of abstraction (and upcoming changes of implementation), we also hide the implementation of type (α, β) `result` behind two functions `return` (for successes) and `throw` (for failures):

Listing 1.4. Monadic error management

```
let return   x = Ok x
let throw    x = Error x

let rec eval = function
   Value f        → return f
 | Div  (x, y) → perform x' ⟵ eval x ; y' ⟵ eval y ;
   if y' = 0. then throw "Division by 0"
   else              return (x' /. y')
 | Add   (x, y) → perform x' ⟵ eval x ; y' ⟵ eval y ;
   return (x' +. y')
```

This new definition of `eval` is arguably as easy to read as the version of Listing 1.2. As we have decided to abstract away type `result`, we need functions to run a computation and determine it success. We call respectively these functions `catch` (for catching one error case) and `attempt` (for catching all error cases and entering/leaving the error monad):

Listing 1.5. Catching errors

```
let catch ~handle = function Ok x      → Ok x
                          | Error e → handle e

let attempt ~handle = function Ok    x → x
                            | Error e → handle e
```

Here we use a *labeled argument* `~handle` (instead of just `handle`) so that it can be called by explicitly stating its name, and even reordered:

```
attempt normal_case ~handle:error_case
```

We may now group all the functions of the error monad as one module `Error_monad` with the following signature:

Listing 1.6. A module for the error monad

```
type (+α, +β) result
val return: α → (α, β) result
val throw : β → (α, β) result
val bind  : (α, β) result → (α → (γ, β) result) →
   (γ, β) result
val catch  : handle:(β → (α, γ) result) →
   (α, β) result → (α, γ) result
val attempt: handle:(β → α) → (α, β) result → α
```

The + in (+α, +β) `result` specify that this type is covariant, *i.e.* an expression of type (α, β) `result` can be generalized even if it is not a value [8].

This adoption of monadic-style error reporting proves sufficient to convey polymorphic type information. As we will see in the next section, we may take advantage of this to improve flexibility of our library even further.

To convince ourselves of the capabilities of the error monad in terms of polymorphism, let us write a toy implementation of (persistent) association lists. The signature contains two functions: `find` to retrieve the value associated to a key, and `add` to add an association to the list. Operation `find k l` fails when nothing is associated to `k` in `l`, while `add k u l` fails when there is already a value `v` associated to `k` in `l`. In both case the key `k` is used as the error report.

We give the implementation of this module in Listing 1.7, together with its signature in Listing 1.8. As expected, the type of errors thrown by `add` and `find` depends on the type of the input.

Listing 1.7. Association list with polymorphic exceptions

```
type (α,β) assoc = (α*β) list

let empty = []

let rec add k u = function
   []              → return [k,u]
| (x, v as a)::l → if k=x then throw k
                    else        perform l' ⟵ add k u l ;
                                return (a::l')
let rec find k = function
   []           → throw k
| (x,v)::l → if k=x then return v
                else        find k l
```

Listing 1.8. Association list signature

```
type (α,β) assoc
val empty: (α,β) assoc
val add   : α → β → (α,β) assoc → ((α,β) assoc , α) result
val find : α → (α,β) assoc → (β,α) result
```

In the rest of the paper, we shall concentrate on the `eval` example. However, should need arise, this example could be refined similarly.

3 Representing errors

While Listing 1.4 presents a code much more usable than that of Listing 1.1 and while this listing is type-safe, the awful truth is that this safety hides a fragility, due to the use of "magic" character strings to represent the nature of errors – here, `"Division by 0"`, a constant which the type-checker cannot take into account when attempting to guarantee coverage. Unfortunately, this fragility is shared by elements of both OCaml's, SML's or Haskell's standard libraries.

Now, of course, it may happen that we need to represent several possible errors cases, along with some context data. For instance, during the evaluation of simple arithmetic exceptions, arithmetic overflow errors could also arise. For debugging purposes, we may even decide that each error should be accompanied by the detail of the expression which caused the error and that overflows should be split between overflows during addition and overflows during division.

The error monad of Section 2 handles propagation for any type β of errors. We shall investigate how to choose this type to handle those cases.

3.1 Errors as heavy-weight sums

The first and most obvious choice is to represent errors as sum types. For our running example, we could write

Listing 1.9. Simple arithmetic errors

```
type cause_of_overflow = Addition | Division
type eval_error = Division_by_zero of expr
                | Overflow  of expr * cause_of_overflow
```

Now, as our error monad lets us transmit arbitrary error information along with the error itself, we may rewrite `eval` so as to take advantage of `eval_error` instead of `string`, without having to declare a new exception constructor or to rewrite the interface or implementation of the error monad:

Listing 1.10. Monadic error management with sum types

```
let ensure_finite f e cause =
  if is_infinite f then throw (Overflow(e, cause))
  else                      return f

let rec eval = function
  Value f          → return f
| Div (x,y) as e → perform x' ⟵ eval x ; y' ⟵ eval y ;
    if y' = 0. then throw (Division_by_zero e)
    else              ensure_finite (x' /. y') e Division
| Add (x,y) as e → perform x' ⟵ eval x ; y' ⟵ eval y ;
    ensure_finite (x' +. y') e Addition
(*val eval: expr → (float, eval_error) result*)
```

While this solution improves on the original situation, it is not fully satisfying. Indeed, it is quite common to have several functions share some error cases but not all. For instance, a basic visual 10-digit calculator and a scientific plotter may share on a common arithmetic library. Both evaluators use division and may suffer from divisions by zero. However, only the scientific plotter defines logarithm and may thus suffer from logarithm-related errors. Should the error-reporting mechanism of the library be defined as one heavy-weight sum type, the visual calculator would need to handle all the same error cases as the scientific plotter. OCaml's built-in pattern-matching coverage test will therefore require

all error cases to be managed, even though the functions which may trigger these error cases are never invoked by the visual calculator.

The alternative is to use disjoint possible errors for distinct functions. However, this choice quickly leads to composability nightmares. Since a division by zero and a logarithm-error are members of two disjoint types, they need to be injected manually into a type `division_by_zero_or_log_error`, defined as a sum type, for use by the scientific plotter. While possible, this solution is cumbersome to generalize and tends to scale very poorly for large projects, especially during a prototyping phase. These composability troubles also appear as soon as two different libraries use disjoint types to represent errors: arithmetic errors, disk errors or interface toolkit errors, for instance, must then be injected into an awkward common type of errors, and projected back towards smaller types of errors as need arises.

3.2 Errors as lightweight composition of sums

Another approach, commonly seen in the Haskell world, and actually not very different from the second choice just mentioned, is to define a more general type along the lines of

Listing 1.11. Haskell-style `either` type

```
type (α, β) either = Left of α | Right of β
```

With such a data structure, building lightweight compositions of error cases becomes a trivial task. However, these lightweight compositions are also an easy recipe for obtaining unreadable constructions consisting in trees of `either` and tuples. That is, attempting to convert `eval` to use only such lightweight types typically results in the following expression, and its rather convoluted type:

Listing 1.12. Monadic error management with lightweight `either`

```
let ensure_finite f message =
  if is_infinte f then throw message
  else                    return f

let rec eval = function
    Value f          → return f
  | Div (x,y) as e → perform x' ←── eval x ; y' ←── eval y ;
     if y' = 0. then throw (Right e)
     else ensure_finite (x' /. y') (Left (Left e))
  | Add (x,y) as e → perform x' ←── eval x ; y' ←── eval y ;
     ensure_finite (x' +. y') (Left (Right e))
(* val eval : expr →
 ( float,( (expr, expr) either,expr) either ) result *)
```

While it is possible to avoid such chains of `either` by combining this approach with the manual definition of new types – perhaps abstracted away behind modules – the result remains unsatisfactory in terms of comprehension and falls far from solving the composability nightmare.

3.3 Errors as extensible types

Another alternative would be the use of *extensible types*, as featured in Alice ML [16]. More generally, one such type is available in languages of the ML family: native exceptions. Instead of our current type `eval_error`, and with the same code of `eval`, we could therefore define two native exceptions, whose role is to be raised monadically:

```
exception Division_by_zero of expr
exception Overflow         of expr * cause_of_overflow
```

If, at a later stage, the set of exceptions needs to be extended to take into account, say, logarithm errors, we may implement the new case with the following definition:

```
exception Logarithm_error  of expr
```

Better even, this solution proves compatible with the existing native exception system and permits trivial conversion of native exceptions for use with the error monad:

```
let attempt_legacy ~handle f arg = try   f arg
                                   with  e → handle e
(*val attempt_legacy: handle:(exn → β) → (α → β) → α → β*)
```

At this point, a first weakness appears: while the addition of brand new error cases such as `Logarithm_error` is a trivial task, extending `cause_of_overflow` is impossible unless we find a way to define `cause_of_overflow` as an extensible type. Assuming that we have a way to express several distinct extensible types, perhaps by using an hypothetical encoding with phantom types, we are still faced with a dilemma: should all errors be represented by items of the same type `exn` or should we use several disjoint extensible types? The question may sound familiar, as we have already been faced with the same alternative in the case of heavy-weight sum types. As it turns out, and for the same reasons, neither choice is acceptable: sharing one type gets into the way of coverage guarantees, while splitting into several types leads, again, to composability nightmares.

Or does it? After all, OCaml does contain an additional kind of types, close cousin to extensible sum types, but with much better flexibility: Polymorphic Variants [7].

3.4 Errors as polymorphic variant

Polymorphic variants represent a kind of lightweight sum types designed to maximize flexibility. Indeed, the main traits of polymorphic variants are that

1. no declaration is necessary – rather, definition is inferred from usage
2. declaration is possible, for specification and documentation purposes
3. open variants may be composed automatically into a larger variant
4. constructors may be shared between unrelated polymorphic variants.

When used to represent errors, trait 1. lets us concentrate on the task of building the algorithm, without having to write down the exact set of errors before the prototyping phase is over. Trait 2. proves useful at the end of the prototyping phase, to improve error-checking of client code and documentation, while 3. lets OCaml infer the set of errors which may be triggered by an expression – and check completeness of the error coverage, just as it would do for heavy-weight sum types. Finally, trait 4. lets us define functions which may share some – but not necessarily all – error cases.

Before rewriting the full implementation of `eval`, let us build a smaller example. The following extract defines an expression expression `div_by_zero` which `throw`s a division by zero with information `Value 0.`:

```
let div_by_zero = throw ('Division_by_zero (Value 0. ))
(* val div_by_zero : (α,
    [> 'Division_by_zero of expr ]) result *)
```

The type of `div_by_zero` mentions that it may have any result α, much like raising ML-exceptions produce results of type α, and that it may throw an error consisting in an *open* variant, marked by constructor 'Division_by_zero, and containing an `expr`.

Similarly, we may define an expression with the ability to cause an overflow during division, much as we could with a heavy-weight sum type:

```
let overflow_div = throw ('Overflow ('Division (Value 0.)))
(* val overflow_div : (_α, _[> 'Overflow of
    _[> 'Division of expr ] ] ) result *)
```

Finally, both expressions may be composed into an expression which might cause either a division by zero or an overflow during division, resulting in:

```
if true then div_by_zero
else        overflow_div
(* (_α, _[> 'Division_by_zero of expr
        | 'Overflow of _[> 'Division of expr ] ]) result*)
```

As we see from the inferred type of this expression, the result of the composition may produce either results (of any type) or errors marked either by 'Division_by_zero (and accompanied by an `expr`) or by 'Overflow (and accompanied by another tag 'Division, itself accompanied by an `expr`). This error signature remains *open*, which allows us to add further error cases.

As expected, converting `eval` to polymorphic variants is straightforward:

Listing 1.13. Monadic error management with polymorphic variants

```
let rec eval = function
   Value f          → return f
 | Div (x,y) as e → perform x' ←— eval x ; y' ←— eval y ;
    if y' = 0. then throw ('Division_by_zero e)
    else ensure_finite (x' /. y') ('Overflow ('Division e))
 | Add (x,y) as e → perform x' ←— eval x ; y' ←— eval y ;
     ensure_finite (x' +. y') ('Overflow ('Addition e))
```

```
(*val eval :
 expr → ( float,  [> ‘Division_by_zero of expr
    | ‘Overflow of [> ‘Addition of expr
                      | ‘Division of expr ] ] ) result *)
```

As we hoped, with polymorphic variants, we do not have to manually label error cases. Rather, the compiler may infer error cases from the source code. As this inferred information appears in the type of our function, coverage may be proved by the type-checker. Therefore, we may write:

```
let test1 e = attempt (eval e) ~handle:(function
    ‘Division_by_zero _→print_string ”Division by 0”; 0.
  | ‘Overflow _         →print_string ”Overflow”; 0.)
(*val test1 : expr → float*)
```

On the other hand, the following extract fails to compile:

```
# let test2 e = attempt (eval e) ~handle:(
  function ‘Overflow _ → print_string ”Overflow”; 0. )

  function ‘Overflow _ →
           ^^^^^^^^^^^^
This pattern matches values of type
  [< ‘Overflow of α ]
but is here used to match values of type
  [> ‘Division_by_zero of expr
    | ‘Overflow of [> ‘Addition of expr
                      | ‘Division of expr ] ]
```

In addition, the composability of polymorphic variants, which we have demonstrated, means that we do not have to decide whether to put all the error cases defined by a library in one common type or to split them among several disjoint types: barring any conflicting name or any specification which we may decide to add to prevent composition, there is no difference between one large polymorphic variant type and the automatically inferred union of several smaller ones.

Note that this choice of polymorphic variants does not alter the signature of our module, as featured in Listing 1.6. In particular, a consequence is that function **catch** can be used to eliminate one (or more) variant type exception while propagating others:

```
let ignore_overflow e = catch (eval e) ~handle:(function
    ‘Overflow _               → return 0.
  | ‘Division_by_zero _ as ex → throw ex)
(*val ignore_overflow: expr →
    (float, [> ‘Division_by_zero of expr]) result*)
```

Unfortunately, due to limitations on type inference of polymorphic variants, this elimination requires manual annotations and *a priori* knowledge of both the types of error cases which must be handled and the types of error cases which should be propagated:

```
let ignore_overflow e = catch (eval e) ~handle:(function
```

```
   'Overflow _                    → return 0.
|  _ as ex                       → throw ex)
(*val ignore_overflow: expr →
   ( float,   [> 'Division_by_zero of expr
              |  'Overflow of [> 'Addition of expr
                              | 'Division of expr ] ] ) result *)
```

While this limitation is perhaps disappointing, the upside is that, as we will now see, polymorphic variants and a little syntactic sugar may carry us farther than usual ML-style exceptions.

3.5   From polymorphic variants to exception hierarchies

We have just demonstrated how polymorphic variants solve the problem of composing error cases. Actually, our examples show a little bit more: we have not only defined two kinds of errors (divisions by zero and overflows), but also defined two sub-cases of errors (overflow due to addition and overflow due to division).

Passing the right parameters to function `attempt`, we may choose to consider all overflows at once, as we have done in our latest examples, or we may prefer to differentiate subcases of overflows:

Listing 1.14. Matching cases and subcases

```
let test3 e = attempt (eval e) ~handle:(function
   'Division_by_zero _   → print_string "Division by 0"; 0.
|  'Overflow 'Addition _ → print_string "+ overflows"; 0.
|  'Overflow _            → print_string "Other overflow"; 0.)
```

In other words, while we have chosen to present sub-cases as additional information carried by the error, we could just as well have decided to consider them elements of a small hierarchy:

  – division by zero is a class of errors ;
  – overflow is a class of errors ;
  – overflow through addition is a class of overflows ;
  – overflow through division is a class of overflows.

From this observation, we may derive a general notion of classes of errors, without compromising composability and coverage checking.

Before we proceed, we need to decide exactly what an exception class should be. If it is to have any use at all, it should be possible to determine if an exception belongs to a given class by simple pattern-matching. In order to preserve our results and the features used up to this point, an exception class should be a data structure, defined by one or more polymorphic variant constructors and their associated signature, as well as some error content. In addition, for exception classes to be useful, it must be possible to specify a subtyping relation between classes. We also need to ensure consistency between the error content of classes related by subtyping. Finally, we should be able to define new classes and subclasses without having to modify the definition of existing code.

To achieve all this, we encode classes using a method comparable to tail polymorphism [2] with polymorphic variants[5]. Where classical uses of tail polymorphism take advantage of either algebraic data-types or records, though, the use of polymorphic variants preserves extensibility.

We first introduce a chain-link record, whose sole use is to provide human-readable field names `sub` and `content`. Field `sub` is used to link a class to its super-class, while field `content` serves to record the class-specific additional error information which the programmer wishes to return along with the error:

```
type (α, β) ex =  {
   content: α;
   sub:      β option;
} constraint β = [> ]
```

The `constraint` β = `[> ]` line binds β to being a polymorphic variant row.

From this, assuming for the course of this example that division by zero is a top-level class of exceptions, we may produce the following constructor:

```
let division_by_zero_exc ?sub content =
 'Division_by_zero {
    content = content;
    sub     = sub;  }
(*val ?sub:([>  ] as α) → β →
    [> 'Division_by_zero of (β, α) ex]*)
```

Argument `content` is self-documented, while `sub` will serve for subclasses to register the link. `?sub` means that it is *optional*, it is realised as a labeled argument of an `option` type.

Similarly, we may now define overflow:

```
let overflow_exc ?sub content =
 'Overflow {
    content = content;
    sub     = sub; }
```

Since we decided to describe overflow during addition as a subclass of overflow, we may define its constructor by chaining a call to `overflow_exc`, passing the new chain-link as argument.

```
let overflow_addition ?sub content =
 overflow_exc ~sub:('Addition {
    content = ();
    sub     = sub;
 }) content
```

Or, equivalently, with a small piece of syntactic sugar introduced to increase readability of exception definitions:

```
let exception overflow_division content =
    Overflow content; Division ()
```

---

[5] A similar idea has been suggested in the context of Haskell [11] but discarded as a "very interesting, but academic" and a "failed alternative".

The changes to the library are complete. Indeed, one simple record type is sufficient to move from polymorphic variants to polymorphic variants with hierarchies. To confirm our claim that we preserve composability and coverage guarantees, let us revisit `eval` and our test cases.

Adapting `eval` to our hierarchy is just the matter of replacing concrete type constructors with abstract constructors:

Listing 1.15. Eval with hierarchies

```
let rec eval = function
  Value f          → return f
| Div (x,y) as e → perform x' ⟵ eval x ; y' ⟵ eval y ;
    if y' = 0. then throw (division_by_zero_exc e)
    else ensure_finite (x' /. y') (overflow_division e)
| Add (x,y) as e → perform x' ⟵ eval x ; y' ⟵ eval y ;
    ensure_finite (x' +. y') (overflow_addition e)
(* val eval : expr →
   (float,
    [> 'Division_by_zero of (expr, α) ex
    |  'Overflow of
      (expr,
       [> 'Addition of (unit, β) ex
       |  'Division of (unit, γ) ex ])
     ex ]) result *)
```

While the type information is painful to read – and could benefit from some improved pretty-printing – it accurately reflects the possible result of `eval`, the nature of exceptions and subexceptions and their contents.

Adapting the test of Listing 1.14 to our more general framework, we obtain the following extract, slightly more complex:

```
let test4 e = attempt (eval e) ~handle:(function
  'Division_by_zero _    →
  print_string "Division by 0"; 0.
| 'Overflow {sub = Some ('Addition _)} →
  print_string "+ overflows"; 0.
| 'Overflow _            →
  print_string "Other overflow"; 0. )
```

For convenience, we introduce another layer of syntactic sugar, marked by a new keyword `attempt`, and which provides a simpler notation for exception patterns. With this sugar, we may rewrite the previous listing as

```
let test5 e = attempt eval e with
  Division_by_zero _      → print_string "Division by 0" ; 0.
| Overflow _; Addition _ → print_string "+ overflows"    ; 0.
| Overflow _              → print_string "Other overflow"; 0.
```

While we do not demonstrate this due to space limitations, this syntactic sugar proves also useful to add optional post-treatment for successes, for failures and for any result (à la `finally`).

With this simplified notation, let us demonstrate coverage guarantees by omitting the case of overflow division:

```
let test6 e = attempt eval e with
   Division_by_zero _      → print_string "Division by 0"; 0.
 | Overflow _; Addition _ → print_string "+ overflows"  ; 0.
```

The following error message demonstrates that the type-checker has correctly detected the missing case (and integrates well with our syntax extension). The solution is suggested at the end of the message:

Listing 1.16. Missing subcase (error message)

```
   Division_by_zero _     →
   ^^^^^^^^^^^^^^^^^^^
This pattern matches values of type
  [< 'Division_by_zero of α
   |  'Overflow of (expr,
               [< 'Addition of β ]) ex ]
but is here used to match values of type
 [> 'Division_by_zero of (expr, _) ex
  |   'Overflow of
   (expr,
     [> 'Addition of (unit, γ) ex
     |  'Division of (unit, δ) ex ])
   ex ]
The first variant type does not allow tag(s) 'Division
```

Similarly, our encoding lets the type-checker spot type or tag errors in exception-matching, as well as provide warnings in case of some useless catch clauses. We do not detail these error/warning messages, which are not any more readable than the one figuring in Listing 1.16, and which could just as well benefit from some customized pretty-printing.

3.6   Bottom line

In this section, we have examined a number of possible designs for error reports within the error monad. Some were totally unapplicable, some others were impractical. As it turns out, by using polymorphic variants, we may achieve both inference of error cases, composability of error cases and simple definition of hierarchies of error classes, while retaining the ability of the type-checker to guarantee coverage of all possible cases. All of this is implemented in a meager 29 lines of code, including the module signature.

At this point, we have obtained all the features we intended to implement. Our next step is to study the performance cost – and to minimize it, if possible.

4   Measuring performance

According to our experience, when hearing about monads, typical users of OCaml tend to shrug and mutter something about breaking performance too much to be as useful as built-in exceptions. Is that true?

At this point, we set out to measure the speed of various representations of errors in OCaml and to search for alternative implementations and variants of the error monad which would let us improve performances. In the course of this quest, we tested several dozens of versions, using sometimes purely functional techniques, mostly imperative primitives, type-unsafe optimizations, compile-time rewriting. Due to space constraints, we can only present a short summary of our results in this document. A more complete overview will be made available as a technical report.

The numbers are presented in Figure 1. The six techniques investigated are

ad-hoc management the technique used in Listing 1.1
native exceptions the technique used in Listing 1.2
pure error monad the error monad demonstrated so far
references and exceptions an implementation of the error monad using native exceptions to handle error propagation:

```
exception Raised
type (α, β) result = β option ref → α
let attempt ~handle v =
 let  result = ref None in
 try  v result
 with Raised → match !result with
     None    → assert false (* Unused *)
   | Some e → handle e
let bind m k r = k (m r) r
let return x _ = x
let throw  x b = b := Some x; raise Raised
```

Results are implemented as functions which raise a native exception when an error is found. As native exceptions in OCaml cannot be polymorphic, we pass the error message in a reference, that we give as an argument to the result-function. The reference actually contains a $β$ `option` to be able to initialize it.

phantom types and exceptions another implementation of the error monad using exceptions. Instead of using a reference as a "channel" to convey the error message, we make use of the type-unsafe `Obj.obj`. And make use of phantomtypes [6] to make sure that it is only applied in safe situations:

```
exception Raised of Obj.t
type (α, β) result = unit → α constraint β = [> ]
let attempt (v:(_,β) result) ~handle =
 try  v ()
 with Raised r → handle (Obj.obj r : β)
let bind m k () = k ( m () ) ()
let return x () = x
let throw  x () = raise (Raised (Obj.repr x))
```

pure error monad with rewriting a variation on the error monad implemented in Section 2. The monadic `bind` primitive is meant to be used mostly in the form `p ⟵ m; e`, a sort of `let ... in` construction which

is interpreted as `bind m (fun p -> e)`. This produces an unnecessary intermediary closure which OCaml compiler does not optimize out. We adress that issue by extending the monad primitives with a preprocessor function `rewrite_bind`:

```
val rewrite_bind: m:Ast.expr → p:Ast.patt →
                  e:Ast.expr → Ast.loc  → Ast.expr
```

We shall use this function directly as the interpretation of p ⟵ m; e, instead of building a term out of traditionnal monadic combinators.

```
let rewrite_bind ~m ~p ~e _loc =
  <:expr< match $m$ with
            | Ok $p$ → $e$
            | Error err → Error err >>
```

Notice that using this function forces to expose the implementation of (α,β) `result`. An alternative solution would be instead to expose another type (α,β) `_result` which happens to be the implementation of (α,β) `result`. Then `rewrite_bind` can convert between both using `Obj.magic`. The beauty of this approach is that even if after preprocessing the code contains unsafe type-casts, these are inserted by the preprocessor at safe places, with no need to ever break the monad abstraction[6].

To obtain this benchmark, we implemented using each technique

an arithmetic evaluator errors are rather uncommon, being raised only in case of division by zero (300 samples)

the n queens problem only one pseudo-error is raised, when a solution has been found (5 samples)

union of applicative sets of integers pseudo-errors are raised very often to mark the fact that no change is necessary to a given set (300 samples).

Every test has been performed with OCaml 3.10.2, under Linux, on native code compiled for the 32bit x86 platform, with maximal inlining, executed 15 times, after a major cycle of garbage-collection, with the best and the worst result discarded. The results are presented as a percentage of the number of samples in which the execution time falls within given bounds:

Very good Execution time of the sample is within 5% of the execution time of the fastest implementation for this test (the "shortest execution time")

Good Within 5-15% of the shortest execution time.

Acceptable Within 15-30% of the shortest execution time.

Slow Within 30-50% of the shortest execution time.

Bad At least 50% longer than the shortest execution time.

For information, we also provide

Average Average of ratio $\dfrac{\text{duration of test}}{\text{shortest execution time}}$.

Deviation Standard deviation of $\dfrac{\text{duration of test}}{\text{shortest execution time}}$.

	Evaluator	Queens	Union		Evaluator	Queens	Union
Ad-hoc error management				Reference and native exceptions			
Very good	56%	40 %	18%	Very good	0%	0 %	0%
Good	26%	60 %	43%	Good	7%	0 %	0%
Acceptable	12%	0 %	35%	Acceptable	33%	0 %	0%
Slow	3%	0 %	4%	Slow	41%	0 %	0%
Bad	3%	0 %	0%	Bad	19%	100%	100%
Average	1.06	1.05	1.13	Average	1.35	1.75	2.26
Deviation	0.12	0.04	0.10	Deviation	0.20	0.06	0.23
Native exceptions				Phantom types and native exceptions			
Very good	70%	100%	100%	Very good	1%	0 %	0%
Good	16%	0 %	0%	Good	8%	0 %	0%
Acceptable	12%	0 %	0%	Acceptable	39%	0 %	0%
Slow	2%	0 %	0%	Slow	35%	0 %	0%
Bad	0%	0 %	0%	Bad	17%	100%	100%
Average	1.06	1.00	1.00	Average	1.35	1.73	2.22
Deviation	0.13	0.00	0.00	Deviation	0.22	0.06	0.22
Pure error monad				Pure error monad with rewriting			
Very good	37%	0 %	0%	Very good	54%	0 %	0%
Good	35%	20 %	0%	Good	28%	100%	0%
Acceptable	18%	60 %	14%	Acceptable	12%	0%	5%
Slow	7%	20 %	71%	Slow	5%	0 %	56%
Bad	3%	0 %	15%	Bad	1%	0 %	38%
Average	1.12	1.24	1.48	Average	1.07	1.07	1.48
Deviation	0.14	0.02	0.14	Deviation	0.15	0.01	0.14

Fig. 1. Testing the performance of the error monad

The first conclusion we draw from our measurements is that the pure error monad is inappropriate as a mechanism for optimising returns. While this is unsurprising, we also notice that the speed of the pure monad is actually quite reasonable when it is used to deal with errors, and can be largely improved by plugging-in a little compiler support. Further experiments with improvements, which go beyond the scope of this document, hint that slightly more complex rewriting rules can go even further – and not just for error monads. By opposition, our every attempt to domesticate native exceptions into a construction which could be checked for complete coverage incurred an impressive slowdown which made them useless. Our most successful attempts at native exceptions, which also required plugged-in compiler support, remained a little slower than the pure error monad with rewriting.

At this point, the fastest implementation of our library consists in the pure error monad (29 lines), compile-time optimisations (49 lines) and some (larger) syntactic sugar. To handle circumstances in which exceptions are used as op-

---

[6] using this technique we have also built a variant of the "phantom types and exceptions" approach, which it proved to be less efficient

timised returns, we complete this with a 16 lines-long module, which provides a mechanism of local exceptions with polymorphism, while retaining the speed of native exceptions and adding a little safety. This module can also make use of hierarchical exceptions but is neither safe nor generic enough to replace the error monad.

5 Conclusion

We have demonstrated how to design an error-reporting mechanism for OCaml extending the exception semantics of ML, without altering the language. With respect to OCaml's built-in exception mechanisms, our work adds static checks, polymorphic error reports, hierarchies and support for locally-defined exceptions and relaxes the need of declaring error cases, while retaining a readable syntax and acceptable performance.

To obtain this richer mechanism, we make use of monads, polymorphic variants and code rewriting and demonstrate the folk theorem of the OCaml community that polymorphic variants are a more generic kind of exceptions. We have also attempted to optimize code through numerous techniques, and succeeded through the use of compile-time domain-specific code generators.

Related works Other works have been undertaken with the objective of making exceptions safer or more flexible. Some of these approaches take the form of compile-time checkers such as OCamlExc [14] or similar works for SML [20]. These tools perform program analysis and thus need to evolve whenever the language's semantic does; their maintenance can be quite involved. Similarly, the Catch tool for Haskell [13] uses abstract interpretation to provide guarantees that pattern matches of a program (including pattern-matching upon errors) suffice to cover all possible cases, even when individual pattern-matches are not exhaustive. All these tools retain the exception mechanism of the underlying language and therefore add no new feature, in particular no hierarchies of error classes.

Other efforts are closer to our approach. In addition to the very notion of monads [19], the Haskell community has seen numerous implementations of extensible sets of exceptions, either through monad transformers or dynamic type reflection. Hierarchical exceptions [11] through typeclass hierarchies and dynamic type reflection have also been implemented for Haskell. These choices could have been transposed and sometimes improved into OCaml. We decided to avoid monad transformers in the simple case of error reporting, as these too often require manual definition and manual composition of program-specific or library-specific error cases. Similarly, several variants on run-time type information are possible in OCaml, either with dynamic type reflection comparable to Haskell's `Data.Typeable`, or combinations of view patterns and either dynamically typed objects or lightweight extensible records, all of which have been implemented for OCaml. However, we preferred avoiding these dynamic typing solutions which, as their Haskell counterpart, forbid any automatic coverage-check. Yet another

encoding of hierarchies has been demonstrated for ML languages, through the use of phantom types [6]. While this work is very interesting, our experiments seem to show that the use of this encoding for exceptions leads to a much less flexible and composable library, in which new sub-kinds of errors cannot be added post-facto to an existing design.

Another combination of native exceptions and monad-like combinators for fast error reporting has been designed in the context of ML [15]. While benchmarks obtained with this discipline indicate better performance than what we achieve, this work aims only at reducing redundant error messages and does not improve the flexibility or safety of error management. This difference in purpose allows an efficient mix of native exceptions and monadic ones. Perhaps more interestingly, the author demonstrates a set of error-related function types which may not be obtained with pure monads, such as:

$$((\alpha \rightarrow \beta) \rightarrow \gamma \rightarrow \delta) \rightarrow (\alpha \rightarrow (\beta, \epsilon) \ \texttt{result}) \rightarrow \gamma \rightarrow (\delta, \epsilon) \ \texttt{result}$$

This combinator, which extends traditional function application to handle erroneous arguments, requires native exceptions, and hence cannot be implemented in our pure monad. It may however be implemented trivially with the complementary library we designed for local exceptions.

Numerous other works focus on performance in ML languages and their kin. In particular, the Glasgow Haskell Compiler is usually able to efficiently inline simple functions or rewrite simple compositions – as the OCaml compiler can do neither automatically in our case, this is what we implement manually to optimize away needless abstractions. As for the technique we employ for performing this inlining, it is essentially a specialized form of multi-stage compilation, as available in MetaOCaml [5] or outside the ML community [9]. In particular, our use of specialized code rewriters to avoid the cost of abstraction is an idea also found in MetaOCaml-based works [3].

Future works While in this document we have investigated only the problem of exception management, our work has yielded ideas which we believe may be applied to the broader domains of effect regions [17]. Indeed, we have started working on an implementation for OCaml through a combination of libraries and syntactic redefinitions. While early experiments seem to indicate that the limitations of type inference on polymorphic variants will limit inference of effect regions, we hope our implementation incurs only a negligible runtime penalty and allows comfortable interactions with existing libraries.

Part of our work on exceptions will be included in OCaml Batteries Included, the community-maintained standard library replacement for OCaml. Most of the modules of this library are being designed to permit monadic error management.

Acknowledgements

References

1. Matthias Blume, Umut Acar, and Wonseok Chae. Exception handlers as extensible cases. In *APLAS '08: 6th Asian Symposium on Programming Languages and Systems*, pages 273–289, Berlin, Heidelberg, 2008. Springer-Verlag.
2. F. Warren Burton. Type extension through polymorphism. *ACM Trans. Program. Lang. Syst.*, 12(1):135–138, 1990.
3. Jacques Carette and Oleg Kiselyov. Multi-stage programming with functors and monads: Eliminating abstraction overhead from generic code. In *GPCE*, pages 256–274, 2005.
4. Jacques Carette, Lydia E. van Dijk, and Oleg Kiselyov. Syntax extension for monads in ocaml. `http://www.cas.mcmaster.ca/~carette/pa_monad/`.
5. Krzysztof Czarnecki, John T. O'Donnell, Jörg Striegnitz, and Walid Taha. Dsl implementation in metaocaml, template haskell, and c++. In *Domain-Specific Program Generation*, pages 51–72, 2003.
6. Matthew Fluet and Riccardo Pucella. Phantom types and subtyping. In *TCS '02: Proceedings of the IFIP 17th World Computer Congress - TC1 Stream / 2nd IFIP International Conference on Theoretical Computer Science*, pages 448–460, Deventer, The Netherlands, The Netherlands, 2002. Kluwer, B.V.
7. Jacques Garrigue. Programming with polymorphic variants. In *ML Workshop*, 1998.
8. Jacques Garrigue. Relaxing the value restriction. In *International Symposium on Functional and Logic Programming, Nara, LNCS 2998*. Springer-Verlag, 2004.
9. Samuel Z. Guyer and Calvin Lin. An annotation language for optimizing software libraries. In *DSL*, pages 39–52, 1999.
10. Richard C. Holt and David B. Wortman. A sequence of structured subsets of pl/i. *SIGCSE Bull.*, 6(1):129–132, 1974.
11. Simon Marlow. An extensible dynamically-typed hierarchy of exceptions. In *Haskell '06: ACM SIGPLAN workshop on Haskell*. ACM Press, September 2006.
12. Robin Milner, Mads Tofte, and David Macqueen. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1990.
13. Neil Mitchell and Colin Runciman. A static checker for safe pattern matching in Haskell. In *Trends in Functional Programming*, volume 6. Intellect, February 2007.
14. François Pessaux. *Détection statique d'exceptions non rattrapées en Objective Caml*. PhD thesis, Université Pierre & Marie Curie - Paris 6, 2000.
15. Norman Ramsey. Eliminating spurious error messages using exceptions, polymorphism, and higher-order functions. *The Computer Journal*, 42(5):360–372, 1999.
16. Andreas Rossberg, Didier Le Botlan, Guido Tack, Thorsten Brunklaus, and Gert Smolka. Alice through the looking glass. *Trends in Functional Programming*, 5:77–96, 2006.
17. Jean-Pierre Talpin and Pierre Jouvelot. The type and effect discipline. *Inf. Comput.*, 111(2):245–296, 1994.
18. David Teller, Arnaud Spiwack, and Till Varoquaux. Catch me if you can. Software package available at `http://www.univ-orleans.fr/lifo/Members/David.Teller/software/exceptions/catch_0_2.tgz`.
19. Philip Wadler. The essence of functional programming. In *POPL '92: Proceedings of the 19th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 1–14, New York, NY, USA, 1992. ACM.
20. Kwangkeun Yi and Sukyoung Ryu. A cost-effective estimation of uncaught exceptions in standard ml programs. *Theor. Comput. Sci.*, 277(1-2):185–217, 2002.