Type-based analysis of uncaught exceptions

François Pessaux, Xavier Leroy

► To cite this version:

François Pessaux, Xavier Leroy. Type-based analysis of uncaught exceptions. POPL 1999: 26th symposium Principles of Programming Languages, ACM, Jan 1999, San Antonio, United States. pp.276 - 290, 10.1145/292540.292565 . hal-01499959

HAL Id: hal-01499959

<https://hal.inria.fr/hal-01499959>

Submitted on 1 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Type-based analysis of uncaught exceptions

François Pessaux Xavier Leroy
INRIA Rocquencourt*

Abstract

This paper presents a program analysis to estimate uncaught exceptions in ML programs. This analysis relies on unification-based type inference in a non-standard type system, using rows to approximate both the flow of escaping exceptions (a la effect systems) and the flow of result values (a la control-flow analyses). The resulting analysis is efficient and precise; in particular, arguments carried by exceptions are accurately handled.

1 Introduction

Many modern programming languages such as Ada, Modula-3, ML and Java provide built-in support for exceptions: raising an exception at some program point transfers control to the nearest handler for that exception found in the dynamic call stack. Exceptions provide safe and flexible error handling in applications: if an exception is not explicitly handled in a function by the programmer, it is automatically propagated upwards in the call graph until a function that “knows” how to deal with the exception is found. If no handler is provided for the exception, program execution is immediately aborted, thus pinpointing the unexpected condition during testing. This stands in sharp contrast with the traditional C-style reporting of error conditions as “impossible” return values (such as null pointers or the integer -1): in this approach, the programmer must write significant amount of code to propagate error conditions upwards; moreover, it is very easy to ignore an error condition altogether, often causing the program to crash much later, or even complete but produce incorrect results.

The downside of using exceptions for error reporting and as a general non-local control structure is that it is very easy to forget to catch an exception at the right place, i.e. to handle an error condition. ML compilers generate no errors or warnings in this case, and the programming mistake will only show up during testing. Exhaustive testing of applications is difficult, and even more so in the case of error conditions that are infrequent or hard to reproduce. Our experience with large ML applications is that uncaught exceptions are the most frequent mode of failure.

*Authors’ address: INRIA Rocquencourt, projet Cristal, B.P. 105, 78153 Le Chesnay, France. E-mail: Francois.Pessaux@inria.fr, Xavier.Leroy@inria.fr. This work has been partially supported by CNET, France Télécom.

To address this issue, languages such as Modula-3 and Java require the programmer to declare, for each function or method, the set of exceptions that may escape out of it. Those declarations are then checked statically during type-checking by a simple intraprocedural analysis. This forces programmers to be conscious of the flow of exceptions through their programs.

Declaring escaping exceptions in functions and method signatures works well in first-order, monomorphic programs, but is not adequate for the kind of higher-order, polymorphic programming that ML promotes. Consider the `map` iterator on lists. In Modula-3 or Java, the programmer must declare a set E of exceptions that the function argument to `map` may raise; `map`, then, may raise the same exceptions E . But E is fixed arbitrarily, thus preventing `map` from being applied to functions that raise exceptions not in E . The genericity of `map` can be restored by taking for E the set of all possible exceptions, but then the precision of the exception analysis is dramatically decreased: all invocations of `map` are then considered as potentially raising any exception. (Similar problems arise in highly object-oriented Java programs using e.g. container classes and iterators intensively.) To deal properly with higher-order functions, a very rich language for exception declarations is required, including at least exception polymorphism (variables ranging over sets of exceptions) and arbitrary unions of exception sets. (See section 2 for a more detailed discussion.) We believe that such a complex language for declaring escaping exceptions is beyond what programmers are willing to put up with.

The alternative that we follow in this paper is to infer escaping exceptions from unannotated ML source code. In other terms, we view the problem of detecting potentially uncaught exceptions as a static debugging problem, where static analyses are applied to the programs not to make them faster via better code generation, but to make them safer by pinpointing possible run-time failures. This approach has several advantages with respect to the Modula-3/Java approach: it blends better with ML’s type inference; it does not change the language and supports the static debugging of “legacy” applications; it allows the use of complex approximations of exception sets, as those need not be written by the programmer (within reason – the results of the analysis must still be understandable to the programmer). Finally, the exception inference needs not be fully compatible with the ML module system: a whole program analysis can be considered (again within reason – analysis time should remain practical).

Several exception analyses for ML have been proposed [8, 36, 35, 3, 4], some based on effect systems, some on control-flow analyses, some on combinations of both (see section 6 for a detailed discussion). The analysis presented in this paper attempts to combine the efficiency of effect systems with the precision of flow analyses. It is based on unification and non-standard type inference algorithms that have

excellent running time and we hope should scale well to large applications. At the same time, our analysis is still fairly precise; in particular, it approximates not only the names of the escaping exceptions, but also the arguments they carry – a feature that is essential to analyze precisely many existing ML programs. This constitutes the main technical contribution of this paper: integrate in the same unification-based framework both approximation of exception effects in the style of effect systems [28], and approximation of sets of values computed at each program point in the style of flow analyses and soft typing [26, 32]. Finally, our analysis has been implemented to cover the whole Objective Caml language – not only core ML, but also datatypes, objects, and the module system. We present some preliminary experimental results obtained with our implementation.

The remainder of this paper is organized as follows. Section 2 lists the main requirements for an ML exception analysis. Section 3 presents the non-standard type system we use for exception analysis. Extension to the full Objective Caml language is discussed in section 4; experimental results obtained with our implementation, in section 5; and related work, in section 6. Concluding remarks can be found in section 7.

2 Design requirements

In this section, we list the main requirements for an effective exception analysis for ML, and show that they go much beyond what can be expressed by exception declarations in Modula-3 or Java. Existing exception analyses have addressed some of these requirements, but none addresses all.

2.1 Handling higher-order functions precisely

The exception behavior of higher-order functions depends on the exceptions that can be raised by their functional arguments. A form of polymorphism over escaping exceptions is thus needed to analyze higher-order functions precisely. Consider the `map` iterator over lists mentioned in introduction. An application `map f l` may raise whatever exception the `f` argument may raise. Writing $\tau \xrightarrow{\varphi} \tau'$ for the annotated type of functions from type τ to type τ' whose set of potentially escaping exception is φ , the behavior of `map` is captured by the following annotated type scheme:

$$\text{map} : \forall \alpha, \beta, \varphi. (\alpha \xrightarrow{\varphi} \beta) \xrightarrow{\emptyset} (\alpha \text{ list} \xrightarrow{\varphi} \beta \text{ list})$$

where α, β range over types and φ ranges over sets of exceptions. In general, the escaping exceptions for a higher-order function are combinations $\varphi_1 \cup \dots \cup \varphi_n \cup \{C_1; \dots; C_n\}$ where the φ_i are variables representing the escaping exceptions for functional arguments and the C_j are exception constants. For instance, we have the following annotated type for function composition $\lambda f. \lambda g. \lambda x. f(g(x))$:

$$\forall \alpha, \beta, \gamma, \varphi, \psi. (\alpha \xrightarrow{\varphi} \beta) \xrightarrow{\emptyset} (\gamma \xrightarrow{\psi} \alpha) \xrightarrow{\emptyset} \gamma \xrightarrow{\varphi \cup \psi} \beta$$

Given the frequent use of higher-order functions in ML programs, an exception analysis for ML must handle them with precision similar to what the annotated types above suggest.

Similar issues arise when functions are stored into data structures such as lists or hash tables (as in callback tables for instance). The exception analysis should keep track of the union of the exceptions that can be raised by functions contained in the structure. It is not acceptable to say that

any exception can be raised by applying a function retrieved from the structure.

2.2 Handling exceptions as first-class values

In ML and Java, exceptions are first-class values: exception values can be built in advance and passed through functions before being raised. Consider for instance the following contrived example:

```
let test = λexn. try raise(exn) with E → 0
```

The exception behavior of this function is that `test exn` raises the exception contained in the argument `exn`, except when `exn` is actually the exception `E`, in which case no exception escapes out of `test`. We seek exception analyses precise enough to capture this behavior.

It is true that the first-class character of exception values is rarely, if ever, used in actual ML programs. However, there is one important idiom where an exception value appears: finalization. Consider:

```
let f = λx. try g(x)
           with E → 0
           | exn → finalization code; raise(exn)
```

Assuming `g` can raise exceptions `E` and `E'`, the exception analyzer should recognize that the `exn` exception variable can only take the value `E'`, thus the `raise(exn)` that re-raises the exception after finalization can only raise `E'`, and so does the function `f` itself.

2.3 Keeping track of exception arguments

ML exceptions can optionally carry arguments, just like all other data type constructors. This argument can be tested in the `with` part of an exception handler, using pattern-matching on the exception value, so that only certain exceptions with certain arguments are caught. Consider the following example:

```
exception Failure of string
let f = λx. if ... then ... else raise(Failure "f")
let g = λx. try f(x) with Failure "f" → 0
```

An exception analysis that only keeps track of the exception head constructors (i.e. `Failure` above) but not of their arguments (i.e. the string `"f"` above) fails to analyze this example with sufficient precision: the analysis records that function `f` may raise the `Failure` exception, hence it considers that the application `f(x)` in `g` may raise `Failure` with any argument. Since the exception handler traps only `Failure "f"`, the analyzer concludes that `g` may raise `Failure`, while in reality no exception can escape `g`.

This lack of precision can be brushed aside as “unimportant” and “bad programming style anyway”. Indeed, the programmer should have declared a specific constant exception `Failure_f` to report the error in `f`, rather than rely on the general-purpose `Failure` exception. However, code fragments similar to the example above appear in legacy Caml applications that we would like to analyze. More importantly, there are also legitimate uses of exceptions with parameters. For instance, the Caml interface to Unix system calls uses the following scheme to report Unix error conditions:

```
type unix_error = EACCES | ENOENT | ENOSPC | ...
(* enumerated type with 67 constructors
   representing Unix error codes *)
exception Unix_error of unix_error
```

This allows user code to trap all Unix errors at once (`try ... with Unix_error(_) -> ...`), and also to trap particular errors (`try ... with Unix_error(ENOENT) -> ...`). Replacing `Unix_error` by 67 distinct exceptions, one for each error code, would make the former very painful. It is desirable that the exception analysis be able to show that certain `Unix_error` exceptions with arguments representing common errors (e.g. `Unix_error(ENOENT)`, “no such file”) are handled in the program and thus do not escape, while we can accept that other `Unix_error` exceptions representing rare errors are not handled in the program and may escape.

The problem with exception arguments is made worse by the availability (in the Caml standard library at least) of predefined functions to raise general-purpose exceptions such as `Failure` above. Indeed, the example with `Failure` above is more likely to appear under the following form:

```
exception Failure of string
let failwith = λmsg. raise(Failure msg)
let f = λx. if ... then ... else failwith("f")
let g = λx. try f(x) with Failure "f" -> 0
```

Precise exception analysis in this example requires tracking the string constant “f” not only when it appears as immediate argument to the `Failure` exception constructor, but also when it is passed to the function `failwith`. Hence the exception analysis must also include some amount of data flow analysis, not limited to exception values.

2.4 Running faster than control-flow analyses

All the requirements we have listed so far point towards control-flow analyses in the style of Shiver’s *k*-CFA [26] or Heintze’s SBA [9]. Control-flow analyses provide an approximation of the set of values that can flow to each program point. It is entirely straightforward to extend them to approximate also the set of escaping exceptions at each program point at the same time as they approximate the set of result values. Alternatively, the exception analysis can be run as a second pass of dataflow analysis exploiting the results of control-flow analysis [35], although this results in some loss of precision, as the control flow can be determined more accurately if exception information is available. This exception analysis benefits from the relatively precise approximation of values provided by the control-flow analysis, especially as far as exception arguments are concerned.

Our first implementation of an exception analyzer for Objective Caml was indeed based on control-flow analysis: 0-CFA initially, then Jagannathan and Wright’s “polymorphic splitting” [12]. Our practical experience with this approach was mixed: the precision of the exception analysis was satisfactory (at least with polymorphic splitting), but the speed of the analysis left a lot to be desired. In particular, we observed quadratic behavior on several examples, indicating that the analysis would not scale easily to large programs¹. Although sophisticated techniques have been developed to speed up program analyses based of set inclusion constraints such as CFA and SBA [2, 6, 5, 19], it is still an open problem whether those analyses can scale to 100,000-line programs.

¹ The complexity of 0-CFA alone is $O(n^3)$, where n is the size of the whole program. We did not observe cubic behavior on our tests, however. Quadratic behavior arises in the following not uncommon case: assume that a group of functions of size $k = O(n)$ recurses over a list of $m = O(n)$ elements given in extension in the program source. At least m iteration of the analysis is required before fixpoint is reached on the parameters and results of the functions. Since each iteration takes time proportional to k , the time of the analysis is $O(n^2)$.

For these reasons, we decided to abandon analyses based on CFA or more generally set inclusion constraints, and settled for less precise but faster analyses based on equality constraints and unification.

3 A type system for exception analysis

In the style of effect systems [16, 28], our exception analysis is presented as a type inference algorithm for a non-standard type system. The type system uses unified mechanisms based on row variables both to keep track of the effects (sets of escaping exceptions) of expressions and to refine the usual ML types by more precise information about the possible values of expressions. In this section, we present first the typing rules for our type system (that is, the specifications for the exception analysis), then type inference issues (the actual analysis).

3.1 The source language

The source language we consider in this paper is a simple subset of ML with integers and exceptions as the only data types, the ability to raise and handle exceptions, and simplified pattern-matching.

Terms:

$a ::= x$	identifier
$ i$	integer constant
$ \lambda x. a$	application
$ a_1(a_2)$	abstraction
$ \text{let } x = a_1 \text{ in } a_2$	the let binding
$ \text{match } a_1 \text{ with } p \rightarrow a_2 \mid x \rightarrow a_3$	pattern-matching
$ C \mid D(a)$	exception constr.
$ \text{try } a_1 \text{ with } x \rightarrow a_2$	exception handler

Patterns:

$p ::= x$	variable pattern
$ i \mid C$	constant patterns
$ D(p)$	constructed pattern

The `match  $a_1$  with  $p \rightarrow a_2 \mid x \rightarrow a_3$` performs pattern-matching on the value of a_1 ; if it matches the pattern p , the branch a_2 is evaluated; otherwise, a_3 is evaluated. Multi-case pattern matchings can be expressed by cascading `match` expressions. The `try  $a_1$  with  $x \rightarrow a_2$` construct evaluates a_1 ; if an exception is raised, its value is bound to x and a_2 is evaluated. There is no syntactic form for raising an exception; instead, we assume predefined a `raise` function in the environment. The `try` construct catches all exceptions; catching only a given exception C is performed by:

```
try a1 with x → match x with C → a2 | y → raise(y)
```

The dynamic semantics for this language is given by the reduction rules in figure 1, in the style of [33]. Values, evaluation contexts, and evaluation results are defined as:

Values:

```
v ::= i | C | D(v) | λx.a | raise
```

Evaluation contexts:

```
Γ ::= [] | Γ(a) | v(Γ) | D(Γ)
      | let x = Γ in a
      | match Γ with p → a2 | x → a3
      | try Γ with x → a
```

$$\begin{aligned}
(\lambda x.a)(v) &\Rightarrow a\{x \leftarrow v\} \\
\text{let } x = v \text{ in } a &\Rightarrow a\{x \leftarrow v\} \\
\text{match } v \text{ with } p \rightarrow a_2 \mid x \rightarrow a_3 &\Rightarrow \sigma(a_2) \text{ if } \sigma = M(v, p) \text{ is defined} \\
\text{match } v \text{ with } p \rightarrow a_2 \mid x \rightarrow a_3 &\Rightarrow a_3\{x \leftarrow v\} \text{ if } M(v, p) \text{ is undefined} \\
\text{try } v \text{ with } x \rightarrow a_2 &\Rightarrow v \\
(\text{raise } v)(a) &\Rightarrow \text{raise } v \\
(\lambda x.a)(\text{raise } v) &\Rightarrow \text{raise } v \\
D(\text{raise } v) &\Rightarrow \text{raise } v \\
\text{let } x = \text{raise } v \text{ in } a &\Rightarrow \text{raise } v \\
\text{match raise } v \text{ with } p \rightarrow a_2 \mid x \rightarrow a_3 &\Rightarrow \text{raise } v \\
\text{try raise } v \text{ with } x \rightarrow a_2 &\Rightarrow a_2\{x \leftarrow v\} \\
\Gamma[a] &\Rightarrow \Gamma[a'] \text{ if } a \Rightarrow a'
\end{aligned}$$

The pattern-matching function $M(v, p)$:

$$M(v, x) = \{x \leftarrow v\} \quad M(i, i) = id \quad M(C, C) = id \quad M(D(v), D(p)) = M(v, p)$$

Figure 1: Reduction rules

Evaluation results:

$$r ::= v \mid \text{raise } v$$

A result of v indicates normal termination with return value v ; a result of $\text{raise } v$ indicates an uncaught exception v .

3.2 The type algebra

The type system uses the following type algebra:

Type expressions:

$$\begin{aligned}
\tau &::= \alpha && \text{type variable} \\
&| \text{int}[\varphi] && \text{integer type} \\
&| \text{exn}[\varphi] && \text{exception type} \\
&| \tau_1 \xrightarrow{\varphi} \tau_2 && \text{function type}
\end{aligned}$$

Type schemes:

$$\sigma ::= \forall \alpha_i, \rho_j, \delta_k. \tau$$

Rows:

$$\begin{aligned}
\varphi &::= \rho && \text{row variable} \\
&| \top && \text{all possible elements} \\
&| \varepsilon; \varphi && \text{the element } \varepsilon \text{ plus whatever is in } \varphi
\end{aligned}$$

Row elements:

$$\begin{aligned}
\varepsilon &::= i : \pi && \text{integer constant} \\
&| C : \pi && \text{constant exception} \\
&| D(\tau) && \text{parameterized exception}
\end{aligned}$$

Presence annotations:

$$\begin{aligned}
\pi &::= \text{Pre} && \text{element is present} \\
&| \delta && \text{presence variable}
\end{aligned}$$

As in effect systems, our function types $\tau_1 \xrightarrow{\varphi} \tau_2$ are annotated by the latent effect φ of the function, that is, the set of exceptions that may be raised during application of the function. In addition, the base types $\text{exn}[\varphi]$ and $\text{int}[\varphi]$ are also annotated by sets of exceptions and integers respectively. Those sets refine the ML types exn and int by restricting the values that an expression of type $\text{exn}[\varphi]$ or $\text{int}[\varphi]$ can have.

Sets of exceptions or integers are represented by *rows* similar to those used for typing extensible records [31, 22, 24]. A row is either $\top$, meaning that all values of the type are possible (we do not have any more precise information), or a sequence of row elements $\varepsilon_1 \dots \varepsilon_n$ terminated by a *row variable* ρ . We impose the following equational theory on rows to express that the order of elements in a row does not matter (equation 1), and that $\top$ is absorbing (equation 2):

$$\varepsilon_1; \varepsilon_2; \varphi = \varepsilon_2; \varepsilon_1; \varphi \quad (1)$$

$$i : \text{Pre}; \top = \top \quad (2)$$

The absorption equation 2 applies only to integer row elements because we intend $\top$ to be used only in rows annotating the int type. (The kinding rules below enforce this invariant.) A $\top$ symbol is required for base types such as int , which have an infinite (or at least very large) signature. It is not required for datatypes such as exn , which have a finite signature: a row enumerating all possible constructors can be used instead (this is discussed in section 4.1.4 below). Moreover, combining $\top$ and rows containing parameterized constructors raises technical problems²; we prefer to avoid the difficulty by restricting $\top$ to rows containing only integer elements.

Rows and row variables support both polymorphism over sets and a form of set union in a unification framework. For instance, the two rows $\varepsilon_1; \rho_1$ and $\varepsilon_2; \rho_2$, which informally represent the sets $\{\varepsilon_1\}$ and $\{\varepsilon_2\}$ respectively, unify into the row $\varepsilon_1; \varepsilon_2; \rho$ representing the set $\{\varepsilon_1; \varepsilon_2\}$ via the substitution $\{\rho_1 \leftarrow (\varepsilon_2; \rho); \rho_2 \leftarrow (\varepsilon_1; \rho)\}$.

A row element ε is either an integer constant i , a constant exception constructor C , or a parameterized exception constructor $D(\tau)$ carrying the annotated type τ of its argument. To maintain crucial kinding invariants (see below),

²The obvious absorption equation $D(\tau); \top = \top = D(\beta); \top$ is unsound, as it allows deductions such as $D(\alpha); \top = \top = D(\beta); \top$, which lead to inconsistent typings. If ML had subtyping and a supertype $\top$ of all types, a correct equation would be $D(\tau); \top = \top$. This equation allows $\top$ to absorb any $D(\tau)$ (because $D(\tau); \top <: D(\tau); \top = \top$), but only allows expansion of $\top$ into $D(\tau)$; $\top$, meaning correctly that no information is available on the argument of D .

$\vdash \rho :: K(\rho)$	$\vdash \top :: \text{INT}(S)$	$\frac{i \notin S \quad \vdash \varphi :: \text{INT}(S \cup \{i\})}{\vdash (i; \pi; \varphi) :: \text{INT}(S)}$	$\frac{C \notin S \quad \vdash \varphi :: \text{EXN}(S \cup \{C\})}{\vdash (C; \pi; \varphi) :: \text{EXN}(S)}$
		$\frac{D \notin S \quad \vdash \varphi :: \text{EXN}(S \cup \{D\}) \quad \vdash \tau \text{ wf}}{\vdash (D(\tau); \varphi) :: \text{EXN}(S)}$	
$\vdash \alpha \text{ wf}$	$\frac{\vdash \varphi :: \text{INT}(\emptyset)}{\vdash \text{int}[\varphi] \text{ wf}}$	$\frac{\vdash \varphi :: \text{EXN}(\emptyset)}{\vdash \text{exn}[\varphi] \text{ wf}}$	$\frac{\vdash \tau_1 \text{ wf} \quad \vdash \varphi :: \text{EXN}(\emptyset) \quad \vdash \tau_2 \text{ wf}}{\vdash \tau_1 \xrightarrow{\varphi} \tau_2 \text{ wf}}$

Figure 2: Kinding rules

the constant row elements (i and C) also carry a *presence annotation*, written π . A presence annotation can be either **Pre**, meaning that the element is present in the set denoted by the row expression; or a *presence variable* δ meaning that the element is actually not present in the set denoted by the row expression, but may be considered as present in order to satisfy unification constraints.

Examples: The type $\text{int}[\top]$ denotes all integer values. The type of integer addition is

$$\forall \rho_1, \rho_2, \rho_3, \rho_4. \text{int}[\rho_1] \xrightarrow{\rho_2} \text{int}[\rho_3] \xrightarrow{\rho_4} \text{int}[\top]$$

(no effects, no information known on the return value).

The type scheme $\forall \rho. \text{int}[1 : \text{Pre}; 2 : \text{Pre}; \rho]$ stands for the set $\{1; 2\}$ and is the type of integer expressions that can only evaluate to 1 or to 2. A universally quantified row variable ρ that occurs only positively in a type should be read as denoting the empty set of elements, for the same reasons that $\forall \alpha. \alpha$ denotes an empty set of values.

The type scheme $\forall \rho, \delta. \text{int}[1 : \delta; 2 : \text{Pre}; \rho]$ stands for the set $\{2\}$. Although 1 is mentioned in the row, it should not be considered present in the set, since its presence annotation δ is universally quantified and occurs only positively.

The type scheme $\forall \rho, \rho'. \text{exn}[D(\text{int}[3 : \text{Pre}; 4 : \text{Pre}; \rho]); \rho']$ stands for the set of exceptions $\{D(3); D(4)\}$.

The **raise** predefined function has the following type scheme: $\forall \alpha, \rho. \text{exn}[\rho] \xrightarrow{\rho} \alpha$. It expresses that an application of **raise** never returns and raises exactly the exceptions that it receives as argument.

Kinding of rows: To simplify the formulation of the typing rules and to ensure the existence of principal unifiers and principal typings, we require the following four structural invariants on rows:

1. A given integer constant or exception constructor should occur at most once in a row (for instance, $(D(\tau); D(\tau'); \varphi)$ is not well-formed).
2. A row variable ρ is preceded by the same set of integer constants and exception constructors in all row expressions where it occurs (for instance, we cannot have both $(1 : \text{Pre}; \rho)$ and $(2 : \text{Pre}; \rho)$ in the same derivation).
3. A row φ annotating an integer type $\text{int}[\varphi]$ can only contain integer elements i .

4. A row φ annotating an exception type $\text{exn}[\varphi]$ or a function type $\tau_1 \xrightarrow{\varphi} \tau_2$ can only contain constant or parameterized constructors C, D and must not end with $\top$.

Invariants (1) and (2) are well known from earlier work on record types [24]. Invariants (3) and (4) are more unusual. They ensure a clear separation between annotations of **int** types (composed of integer elements and possibly $\top$) and annotations of the **exn** types (composed of constructors and no $\top$). Since $\top$ absorbs only integer elements (equation 2), we do not want it to occur in rows containing exception constructors C, D .

Following [24, 18], we use *kinds* to enforce the invariants above. Our kinds κ are composed of a tag (either **INT** or **EXN**) and a set of constants and constructors:

$$\kappa ::= \text{INT}(\{i_1, \dots, i_n\}) \mid \text{EXN}(\{C_1, \dots, C_p, D_1, \dots, D_q\})$$

The constants and constructors appearing in the set part of a kind are those constants and constructors that must not appear in rows of that kind (because they already appear in elements concatenated before those rows). We assume given a global mapping K assigning kinds to row variables, and such that for each κ there are infinitely many variables of that kind (i.e. $K^{-1}(\kappa)$ is infinite). The kinding rules are shown in figure 2. They define the two judgements $\vdash \varphi :: \kappa$ (row φ has kind κ) and $\vdash \tau \text{ wf}$ (type τ is well-formed).

3.3 The typing rules

Figure 3 shows the typing rules for our system. They define the judgement $E \vdash a : \tau / \varphi$, where E is the typing environment, a the term to type, τ the type of values that a may evaluate to, and φ the set of exceptions that may escape during the evaluation of a . All types appearing in the rules are assumed to be well-kinded. We assume that typing starts in the initial environment $E_0 = \{\text{raise} : \forall \alpha, \rho. \text{exn}[\rho] \xrightarrow{\rho} \alpha\}$.

The rules for variables and **let** bindings (rules 1 and 5) are standard, except that we generalize over all three kinds of type variables. For variables as well as other language constructs that never raise exceptions (rules 1, 2, 3, 7), the φ component of the result is unconstrained and can be chosen as needed to satisfy equality constraints in the remainder of the typing derivation.

The rules for function abstraction and application (rules 3 and 4) are the usual rules for effect systems. For abstraction, the effect of the function body becomes the latent effect of the function type. For application $a_1(a_2)$, we require that the same set φ of exceptions occurs as effect of a_1 , latent effect of the function denoted by a_1 , and effect of a_2 . This corresponds in our unification setting to taking the union of those three effects.

Typing of expressions:

$$\begin{array}{c}
\frac{\tau \leq E(x)}{E \vdash x : \tau} \quad (1) \qquad \frac{\vdash \varphi' :: \text{INT}(\{i\}) \quad \vdash \varphi :: \text{EXN}(\emptyset)}{E \vdash i : \text{int}[i : \text{Pre}; \varphi'] / \varphi} \quad (2) \qquad \frac{\vdash \tau_1 \text{ wf} \quad E \oplus \{x : \tau_1\} \vdash a : \tau_2 / \varphi' \quad \vdash \varphi :: \text{EXN}(\emptyset)}{E \vdash \lambda x. a : (\tau_1 \xrightarrow{\varphi'} \tau_2) / \varphi} \quad (3) \\
\\
\frac{E \vdash a_1 : (\tau' \xrightarrow{\varphi} \tau) / \varphi \quad E \vdash a_2 : \tau' / \varphi}{E \vdash a_1(a_2) : \tau / \varphi} \quad (4) \qquad \frac{E \vdash a_1 : \tau_1 / \varphi \quad E \oplus \{x : \text{Gen}(\tau_1, E, \varphi)\} \vdash a_2 : \tau / \varphi}{E \vdash \text{let } x = a_1 \text{ in } a_2 : \tau / \varphi} \quad (5) \\
\\
\frac{E \vdash a_1 : \tau_1 / \varphi \quad \vdash p : \tau_1 \Rightarrow E' \quad \vdash \tau_1 - p \rightsquigarrow \tau_2 \quad E \oplus E' \vdash a_2 : \tau / \varphi \quad E \oplus \{x : \tau_2\} \vdash a_3 : \tau / \varphi}{E \vdash \text{match } a_1 \text{ with } p \rightarrow a_2 \mid x \rightarrow a_3 : \tau / \varphi} \quad (6) \\
\\
\frac{\vdash \varphi' :: \text{EXN}(\{C\}) \quad \vdash \varphi :: \text{EXN}(\emptyset)}{E \vdash C : \text{exn}[C : \text{Pre}; \varphi'] / \varphi} \quad (7) \qquad \frac{\tau \leq \text{TypeArg}(D) \quad E \vdash a : \tau / \varphi \quad \vdash \varphi' :: \text{EXN}(\{D\}) \quad \vdash \varphi :: \text{EXN}(\emptyset)}{E \vdash D(a) : \text{exn}[D(\tau); \varphi'] / \varphi} \quad (8) \\
\\
\frac{E \vdash a_1 : \tau / \varphi' \quad E \oplus \{x : \text{exn}[\varphi']\} \vdash a_2 : \tau / \varphi}{E \vdash \text{try } a_1 \text{ with } x \rightarrow a_2 : \tau / \varphi} \quad (9)
\end{array}$$

Typing of patterns:

$$\begin{array}{c}
\vdash x : \tau \Rightarrow \{x : \tau\} \quad (10) \qquad \vdash i : \text{int}[i : \text{Pre}; \varphi] \Rightarrow \{\} \quad (11) \qquad \vdash C : \text{exn}[C : \text{Pre}; \varphi] \Rightarrow \{\} \quad (12) \\
\\
\frac{\tau \leq \text{TypeArg}(D) \quad \vdash p : \tau \Rightarrow E}{\vdash D(p) : \text{exn}[D(\tau); \varphi] \Rightarrow E} \quad (13)
\end{array}$$

Pattern subtraction:

$$\begin{array}{c}
\frac{\vdash \tau' \text{ wf}}{\vdash \tau - x \rightsquigarrow \tau'} \quad (14) \qquad \vdash \text{int}[i : \text{Pre}; \varphi] - i \rightsquigarrow \text{int}[i : \pi; \varphi] \quad (15) \qquad \vdash \text{exn}[C : \text{Pre}; \varphi] - C \rightsquigarrow \text{exn}[C : \pi; \varphi] \quad (16) \\
\\
\frac{\vdash \tau - p \rightsquigarrow \tau'}{\vdash \text{exn}[D(\tau); \varphi] - D(p) \rightsquigarrow \text{exn}[D(\tau'); \varphi]} \quad (17)
\end{array}$$

Instantiation and generalization:

$\tau' \leq \forall \alpha_i \rho_j \delta_k. \tau$ iff there exists τ_i, φ_j, π_k such that $\tau' = \tau \{ \alpha_i \leftarrow \tau_i, \rho_j \leftarrow \varphi_j, \delta_k \leftarrow \pi_k \}$ and $\vdash \tau_i \text{ wf}$ and $\vdash \varphi_j :: K(\rho_j)$.
 $\text{Gen}(\tau, E, \varphi)$ is $\forall \alpha_i \rho_j \delta_k. \tau$ where $\{\alpha_i, \rho_j, \delta_k\} = FV(\tau) \setminus (FV(E) \cup FV(\varphi))$.

Figure 3: The typing rules

For integer constants and exception constructors (rules 2, 7 and 8), we record the actual value of the expression in the approximation part of the type **int** or **exn**. For instance, the type of i must be of the form $\text{int}[i : \text{Pre}; \varphi]$, forcing $i : \text{Pre}$ to appear in the type of the expression. In rules 8 and 13, we write $\text{TypeArg}(D)$ for the type scheme of the argument of constructor D , e.g. $\text{TypeArg}(D) = \forall \rho. \text{int}[\rho]$ for an integer-valued exception D .

For an exception handler **try** a_1 **with** $x \rightarrow a_2$ (rule 9), the effect φ_1 of a_1 is injected in the type $\text{exn}[\varphi_1]$ assumed for x in a_2 .

The most interesting rule is rule 6 for the **match** construct. This rule is crucial to the precision of our exception analysis. When typing **match** a_1 **with** $p \rightarrow a_2 \mid x \rightarrow a_3$, we want to reflect the fact that the second alternative ($x \rightarrow a_3$) is selected only when the first alternative ($p \rightarrow a_2$) does not match the value of a_1 . In other terms, the type of values that can “flow” to x in the second alternative is not the type of the matched value a_1 , but the type of a_1 from which we have excluded all values matching the pattern p in the first alternative.

To achieve this, rules 14–17 define the pattern subtraction predicate $\vdash \tau - p \rightsquigarrow \tau'$, meaning that τ' is a correct

type for the values of type τ that do not match pattern p . For a variable pattern $p = x$ (rule 14), all values match the pattern, so it is correct to assume any τ' for the type of the non-matched values. For an integer pattern $p = i$ (rule 15), we force τ to unify with $\text{int}[i : \text{Pre}; \varphi]$, thus exposing in φ the set of all possible values of type τ that are different from i . Then, we take $\tau' = \text{int}[i : \pi; \varphi]$ for a suitable π . In particular, if that π is unconstrained in the remainder of the derivation, we can take π to be a fresh presence variable δ , thus reflecting that i is not among the possible values of type τ' . The rules for exception patterns (rules 16 and 17) are similar. If the exception has an argument, instead of changing a presence annotation, we recursively subtract in the type of the argument of the exception.

It is easy to see that the typing rules preserve the kinding invariants: if E is well-kinded and $E \vdash a : \tau / \varphi$, then $\vdash \tau \text{ wf}$ and $\vdash \varphi :: \text{EXN}(\emptyset)$.

3.4 Examples of typings

We now show some typings derivable in our system. These are principal typings identical to those found by our exception analyzer. Consider first a simple handler for one

exception C :

```
try raise(C)
with x → match x with C → 1 | y → raise y
```

The effect of `raise(C)` is $C:\text{Pre}; \rho$. Hence, the type of x is $\text{exn}[C:\text{Pre}; \rho]$. Subtracting the pattern C from this type, we obtain the type $\text{exn}[C:\delta; \rho]$ for y . Hence the effect of the whole `match` expression, and also of the whole `try` expression, is $C:\delta; \rho$. The type is $\text{int}[1:\text{Pre}; \rho']$. Since δ, ρ and ρ' are generalizable and occur only positively, we have established that no exception escapes the expression, and that it can only evaluate to the integer 1.

We now extend the previous example along the lines of the `failwith` example of section 2.3:

```
let failwith = λn. raise(D(n)) in
try failwith(42)
with x → match x with D(42) → 0 | y → raise y
```

We obtain the following intermediate typings:

$$\begin{aligned} \text{failwith} &: \forall \alpha, \rho_1, \rho_2. \text{int}[\rho_1] \xrightarrow{D(\text{int}[\rho_1]); \rho_2} \alpha \\ x &: \text{exn}[D(\text{int}[42:\text{Pre}; \rho_3]); \rho_4] \\ y &: \text{exn}[D(\text{int}[42:\delta; \rho_3]); \rho_4] \end{aligned}$$

Thus we conclude as before that no exception escapes this expression.

For a representative example of higher-order functions, consider function composition:

```
let compose = λf. λg. λx. f(g(x)) in
compose (λy. 0) (λz. raise(C)) 1
```

The type scheme for `compose` is $\forall \alpha, \beta, \gamma, \rho, \rho_1, \rho_2. (\alpha \xrightarrow{\rho} \beta) \xrightarrow{\rho_1} (\gamma \xrightarrow{\rho} \alpha) \xrightarrow{\rho_2} \gamma \xrightarrow{\rho} \beta$. The three occurrences of ρ express the union of the effects of f and g . The application of `compose` above has effect $C:\text{Pre}; \rho_3$.

Concerning exceptions as first-class values, the first example from section 2.2 becomes:

```
let test =
  λexn. try raise(exn)
    with x → match x with C → 1
              | y → raise(y)
in test(C)
```

The type scheme for `test` is $\forall \rho, \rho', \delta. \text{exn}[C:\text{Pre}; \rho] \xrightarrow{C:\delta; \rho} \text{int}[1:\text{Pre}; \rho']$, expressing that the function raises whatever exception it receives as argument, except C . The application `test(C)` has thus type $\text{int}[1:\text{Pre}; \rho_1]$ and effect $C:\delta_2; \rho_2$. Hence no exception escapes. The application `test C'` where C' is another exception distinct from C would have effect $C:\delta_3; C':\text{Pre}; \rho_3$, thus showing that C' may escape.

Finally, here is an (anecdotal) example that is ill-typed in ML, but well-typed in our type system due to the refined typing of pattern-matching:

```
match 1 with x → x | e → raise e
```

Since the first case of the matching is a catch-all, rule 6 lets us assign the type $\text{exn}[\rho']$ for a fresh ρ' to the variable e bound by the second case, even though the matched value is an integer. Hence the expression is well-typed, and moreover we obtain that it has type $\text{int}[1:\text{Pre}; \rho]$ and raises no exceptions (its effect is ρ' for any ρ').

3.5 Type soundness and correctness of the exception analysis

We now establish the correctness of our exception analysis: all uncaught exceptions are predicted by our effect system. This property is closely connected to the type soundness of our system.

Theorem 1 (Subject reduction) *Reduction preserves typing: if $E_0 \vdash a : \tau/\varphi$ and $a \Rightarrow a'$, then $E_0 \vdash a' : \tau/\varphi$*

The proof of subject reduction is mostly standard and follows [33] closely. Detailed proofs of the statements in this paper can be found in the technical report [14]. A key lemma is the following property of pattern subtraction:

Lemma 2 (Correctness of subtraction) *If $E_0 \vdash v : \tau/\varphi$ and $M(v, p)$ is undefined (v does not match pattern p) and $\vdash \tau - p \rightsquigarrow \tau'$, then $E_0 \vdash v : \tau'/\varphi$.*

The correctness of our exception analysis (all uncaught exceptions are detected) is a simple corollary of subject reduction:

Theorem 3 (Correctness of exception analysis)

Let a be a complete program. Assume $E_0 \vdash a : \tau/\varphi$ and $a \Rightarrow^ \text{raise } v$. Then, either $v = C$ and $\varphi = C:\text{Pre}; \varphi'$ for some C and φ' , or $v = D(v')$ and $\varphi = D(\tau'); \varphi'$ and $E_0 \vdash v' : \tau'/\varphi$ for some D, v', τ', φ' . In either case, the uncaught exception v is correctly predicted in the effect φ .*

Type soundness for our non-standard type system follows from the subject reduction property and the following lemma showing that well-typed expressions either reduce to a value or to an uncaught exception, or loop, but never get “stuck”.

Lemma 4 (Progress) *If $E_0 \vdash a : \tau/\varphi$, then either a is a value v , or a is an uncaught exception `raise v`, or there exists a' such that $a \Rightarrow a'$.*

3.6 Principal types and inference of types and exceptions

Just like the ML type system, our type system admits principal types, which can be computed by a simple extension of Milner’s algorithm W , thus implementing the exception analysis.

Theorem 5 (Principal unifiers) *The set of well-kinded types modulo equations (1) and (2) admits principal unifiers. More precisely, there exists an algorithm `mgu` that, for all system Q of well-kinded equations between types, either returns a substitution μ that is a principal solution of Q , or fails, meaning that Q has no solution. Moreover, the substitution μ preserves kinds in the following sense: for all α , $\vdash \mu(\alpha) \text{ wf}$ and for all ρ , $\vdash \mu(\rho) :: K(\rho)$.*

In the theorem above, systems of well-kinded equations are sets $Q = \{\tau_i = \tau'_i; \varphi_j = \varphi'_j; \pi_l = \pi'_l\}$ of equations between types, rows, row elements, and presence annotations such that for all i , $\vdash \tau_i \text{ wf}$ and $\vdash \tau'_i \text{ wf}$, and for all j , there exists a kind κ_j such that $\varphi_j :: \kappa_j$ and $\varphi'_j :: \kappa_j$.

The existence of principal unifiers follows from the fact that our equational theory is syntactic and regular [23]. The algorithm `mgu` is given in appendix A.

Theorem 6 (Principal types) *There exists a type inference algorithm W satisfying the following conditions:*

- (Correctness) If E is well-kinded and $(\tau, \varphi, \theta) = W(E, a)$ is defined, then $\theta(E) \vdash a : \tau / \varphi$.
- (Completeness) If E is well-kinded and there exists a kind-preserving substitution θ' and types τ', φ' such that $\theta'(E) \vdash a : \tau' / \varphi'$, then $(\tau, \varphi, \theta) = W(E, a)$ is defined and there exists a substitution ψ such that $\tau' = \psi(\tau)$ and $\varphi' = \psi(\varphi)$ and $\theta'(v) = \psi(\theta(v))$ for all type, row or presence variable v not used as a fresh variable by algorithm W .

The algorithm W is shown in appendix B.

4 Extension to the full Objective Caml language

In this section, we discuss the main issues in extending the analysis presented in section 3 to deal with the whole Objective Caml language [15].

4.1 Datatypes

User-defined datatypes (sum types) can be approximated in several different ways, depending on the desired trade-off between precision and speed of the analysis. We have considered the four approaches listed below (from most precise to least precise).

4.1.1 Full approximation of datatypes

The first approach applies to datatypes the same treatments as for exceptions: we annotate the type by a row φ approximating the possible values of that type, as constant constructors with presence annotations, and unary constructors with types of arguments. Consider the source-level datatype definition

type $\tilde{\alpha} \ t = C_1 \mid \dots \mid C_n \mid D_1 \text{ of } \mu_1 \mid \dots \mid D_m \text{ of } \mu_m$

where the μ_i are unannotated ML types. The propagation of approximations is captured by the following type schemes assigned to the constructors C_i and D_i :

$$\begin{aligned} C_i & : \forall \tilde{\alpha}, \rho'. \tilde{\alpha} \ t[C_i : \text{Pre}; \rho'] \\ D_i & : \forall \tilde{\alpha}, \vec{\rho}, \rho', \rho''. \tau_i \xrightarrow{\rho''} \tilde{\alpha} \ t[D_i(\tau_i); \rho'] \end{aligned}$$

where τ_i is the annotated type obtained from μ_i by adding distinct fresh row variables taken from $\vec{\rho}$ on every type constructor that carries a row annotation. For instance, given the declaration

type `intlist` = `Nil` | `Cons` of `int` * `intlist`

we assign `Nil` and `Cons` the type schemes

$$\begin{aligned} \text{Nil} & : \forall \rho. \text{intlist}[\text{Nil} : \text{Pre}; \rho] \\ \text{Cons} & : \forall \rho_1, \rho_2, \rho_3, \rho_4. \text{int}[\rho_1] \times \text{intlist}[\rho_2] \xrightarrow{\rho_3} \\ & \quad \text{intlist}[\text{Cons}(\text{int}[\rho_1] \times \text{intlist}[\rho_2]); \rho_4] \end{aligned}$$

Recursive datatypes such as `intlist` above naturally lead to recursive type expressions. Consider:

let `tail` =
 $\lambda x. \text{match } x \text{ with } \text{Cons}(\text{hd}, \text{tl}) \rightarrow \text{tl} \mid 1 \rightarrow 1$

During inference, `tl` and `1` receive types `intlist` $[\rho_1]$ and `intlist` $[\text{Cons}(\text{int}[\rho_2] \times \text{intlist}[\rho_1]); \rho_3]$ respectively. If only finite type expressions are allowed, those two types have no unifier and the program is rejected by the analysis. This is not acceptable, so we extend our type system with recursive (infinite, regular) type expressions. On the example above, we obtain $\mu\alpha. \text{intlist}[\text{Cons}(\text{int}[\rho_2] \times \alpha); \rho_3]$. The extension of our type system with recursive type expressions involves replacing term unification by graph unification in the type inference algorithm, but this causes no technical difficulties.

4.1.2 “Looped” approximations for recursive datatypes

The approximation scheme described above has the undesirable side-effect of recording in the type approximation the whole structure of a data structure given in extension. If the data types involved are recursive, we may end up with very large type approximations. Continuing the `intlist` example above, consider the expression $\ell_n = \text{Cons}(i_1, \text{Cons}(i_2, \dots, \text{Cons}(i_n, \text{Nil}) \dots))$. With the type of `Cons` given above, this expression is given an annotated type that is of depth n and records not only the fact that the list contains the integers $i_1 \dots i_n$ (an information that might be useful to analyze exceptions), but also the fact that the list has length n and that its first element is i_1 , the second i_2 , etc. The latter piece of information is, on practical examples, useless for analyzing exceptions. Moreover, such large approximations slow down the analysis.

A solution to this problem comes from the following remark: as soon as one of those big data structures given in extension is passed to a sufficiently complex function, its big, unfolded annotated type is going to be unified with a recursive type, forcing all the information in the big type to be folded back into a smaller recursive type. For instance, if we pass the list ℓ_n to the `tail` function shown above, the type of the list will be unified into

$$\begin{aligned} \tau_n & = \mu\alpha. \text{intlist}[\text{Cons}(\text{int}[i_1 : \text{Pre}; \dots; i_n : \text{Pre}; \rho_1] \times \alpha); \\ & \quad \text{Nil} : \text{Pre}; \rho_2] \end{aligned}$$

The idea, then, is to force this folding into a recursive type when the data structure is created, by giving recursive, pre-folded types to the data type constructors. This is easily achieved by unifying, in the type of the constructors, all occurrences of the recursively-defined type in argument position with the occurrence of the recursively-defined type in result position. For instance, in the case of the `Cons` constructor of type `intlist`, we start with the type

$$\begin{aligned} & \text{int}[\rho_1] \times \text{intlist}[\rho_2] \xrightarrow{\rho_3} \\ & \quad \text{intlist}[\text{Cons}(\text{int}[\rho_1] \times \text{intlist}[\rho_2]); \rho_4] \end{aligned}$$

as in the previous section, then unify the two underlined `intlist` types, then generalize the free variables, obtaining $\text{Cons} : \forall \rho_1, \rho_3, \rho_4. \text{int}[\rho_1] \times \tau \xrightarrow{\rho_3} \tau$ where τ is $\mu\alpha. \text{intlist}[\text{Cons}(\text{int}[\rho_1] \times \alpha); \rho_4]$. With this type for `Cons`, the list ℓ_n is given the reasonably compact type τ_n shown above.

This technique of “looping” the types of constructors also works for parameterized datatypes, as long as they are regular (the data type constructor is used with the same parameters in the argument types of the constructors). For non-regular datatypes such as

type `'a nreg` = `Leaf` of `'a` | `Node` of `'a list nreg`

the unification of the occurrences of t in the type of `Node` would render that constructor essentially useless. Fortunately, such non-regular data types are extremely rare in actual programs, so we can use full approximations for them without impacting performance.

4.1.3 Adding row parameters to datatypes

An alternative to annotating datatype constructors with rows is to add row parameters to the type constructor reflecting the row annotations on `exn`, `int` and function types contained within the datatype. This technique is used by Fährdrich *et al* [4]. For instance, the ML datatype definition

```
type t = A of int | B of exn | C of t | D of t
```

is turned into

```
type (ρ1, ρ2) t = A of int[ρ1] | B of exn[ρ2]
                | C of (ρ1, ρ2) t | D of (ρ1, ρ2) t
```

Two parameters ρ_1 and ρ_2 were added in order to reflect in the type t the possible values of types `int` and `exn` contained in that type. The type t itself is not annotated by a row recording which constructors A, B, C and D are present in values of that type. The net effect is to forget the structure of terms of type t , while correctly remembering the integers and exception values contained in the structure.

In practice, this solution appears to be slightly less precise and slightly more efficient than full approximations of non-recursive datatypes and looped approximations of recursive datatypes: type expressions are smaller, but in the case of t above, looped approximations can express the fact that a value of type t lack one of the constructors C or D, while this is not captured in the solution based on extra row parameters.

On datatypes that are not annotated by a row, we can no longer perform type subtraction during pattern-matching, since we have no approximation on the structure of values of that type. Hence, we simply consider that subtraction is the identity relation on those datatypes.

4.1.4 Datatypes without any approximations

For maximal speed and minimal precision, we can put no annotations at all on a datatype (neither a row approximation nor extra row parameters). This way, we forget not only the structure of values of that type, but also the exceptions, functions and base values contained in that type. Of course, this forces us to make very pessimistic assumptions on values extracted from a datatype without approximation. For instance, if we extract an integer by pattern-matching on such a datatype, we must give it type `int[⊤]` since it can really be any integer. This is reflected in the types of constructors by putting $\top$ annotations on all annotated types in the constructor argument. In the `intlist` example above, if we choose not to annotate `intlist` at all, we must give its constructors the following types:

```
Nil  : intlist
Cons : ∀ρ1. int[⊤] × intlist  $\xrightarrow{\rho_1}$  intlist
```

This approach assumes that we have $\top$ annotations for all types, while the type system from section 3 only has $\top$ for type `int`. However, we can allow $\top$ to annotate other base types such as `float` and `string`. For exceptions and other datatypes, since there are finitely many constructors,

we can use a (potentially recursive) row enumerating all constructors of the datatype instead of a built-in constant $\top$. In the case of lists, for instance, we can use the following “top row” $T_{\text{list}}(\alpha, \rho)$:

$$T_{\text{list}}(\alpha, \rho) = \mu\rho'. \text{Nil} : \text{Pre}; \text{Cons}(\alpha \times \alpha \text{ list}[\rho']); \rho$$

The annotated type $\tau \text{ list}[T_{\text{list}}(\tau, \rho)]$ correctly represents any list of elements of type τ .

The “no approximation” approach described in this paragraph may look excessively coarse, but is actually quite effective for datatypes that introduce no base types, exception types, nor function types. Prominent examples are the built-in ML types $\alpha \text{ list}$ and $\alpha \text{ array}$, where the α parameter already records all the information we need about list and array elements. For instance, a list of functions from integers to booleans has type $(\text{int}[\varphi_1] \xrightarrow{\varphi_2} \text{bool}[\varphi_3]) \text{ list}$, where φ_2 denotes the union of the effects of all functions present in the list. A function extracted from that list and applied has effect φ_2 , and not any exception as one might naively expect.

4.1.5 Choosing a datatype approximation

The choice between the four datatype analysis strategies described above can be done on a per-datatype basis, depending on the shape of the datatype definition. We have considered several simple heuristics to perform this choice. Our first prototype used full approximations for non-parameterized datatypes, and no approximations for parameterized datatypes. Our current prototype uses full approximations for non-recursive or non-regular datatypes, looped approximations for recursive datatypes, and no approximations for built-in types without interesting structure (arrays and floating-point numbers, for instance). Another factor that we plan to integrate in the heuristic is whether the datatype introduces any exception type, function type, or base type likely to be an exception argument (`string` and `int`, essentially); if not, we could favor the “no approximation” approach.

4.2 Tuples and records

Tuple types are not approximated specially: each component of the tuple type carries its own annotation. For instance, $\text{int}[1:\text{Pre}; 2:\text{Pre}; \rho] \times \text{int}[3:\text{Pre}; 4:\text{Pre}; \rho']$ stands for the set of four pairs $\{1;2\} \times \{3;4\}$. Pattern subtraction on tuple types is not pointwise subtraction, which would lead to incorrect results. Consider the type $\text{int}[1:\text{Pre}; \rho] \times \text{int}[2:\text{Pre}; 3:\text{Pre}; \rho']$. Subtracting pointwise the pattern $(1, 2)$ from this type would lead to type $\text{int}[1:\delta; \rho] \times \text{int}[2:\delta'; 3:\text{Pre}; \rho']$, which is incorrect since the value $(1, 3)$ is no longer in the set. Therefore, the current implementation performs no subtraction on tuples: we take $\vdash (\tau_1 \times \tau_2) - (p_1, p_2) \rightsquigarrow \tau_1 \times \tau_2$. For a more refined behavior, we could perform subtraction on one of the components if all other components are matched against catch-all patterns. For instance, we could take $\vdash (\tau_1 \times \tau_2) - (p_1, x_2) \rightsquigarrow \tau'_1 \times \tau_2$ if $\vdash \tau_1 - p_1 \rightsquigarrow \tau'_1$.

Unlike in SML, records in Caml are declared and matched by name. We analyze them like datatypes, by annotating the name of the record type by a row of a particular form. The row contains exactly one element recording the annotated type of every field. Pattern subtraction for record types behaves as in the case of tuples.

To summarize, the extended type algebra for datatypes, tuples and records is as follows:

Type expressions:

$$\begin{array}{l|l} \tau ::= \dots & \\ \hline \tilde{\tau} \text{ } \mathbf{t}[\varphi] & \text{approximated type constructor} \\ \hline \tilde{\tau} \text{ } \mathbf{t} & \text{non-approximated type constructor} \\ \hline \tau_1 \times \dots \times \tau_n & \text{tuple type} \end{array}$$

Row elements:

$$\varepsilon ::= \dots \mid \{lbl_1 : \tau_1; \dots; lbl_n : \tau_n\}$$

4.3 Mutable data structures

Mutable data structures (references, arrays, records with mutable fields) are trivially handled: it suffices to introduce the standard value restriction on **let**-generalization [34]. This results in a precise approximation of mutable data. For instance, an array of functions has type $(\tau_1 \xrightarrow{\varphi} \tau_2) \text{ array}$, where φ is the union of the latent effects of all functions stored in the array. In contrast, control-flow analyses would lose track of which functions are stored in the array, and thus also of the exceptions they may raise, unless supplemented by a region (aliasing) analysis.

4.4 Objects and classes

Because our system already uses recursive types, OCaml-style objects do not add significant complexity to our framework. We just need to extend the type algebra with object types, that is, polymorphic records of methods [21]. The type of each method is annotated by its latent effect. No extension to rows and row elements are needed. Since there are no object patterns in pattern-matching, pattern subtraction needs not be modified.

The OCaml class language interferes very little with the exception analysis. No significant modifications to the class type-checker are needed.

4.5 Modules and functors

Structures are assigned annotated signatures containing annotated types for the value components. Type abbreviations are currently handled by systematic expansion of their definitions³.

For matching a structure S against a signature Σ , there are two possible semantics. The opaque semantics says that the only things known about the restriction $(S : \Sigma)$ is what Σ publicizes. In our case, since user-provided signatures Σ contain no annotations, this amounts to forgetting the result of the analysis of S and assume $\top$ annotation on all value components of the restricted structure. The transparent semantics simply check that S matches Σ , but the restriction $(S : \Sigma)$ retains all information known about S . We implemented the transparent semantics, as the opaque semantics results in too much information loss. (The opaque semantics also precludes choosing datatype annotations based on the definition of the datatype.)

Similar problems arise with functors. All is known about the parameter of a functor is its syntactic signature. Hence, a naive analysis would assume $\top$ annotation on all components of the functor argument. For better precision, one could use techniques based on conjunctive types such as [25]. Other issues with functors are still unclear, such as

³This might cause performance problems in conjunction with OCaml objects, which relies intensively on type abbreviations to make type expressions more manageable [21]. If this turns out to be a problem, we could also handle abbreviations by adding extra row parameters to the type constructors, as described in [4] and in section 4.1.3.

the generativity of exception declaration in functor bodies, and the impact of the “exception polymorphism” offered by functors (a functor can take one or several exceptions as arguments, and have a different exception behavior depending on whether those arguments are instantiated later with identical or different exceptions).

For simplicity, we chose not to analyze functors when they are defined, but instead expand the functor body at each application and re-analyze the β -reduced body. Although this transformation increases the size of the analyzed source, the Caml programs we are interested in do not use functors intensively and this simple approach to analyzing functors works well in practice.

4.6 Separate analysis

Transparent signature matching precludes “true” separate analysis (where any module can be analyzed separately knowing only the syntactic signatures of the modules it imports). We can still do “bottom-up” separate analysis, however: a module can be analyzed separately provided the implementations of its imports have been analyzed already, and their annotated signatures inferred.

Since annotated signature for a module may contain free row variables (e.g. if the module defines mutable structures), separately analyzing several clients of that module may result in independent instantiations of those free variables. Those instantiations are recorded in the result of the analysis of each module, and reconciled in a final “linking” pass before displaying the results of the analysis.

4.7 Polymorphic recursion

Polymorphic recursion as introduced by Mycroft [17] is not needed to type-check the source OCaml language, but is desirable to enhance the precision of our exception analyzer. With ML-style monomorphic recursion, we obtain false positives on functions that recursively call themselves inside a **try...with**. Consider:

```
let rec f =
  λx. try f(x) with C → () | y → raise y
```

The latent effect inferred for **f** is $C; \rho$ because the effect of **f(x)** is unified with the type of the pattern **C** at a time where the type of **f** is not yet generalized. With polymorphic recursion, we can assign **f** the type scheme $\forall \alpha, \rho. \alpha \xrightarrow{\rho} \text{unit}$ both outside and inside the recursion; it is a fresh instance of that type scheme that gets unified with the type of **C**, thus not polluting the type scheme of **f**.

Although type inference with polymorphic recursion is undecidable [13], there exists incomplete inference algorithms that work very well in practice. We experimented with Henglein’s algorithm [11] and with a home-grown algorithm based on restricted fixpoint iteration and obtained good results.

5 Experimental results

In this section, we present some experimental results obtained with our implementation. Currently, our analyzer implements all extensions described in section 4 except objects⁴. The analyzer is compiled with the OCaml 2.00 native-code compiler and runs on a Pentium II 333 Mhz workstation under Linux.

⁴The analysis of objects and classes was prototyped separately and remains to be merged in our main implementation.

Test program	Size (lines)	Analysis time	Analysis speed (lines per sec.)	OCaml type- checking time
1. Huffman compression	233	0.07/0.08 s	3300/2900 l/s	0.08 s
2. Knuth-Bendix	441	0.14/0.16 s	3200/2800 l/s	0.14 s
3. Docteur (Eliza clone)	556	0.81/0.83 s	680/670 l/s	0.10 s
4. Lexer generator	1169	0.27/0.32 s	4300/3700 l/s	0.20 s
5. Nucleic	2919	1.90/1.88 s	1530/1550 l/s	0.62 s
6. OCaml standard library	3082	2.52/2.52 s	1200/1200 l/s	1.89 s
7. Analyzer of .h files	3088	0.54/0.58 s	5700/5300 l/s	0.27 s
8. Our exception analyzer	12235	10.3/16.1 s	1200/760 l/s	3.86 s
9. The OCaml bytecode compiler	17439	12.6/22.9 s	1400/760 l/s	4.00 s

Figure 4: Experimental results (without polymorphic recursion/with polymorphic recursion)

Analysis speed: Figure 4 gives timings for the analysis of various small to medium-sized OCaml programs. We give timings both without and with polymorphic recursion. For comparison, we also give the time OCaml takes to parse and type-check those programs. (The timings given include parsing and pre-processing as well as analysis time.)

The overall performances are quite good, in the order of 1000–2000 lines of source per second. Programs that contain large data structures given in extension (Nucleic, Docteur) take longer to analyze due to the large size of the rows annotating the types of those data structures. On average, the exception analysis takes twice as much time as OCaml type inference; the ratio ranges between 1 (on simple programs) and 8 (on Docteur, because of the large constant data structures). Polymorphic recursion slows down the analysis somewhat on the largest benchmarks, but the slowdown remains acceptable compared with the increase in precision.

Precision of the analysis: We have manually inspected the output of the analyzer on our benchmark programs. Programs 1, 3, 4, 5 and 7 have a relatively simple exception behavior, and our analysis reports exact results for those programs: there are no false positives except run-time errors such as “division by zero” or “array index out of bounds”, which require extra analyses (or even general program proof) to show that they cannot occur.

For Knuth-Bendix, which has a quite complicated exception structure, 8 exceptions (**Failure** with 8 different string arguments) appearing in the source are correctly reported as non-escaping; 7 exceptions (one **Invalid_argument** and 6 **Failure**) are reported as potentially escaping, and can actually occur in some circumstances. Without polymorphic recursion, the analysis reports two false positives (one **Not_found** and one **Failure**), which correspond to recursive functions containing **try ... with** around recursive calls. Adding polymorphic recursion as discussed in section 4.7 removes one of those false positives. The other one is still there, because our incomplete inference algorithm for polymorphic recursion fails to give a type polymorphic enough to one of the functions. A more precise algorithm such as Henglein’s [11] would probably eliminate the other false positive as well.

The larger examples 8 and 9 exhibit another source of false positives: mutable data structures (references and arrays) containing functions. As mentioned in section 4.3, the row variables appearing in approximations of mutable data structures are not generalized, hence “collect” all exceptions at their use sites. For instance:

```
let r = ref(λx. ...) in
```

```
let f = λy. if cond then !r y
           else raise C
in !r 0
```

r has type $\text{int} \xrightarrow{\rho} \text{int}$ where ρ is not generalized. When typing **f**, the effect of **raise C** is unified with that of **!r y**, hence ρ becomes $C : \text{Pre}; \rho'$ and the application **!r 0** appears to raise **C**.

6 Related work

6.1 Exception analyses for ML

Several exception analyses for ML are described in the literature. Guzmán and Suárez [8] develop a simple type and effect system to keep track of escaping exceptions. Their system does not handle exceptions as first-class values, nor exceptions carrying arguments. The first exception analysis proposed by Yi [36] is based on general abstract interpretation techniques, and runs too slowly to be usable in practice. Later, Yi and Ryu [35] developed a more efficient analysis roughly equivalent to a conventional control-flow analysis to approximate the call graph and the values of exceptions, followed by a data-flow analysis to estimate uncaught exceptions.

Fähndrich and Aiken [3, 4] have applied their BANE toolkit for constraint-based program analyses to the problem of analyzing uncaught exceptions in SML. Their system uses a combination of inclusion constraints (as in control-flow analyses) to approximate the control flow, and equality constraints (unification) between annotated types to keep track of exception values.

To compare performances between [35], [3] and our analyzer, we used two of our benchmarks for which we have a faithful SML translation: Knuth-Bendix and Nucleic. The times reported below are of the form t_1/t_2 , where t_1 is the time spent in exception analysis only, and t_2 is the total program analysis time, including parsing and type-checking in addition to exception analysis.

Test program	Yi-Ryu	BANE	us
Knuth-Bendix	1.2/1.5 s	1.6/2.2 s	0.06/0.14 s
Nucleic	3.8/7.8 s	3.3/7.6 s	1.45/1.86 s

From these figures, our exception analysis seems notably faster. However, there are many external factors that influence the total running times of the analyses (such as the Yi-Ryu and BANE analyses being compiled by SML/NJ while ours is compiled by Objective Caml), so the figures above are not conclusive.

The main difference between the analyses of [35, 3] and ours is the approximation of arguments carried by exceptions: they approximate only exception and function values carried by exceptions, but our analysis is the only one that also approximates exception arguments that are strings, integers, or datatypes. As explained in section 2.3, approximating all arguments of exceptions is crucial to obtain precise analysis of many real applications.

In theory, our unification-based analysis should be less precise than analyses based on inclusion constraints such as [35, 3]: the bidirectional propagation of information performed by unification causes exception effects to “leak” in types where those exceptions cannot actually occur. It is easy to construct artificial examples of such leaks, e.g. by replacing `let`-bound identifiers by λ -bound identifiers. However, those examples do not seem to occur in actual programs. The only leaks we observed in actual programs were related either to deficiencies of our incomplete algorithm for typing polymorphic recursion, or to functions contained inside mutable data structures. On those two cases, [3] obtains more precise results than our analysis.

6.2 Other related work

Our use of rows with row variables and presence annotations to approximate values of base types and sum types is essentially identical to Rémy’s typing of extensible variants [22]. Another application of Rémy’s encoding is the soft typing system for Scheme of [32].

There is a natural connection between exception analysis and type inference for extensible variants: using the well-known functional encoding of exceptions (where each subexpression is transformed to return a value of a variant type, either an exception tag or `NormalResult(v)` where v is the value of the subexpression), estimating uncaught exceptions is equivalent to inferring precise variant types. Pottier [20] outlines an exception analysis thus derived from a type inferencer for ML with subtyping.

Refinement types [7] also introduce annotations on types to characterize subsets of ML’s data types. Our approach is less ambitious than refinement types, in that it does not try to capture “deep” structural invariants of recursive data structures; on the other hand, type inference is much easier.

The principles of effect systems were studied extensively in the early ’90s [16, 28], but few practical applications have been developed since. An impressive application is the region analysis of Tofte *et al.* [30, 29]. Like ours, its precision is improved by typing recursion polymorphically.

Several program analyses based on unification and running in quasi-linear time have been proposed as faster alternatives to more conventional dataflow analyses. Two well-known examples are Henglein’s tagging analysis [10] and Steensgaard’s aliasing analysis [27]. Baker [1] suggests other examples of unification-based analyses.

7 Conclusions and future work

It is often said that unification-based program analyses are faster, but less precise than more general constraint-based analyses such as CFA or SBA. For exception analysis, our experience indicates that a combination of unification, `let`-polymorphism, and polymorphic recursion is in practice almost as precise as analyses based on inclusion constraints. (The only case where our analysis is noticeably less precise than inclusion constraints is when references to functions are used intensively.) The running times of our algorithm seem

excellent (although its theoretical complexity is at least as high as that of ML type inference). In turn, this good efficiency of our analysis allows us to keep more information on exception arguments than the other exception analyses, increasing greatly the precision of the analysis on certain ML programs. Thus, we see an interesting case of “less is more”, where an *a priori* imprecise technology (unification) allows to improve eventually the precision of the analysis.

Some engineering issues remain to be solved before our analysis can be applied to large ML applications. The main practical issue is displaying the results of the analysis in a readable way. The volume of information contained in annotated type expressions can be overwhelming. The programmer should be able to select different levels of display abstracting some of that information.

Acknowledgements

The inference algorithm for polymorphic recursion used in our implementation was designed in collaboration with Pierre Weis. We thank François Pottier and Didier Rémy for interesting discussions.

References

- [1] H. G. Baker. Unify and conquer (garbage, updating, aliasing, ...) in functional languages. In *Lisp and Functional Programming 1990*. ACM Press, 1990.
- [2] M. Fähndrich and A. Aiken. Making set-constraint based program analyses scale. Technical Report 96-917, University of California at Berkeley, Computer Science Division, 1996.
- [3] M. Fähndrich and A. Aiken. Program analysis using mixed term and set constraints. In *Static Analysis Symposium ’97*, number 1302 in LNCS, pages 114–126. Springer-Verlag, 1997.
- [4] M. Fähndrich, J. S. Foster, A. Aiken, and J. Cu. Tracking down exceptions in Standard ML programs. Technical report, University of California at Berkeley, Computer Science Division, 1998.
- [5] M. Fähndrich, J. S. Foster, Z. Su, and A. Aiken. Partial online cycle elimination in inclusion constraint graphs. In *Prog. Lang. Design and Impl. 1998*, pages 85–96. ACM Press, 1998.
- [6] C. Flanagan and M. Felleisen. Componential set-based analysis. In *Prog. Lang. Design and Impl. 1997*. ACM Press, 1997.
- [7] T. Freeman and F. Pfenning. Refinement types for ML. In *Prog. Lang. Design and Impl. 1991*, pages 268–277. ACM Press, 1991.
- [8] J. C. Guzmán and A. Suárez. A type system for exceptions. In *Proc. 1994 workshop on ML and its applications*, pages 127–135. Research report 2265, INRIA, 1994.
- [9] N. Heintze. Set-based analysis of ML programs. In *Lisp and Functional Programming ’94*, pages 306–317. ACM Press, 1994.
- [10] F. Henglein. Global tagging optimization by type inference. In *Lisp and Functional Programming 1992*. ACM Press, 1992.

- [11] F. Henglein. Type inference with polymorphic recursion. *ACM Trans. Prog. Lang. Syst.*, 15(2):253–289, 1993.
- [12] S. Jagannathan and A. Wright. Polymorphic splitting: An effective polyvariant flow analysis. *ACM Trans. Prog. Lang. Syst.*, 20(1):166–207, 1998.
- [13] A. J. Kfoury, J. Tiuryn, and P. Urzyczyn. Type reconstruction in the presence of polymorphic recursion. *ACM Trans. Prog. Lang. Syst.*, 15(2):290–311, 1993.
- [14] X. Leroy and F. Pessaux. Type-based analysis of uncaught exceptions. Research report 3541, INRIA, Nov. 1998. Extended version of this paper.
- [15] X. Leroy, J. Vouillon, D. Doligez, et al. The Objective Caml system. Software and documentation available on the Web, <http://caml.inria.fr/ocaml/>, 1996.
- [16] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *15th symp. Principles of Progr. Lang.*, pages 47–57. ACM Press, 1988.
- [17] A. Mycroft. Polymorphic type schemes and recursive definitions. In *International Symposium on Programming*, number 167 in LNCS, pages 217–228. Springer-Verlag, 1984.
- [18] A. Ohori. A polymorphic record calculus and its compilation. *ACM Trans. Prog. Lang. Syst.*, 17(6):844–895, 1995.
- [19] F. Pottier. A framework for type inference with subtyping. In *Int. Conf. on Functional Progr. 1998*, pages 228–238. ACM Press, 1996.
- [20] F. Pottier. Type inference in the presence of subtyping: from theory to practice. Research report 3483, INRIA, Sept. 1998.
- [21] D. Rémy and J. Vouillon. Objective ML: A simple object-oriented extension of ML. In *24th symp. Principles of Progr. Lang.*, pages 40–53. ACM Press, 1997.
- [22] D. Rémy. Records and variants as a natural extension of ML. In *16th symp. Principles of Progr. Lang.*, pages 77–88. ACM Press, 1989.
- [23] D. Rémy. Syntactic theories and the algebra of record terms. Research report 1869, INRIA, 1993.
- [24] D. Rémy. Type inference for records in a natural extension of ML. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming*. MIT Press, 1993.
- [25] Z. Shao and A. Appel. Smartest recompilation. In *20th symp. Principles of Progr. Lang.*, pages 439–450. ACM Press, 1993.
- [26] O. Shivers. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis CMU-CS-91-145, Carnegie Mellon University, May 1991.
- [27] B. Steensgaard. Points-to analysis in almost linear time. In *23rd symp. Principles of Progr. Lang.*, pages 32–41. ACM Press, 1996.
- [28] J.-P. Talpin and P. Jouvelot. The type and effect discipline. *Inf. and Comp.*, 111(2):245–296, 1994.
- [29] M. Tofte and L. Birkedal. A region inference algorithm. *ACM Trans. Prog. Lang. Syst.*, 1998. To appear.
- [30] M. Tofte and J.-P. Talpin. Region-based memory management. *Inf. and Comp.*, 132(2):109–176, 1997.
- [31] M. Wand. Complete type inference for simple objects. In *Logic in Computer Science 1987*, pages 37–44. IEEE Computer Society Press, 1987.
- [32] A. K. Wright and R. Cartwright. A practical soft type system for Scheme. *ACM Trans. Prog. Lang. Syst.*, 19(1):87–152, 1997.
- [33] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Inf. and Comp.*, 115(1):38–94, 1994.
- [34] A. K. Wright. Simple imperative polymorphism. *Lisp and Symbolic Computation*, 8(4):343–356, 1995.
- [35] K. Yi and S. Ryu. Towards a cost-effective estimation of uncaught exceptions in SML programs. In *Static Analysis Symposium '97*, number 1302 in LNCS, pages 98–113. Springer-Verlag, 1997.
- [36] K. Yi. An abstract interpretation for estimating uncaught exceptions in Standard ML programs. *Sci. Comput. Programming*, 31(1):147–173, 1998.

A The unification algorithm

In this appendix, we give the unification algorithm for our type algebra modulo the two equations (1) and (2). We define the head constructor $H(\varepsilon)$ of a row element ε as follows:

$$H(i:\pi) = i \quad H(C:\pi) = C \quad H(D(\tau)) = D$$

The algorithm handles the left commutativity axiom (equation (1)) like in [24].

$$\text{mgu}(\emptyset) = id$$

Unification between types:

$$\begin{aligned} \text{mgu}(\{\alpha = \alpha\} \cup Q) &= \text{mgu}(Q) \\ \text{mgu}(\{\alpha = \tau\} \cup Q) &= \\ &\quad \text{mgu}(Q\{\alpha \leftarrow \tau\}) \circ \{\alpha \leftarrow \tau\} \\ &\quad \text{if } \alpha \notin FV(\tau) \\ \text{mgu}(\{\tau = \alpha\} \cup Q) &= \\ &\quad \text{mgu}(Q\{\alpha \leftarrow \tau\}) \circ \{\alpha \leftarrow \tau\} \\ &\quad \text{if } \alpha \notin FV(\tau) \\ \text{mgu}(\{\text{int}[\varphi_1] = \text{int}[\varphi_2]\} \cup Q) &= \\ &\quad \text{mgu}(\{\varphi_1 = \varphi_2\} \cup Q) \\ \text{mgu}(\{\text{exn}[\varphi_1] = \text{exn}[\varphi_2]\} \cup Q) &= \\ &\quad \text{mgu}(\{\varphi_1 = \varphi_2\} \cup Q) \\ \text{mgu}(\{\tau_1 \xrightarrow{\varphi_1} \tau'_1 = \tau_2 \xrightarrow{\varphi_2} \tau'_2\} \cup Q) &= \\ &\quad \text{mgu}(\{\tau_1 = \tau_2; \varphi_1 = \varphi_2; \tau'_1 = \tau'_2\} \cup Q) \end{aligned}$$

Unification between rows:

$$\begin{aligned} \text{mgu}(\{\rho = \rho\} \cup Q) &= \text{mgu}(Q) \\ \text{mgu}(\{\rho = \varphi\} \cup Q) &= \\ &\quad \text{mgu}(Q\{\rho \leftarrow \varphi\}) \circ \{\rho \leftarrow \varphi\} \\ &\quad \text{if } \rho \notin FV(\varphi) \\ \text{mgu}(\{\varphi = \rho\} \cup Q) &= \\ &\quad \text{mgu}(Q\{\rho \leftarrow \varphi\}) \circ \{\rho \leftarrow \varphi\} \end{aligned}$$

if $\rho \notin FV(\varphi)$
 $\text{mgu}(\{\top = \top\} \cup Q) = \text{mgu}(Q)$
 $\text{mgu}(\{(i : \pi; \varphi) = \top\} \cup Q) =$
 $\text{mgu}(\{\pi = \text{Pre}; \varphi = \top\} \cup Q)$
 $\text{mgu}(\{\top = (i : \pi; \varphi)\} \cup Q) =$
 $\text{mgu}(\{\pi = \text{Pre}; \varphi = \top\} \cup Q)$
 $\text{mgu}(\{(\varepsilon_1; \varphi_1) = (\varepsilon_2; \varphi_2)\} \cup Q) =$
 $\text{mgu}(\{\varepsilon_1 = \varepsilon_2\} \cup Q)$
 if $H(\varepsilon_1) = H(\varepsilon_2)$
 $\text{mgu}(\{(\varepsilon_1; \varphi_1) = (\varepsilon_2; \varphi_2)\} \cup Q) =$
 $\text{mgu}(\{\varphi_1 = (\varepsilon_2; \rho); \varphi_2 = (\varepsilon_1; \rho)\} \cup Q)$
 if $H(\varepsilon_1) \neq H(\varepsilon_2)$
 and ρ is not free in the l.h.s.
 and has kind $t(S \cup \{H(\varepsilon_1), H(\varepsilon_2)\})$
 where $t(S)$ is the kind of $\varepsilon_1; \varphi_1$ and $\varepsilon_2; \varphi_2$

Unification between row elements:

$\text{mgu}(\{(i : \pi_1) = (i : \pi_2)\} \cup Q) =$
 $\text{mgu}(\{\pi_1 = \pi_2\} \cup Q)$
 $\text{mgu}(\{(C : \pi_1) = (C : \pi_2)\} \cup Q) =$
 $\text{mgu}(\{\pi_1 = \pi_2\} \cup Q)$
 $\text{mgu}(\{D(\tau_1) = D(\tau_2)\} \cup Q) =$
 $\text{mgu}(\{\tau_1 = \tau_2\} \cup Q)$

Unification between presence annotations:

$\text{mgu}(\{\delta = \delta\} \cup Q) = \text{mgu}(Q)$
 $\text{mgu}(\{\delta = \pi\} \cup Q) =$
 $\text{mgu}(Q\{\delta \leftarrow \pi\}) \circ \{\delta \leftarrow \pi\}$ if $\pi \neq \delta$
 $\text{mgu}(\{\pi = \delta\} \cup Q) =$
 $\text{mgu}(Q\{\delta \leftarrow \pi\}) \circ \{\delta \leftarrow \pi\}$ if $\pi \neq \delta$
 $\text{mgu}(\{\text{Pre} = \text{Pre}\} \cup Q) = \text{mgu}(Q)$

If none of the cases above is applicable, $\text{mgu}(Q)$ is undefined.

B The type inference algorithm

The result of the algorithm $W(E, a)$ is the triple (τ, φ, θ) defined by induction on a as follows:

If a is x (with $x \in \text{Dom}(E)$):
 let ρ be a fresh variable of kind $\text{EXN}(\emptyset)$
 take $\tau = \text{Inst}(E(x))$ and $\varphi = \rho$ and $\theta = id$.

If a is i :
 let ρ be a fresh variable of kind $\text{INT}(\{i\})$
 let ρ' be a fresh variable of kind $\text{EXN}(\emptyset)$
 take $\tau = \text{int}[i : \text{Pre}; \rho]$ and $\varphi = \rho'$ and $\theta = id$.

If a is $\lambda x. a_1$:
 let α be a fresh variable
 let $(\tau_1, \varphi_1, \theta_1) = W(E \oplus \{x : \alpha\}, a_1)$
 let ρ be a fresh variable of kind $\text{EXN}(\emptyset)$
 take $\tau = \theta_1(\alpha) \xrightarrow{\varphi} \tau_1$ and $\varphi = \rho$ and $\theta = \theta_1$.

If a is $a_1(a_2)$:
 let $(\tau_1, \varphi_1, \theta_1) = W(E, a_1)$
 let $(\tau_2, \varphi_2, \theta_2) = W(\theta_1(E), a_2)$
 let α be a fresh variable
 let $\mu = \text{mgu}\{\theta_2(\tau_1) = \tau_2 \xrightarrow{\varphi_2} \alpha, \theta_2(\varphi_1) = \varphi_2\}$
 take $\tau = \mu(\alpha)$ and $\varphi = \mu(\varphi_2)$ and $\theta = \mu \circ \theta_2 \circ \theta_1$.

If a is let $x = a_1$ in a_2 :
 let $(\tau_1, \varphi_1, \theta_1) = W(E, a_1)$
 let $(\tau_2, \varphi_2, \theta_2) = W(\theta_1(E) \oplus \{x : \text{Gen}(\tau_1, \theta_1(E), \varphi_1)\}, a_2)$
 let $\mu = \text{mgu}\{\theta_2(\varphi_1) = \varphi_2\}$
 take $\tau = \mu(\tau_2)$ and $\varphi = \mu(\varphi_2)$ and $\theta = \mu \circ \theta_2 \circ \theta_1$.

If a is match a_1 with $p \rightarrow a_2 \mid x \rightarrow a_3$:
 let $(\tau_1, \varphi_1, \theta_1) = W(E, a_1)$
 let $(E', \tau', \psi) = \text{Patsubtr}(p, \tau_1)$
 let $(\tau_2, \varphi_2, \theta_2) = W(\psi(\theta_1(E)) \oplus E', a_2)$
 let $(\tau_3, \varphi_3, \theta_3) = W(\theta_2(\psi(\theta_1(E))) \oplus \{x : \theta_2(\tau')\}, a_3)$
 let $\mu = \text{mgu}\{\theta_3(\tau_2) = \tau_3, \theta_3(\varphi_2) = \varphi_3, \theta_3(\theta_2(\psi(\varphi_1))) = \varphi_3\}$
 take $\tau = \mu(\tau_3)$ and $\varphi = \mu(\varphi_3)$ and $\theta = \mu \circ \theta_3 \circ \theta_2 \circ \psi \circ \theta_1$.

If a is C :
 let ρ be a fresh variable of kind $\text{EXN}(\{C\})$
 let ρ' be a fresh variable of kind $\text{EXN}(\emptyset)$
 take $\tau = \text{exn}[C : \text{Pre}; \rho]$ and $\varphi = \rho'$ and $\theta = id$.

If a is $D(a_1)$:
 let $(\tau_1, \varphi_1, \theta_1) = W(E, a_1)$
 let $\tau_2 = \text{Inst}(\text{TypeArg}(D))$
 let $\mu = \text{mgu}\{\tau_2 = \tau_1\}$
 let ρ be a fresh variable of kind $\text{EXN}(\{D\})$
 let ρ' be a fresh variable of kind $\text{EXN}(\emptyset)$
 take $\tau = \text{exn}[D(\mu(\tau_1)); \rho]$ and $\varphi = \rho'$ and $\theta = \mu \circ \theta_1$.

If a is try a_1 with $x \rightarrow a_2$:
 let $(\tau_1, \varphi_1, \theta_1) = W(E, a_1)$
 let $(\tau_2, \varphi_2, \theta_2) = W(\theta_1(E) \oplus \{x : \text{exn}[\varphi_1]\}, a_2)$
 let $\mu = \text{mgu}\{\theta_2(\tau_1) = \tau_2\}$
 take $\tau = \mu(\tau_2)$ and $\varphi = \mu(\varphi_2)$ and $\theta = \mu \circ \theta_2 \circ \theta_1$.

The auxiliary function $\text{Inst}(\sigma)$ (trivial instance):
 $\text{Inst}(\forall \alpha_i, \rho_j, \delta_k. \tau)$ is $\tau\{\alpha_i \leftarrow \alpha'_i, \rho_j \leftarrow \rho'_j, \delta_k \leftarrow \delta'_k\}$ where $\alpha'_i, \rho'_j, \delta'_k$ are fresh variables such that ρ'_j and ρ_j have the same kind for all j .

The auxiliary function Patsubtr (typing of patterns and pattern subtraction): $\text{Patsubtr}(p, \tau)$ is the triple (E, τ', θ) defined by induction on p as follows:

If p is x :
 let α be a fresh variable
 take $E = \{x : \tau\}$ and $\tau' = \alpha$ and $\theta = id$.

If p is i :
 let ρ be a fresh variable of kind $\text{INT}(\{i\})$
 let $\mu = \text{mgu}\{\tau = \text{int}[i : \text{Pre}; \rho]\}$
 let δ be a fresh presence variable
 take $E = \emptyset$ and $\tau' = \text{int}[i : \delta; \mu(\rho)]$ and $\theta = \mu$.

If p is C :
 let ρ be a fresh variable of kind $\text{EXN}(\{C\})$
 let $\mu = \text{mgu}\{\tau = \text{exn}[C : \text{Pre}; \rho]\}$
 let δ be a fresh presence variable
 take $E = \emptyset$ and $\tau' = \text{exn}[C : \delta; \mu(\rho)]$ and $\theta = \mu$.

If p is $D(p_1)$:
 let $\tau_1 = \text{Inst}(\text{TypeArg}(D))$
 let $(E_1, \tau'_1, \theta_1) = \text{Patsubtr}(p_1, \tau_1)$
 let ρ be a fresh variable of kind $\text{EXN}(\{D\})$
 let $\mu = \text{mgu}\{\tau = \text{exn}[D(\theta_1(\tau_1)); \rho]\}$
 take $E = \mu(E_1)$ and $\tau' = \text{exn}[D(\mu(\tau'_1)); \mu(\rho)]$ and $\theta = \mu \circ \theta_1$.