

Backwards-Compatible Row-Based Exceptions in ML

SIMCHA VAN COLLEM, Radboud University Nijmegen, The Netherlands

PAULO EMÍLIO DE VILHENA, Imperial College London, United Kingdom

ROBBERT KREBBERS, Radboud University Nijmegen, The Netherlands

We introduce a type system that provides strong types for exception tracking in ML-style languages. Our type system employs a rich notion of row polymorphism and subtyping to ensure backwards compatibility, making sure that code without exception tracking continues to work and can be generalized gracefully to support exception tracking. We study the safety and abstraction guarantees of our type system, in particular the role of local exceptions for data abstraction. We formulate these claims using binary logical relations in a novel relational separation logic for exceptions, an independent contribution of this paper. We support a realistic subset of features from ML-style languages, such as extensible variant types, local exceptions, and concurrency. We exercise our type system and logic on a number of challenging examples taken from the OCaml standard library, from one of Jane Street’s OCaml libraries, and from Filinski’s PhD thesis. All our results are mechanized in the Rocq prover using Iris.

CCS Concepts: • **Theory of computation** → **Concurrency**; **Separation logic**; **Control primitives**.

Additional Key Words and Phrases: Exceptions, Types, Concurrency, Relational Separation Logic, Iris, Rocq

ACM Reference Format:

Simcha van Collem, Paulo Emilio de Vilhena, and Robbert Krebbers. 2026. Backwards-Compatible Row-Based Exceptions in ML. *Proc. ACM Program. Lang.* 10, PLDI, Article 225 (June 2026), 25 pages. <https://doi.org/10.1145/3808303>

1 Introduction

Exceptions are widely used for error handling and writing efficient and concise programs using non-local control flow. Exceptions are also error-prone. Programmers should ensure exceptions are handled, to avoid unexpected crashes, and that exception names do not collide, to avoid catching more exceptions than intended. In this paper, we develop a type system that provides programmers with strong types to control the correct usage of exceptions and thereby avoid such programming mistakes. We formulate these claims in a novel relational separation logic for exceptions, the first of this kind. Moreover, we consider a realistic subset of ML-style languages, such as Standard ML and OCaml, and pay special attention to the compatibility issues between the standard typing discipline of such languages and our design. In sum, we focus on the following aspects:

Feature support. We should support two key features of exceptions in ML: the ability to treat exceptions as first-class data through the *extensible variant type* **exn**, which is important, for example, to catch an exception in one thread and re-raise it in another thread; and the ability to declare exceptions locally. **Backwards compatibility.** We should allow programmers to opt out from exception tracking in (parts of) their program. To achieve this, our system should enjoy two properties. First, it should be possible to encode the ordinary ML-style types and typing rules so

Authors’ Contact Information: [Simcha van Collem](mailto:simcha.vancollem@ru.nl), simcha.vancollem@ru.nl, Radboud University Nijmegen, Nijmegen, The Netherlands; [Paulo Emilio de Vilhena](mailto:p.de-vilhena@imperial.ac.uk), p.de-vilhena@imperial.ac.uk, Imperial College London, London, United Kingdom; [Robbert Krebbers](mailto:Robbert.Krebbers@ru.nl), mail@robbertkrebbers.nl, Radboud University Nijmegen, Nijmegen, The Netherlands.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

ACM 2475-1421/2026/6-ART225

<https://doi.org/10.1145/3808303>

that code without exception tracking continues to work, that is, our type system is an *extension* of the version without exception tracking. Second, it should be possible to generalize types one at a time without breaking old code. We call this second property *graceful generalization*. **Safety and abstraction theorems.** Our type system should ensure type and exception safety: a program with a sufficiently strong type does not raise uncaught exceptions. Moreover, it should guarantee *representation independence* [47, 56], which roughly says that, if a program uses a library or function through an abstract interface, then changes to the internal representation of this library or function cannot affect the program’s behavior. Representation independence has been studied extensively, particularly in the context of local state [2, 54, 63], but not in the context of both extensible variant types and exceptions. Finally, the type system should guarantee that, if a program uses an exception abstractly, then locally declared exception names cannot be used to catch this exception.

Our type system. We extend a language in the style of ML with rows [55, 71] and row polymorphism to enable exception tracking. As standard in type systems based on row typing [21, 32, 33, 40], function types are of the shape $\tau \rightarrow_{\rho} \sigma$, where the *row* ρ describes which exceptions can be raised. What is novel about our type system is that it is backwards compatible with ML-style type systems in two ways. First, it is an *extension*, which means that every typing rule from an ordinary ML-style type system can be derived from our generalized typing rules. The key idea to enable this encoding is (1) to introduce a *top row* \top , stating any exception can be raised and thus representing the lack of exception tracking, and (2) to see the ordinary function type $\tau \rightarrow \sigma$ as $\tau \rightarrow_{\top} \sigma$. Second, thanks to carefully designed yet simple subtyping rules, our system allows for a *graceful generalization* to exception tracking, where programs can be given a stronger type with exception tracking, one function at a time, while these types remain subtypes of their original counterparts without exception tracking. Let us illustrate this second aspect through an example:

```

with_return :  $\forall \alpha. ((\alpha \rightarrow \perp) \rightarrow \alpha) \rightarrow \alpha$ 
with_return f  $\triangleq$  let exn Return in try f ( $\lambda x. \mathbf{raise}$  (Return x)) with Return x  $\Rightarrow$  x
mapMexn :  $\forall \alpha, \beta. (\alpha \rightarrow \mathbf{option} \beta) \rightarrow \mathbf{list} \alpha \rightarrow \mathbf{option} (\mathbf{list} \beta)$ 
mapMexn f xs  $\triangleq$  with_return ( $\lambda \mathbf{return}. \mathbf{Some} (\mathbf{List.map} (\lambda x. \mathbf{match} f x \mathbf{with} \mathbf{Some} y \Rightarrow y \mid \mathbf{None} \Rightarrow \mathbf{return} \mathbf{None}) xs)$ )

```

The higher-order function $\mathbf{mapM}_{\text{exn}}$ maps a function f over a list xs using $\mathbf{List.map}$ of type $\forall \alpha, \beta. (\alpha \rightarrow \beta) \rightarrow \mathbf{list} \alpha \rightarrow \mathbf{list} \beta$ from the OCaml standard library. The function f might fail by returning \mathbf{None} , in which case $\mathbf{mapM}_{\text{exn}}$ breaks out of $\mathbf{List.map}$ and returns \mathbf{None} . Instead of using exceptions directly, the implementation relies on $\mathbf{with_return}$ (a simplification of the homonymous function from Jane Street [23]), which uses a *local exception* \mathbf{Return} to implement a *return* functionality, a common feature in imperative languages.

The function $\mathbf{mapM}_{\text{exn}}$ type checks in our system with the specified signature. Because ordinary types are implicitly annotated with \top , this signature does not track exceptions. This raises the question: how do we update $\mathbf{mapM}_{\text{exn}}$ ’s type to track exceptions without breaking existing clients of $\mathbf{mapM}_{\text{exn}}$? A similar question applies to how the types of $\mathbf{List.map}$ and $\mathbf{with_return}$ can be generalized without making $\mathbf{mapM}_{\text{exn}}$ ill-typed. In our system, all these types can be generalized one at a time, without invalidating the types remaining to be generalized. The key idea is that, when generalizing a type to a row-polymorphic one, the original type can still be obtained either by simple specialization of the row-polymorphic type to \top or via subtyping rules. Indeed, after generalizing $\mathbf{List.map}$ ’s type to $\forall \theta, \alpha, \beta. (\alpha \rightarrow_{\theta} \beta) \rightarrow_{\langle \rangle} \mathbf{list} \alpha \rightarrow_{\theta} \mathbf{list} \beta$, for example, the function $\mathbf{mapM}_{\text{exn}}$ remains typeable because we can instantiate θ with \top and subtype the empty row $\langle \rangle$ into \top . The same applies both to the generalization of $\mathbf{with_return}$ ’s type to $\forall \theta, \alpha. (\forall \theta'. (\alpha \rightarrow_{\theta'} \perp) \rightarrow_{\theta, \theta'} \alpha) \rightarrow_{\theta} \alpha$, where the row θ' abstracts the local exception \mathbf{Return} ; and to the generalization of $\mathbf{mapM}_{\text{exn}}$ ’s type to $\forall \theta, \alpha, \beta. (\alpha \rightarrow_{\theta} \mathbf{option} \beta) \rightarrow_{\langle \rangle} \mathbf{list} \alpha \rightarrow_{\theta} \mathbf{option} (\mathbf{list} \beta)$.

Of course, we are not the first to extend ML-style languages with exception tracking or to consider associated compatibility issues. Leroy and Pessaux [40] design an extension of OCaml with a powerful inference engine to provide a fine analysis of exceptions including the values with which they are raised (to support legacy uses of exceptions where error codes are represented by exception arguments). Our focus is on a declarative type system with a rich notion of polymorphism and subtyping, and on the formalization of the system's abstraction guarantees. Gradual typing [61] has also been used to address the migration from unchecked to checked effects [50, 57, 58]. Similar to our top row \top , these systems include a construct to account for the absence of effect tracking, but assign it a different semantics. Our approaches to backwards compatibility are also different. They allow for (dynamically checked) *down casts* to turn a function with a dynamic row into a function with a concrete row. We only allow *up casts* and use a rich notion of polymorphism and subtyping instead of dynamic checks. See §6 for more details.

Local exceptions. A key feature that is missing in prior work are local exceptions. To see why they are challenging, consider the following example:

```
with_return ( $\lambda$ return1. 20 + with_return ( $\lambda$ return2. return1 22))
```

The call to *return*₁ raises an exception that should be caught by the outer **with_return**, so this program should return 22. This behavior relies on **with_return** using a local exception and on the semantics of local exceptions in ML. If both instances of **with_return** used a global exception *Return*, then the names would collide and the program would return 42. Semantically, exceptions in Standard ML [46, §2.6] and OCaml [59] are not represented by their name, but by a fresh exception label that is dynamically generated by the exception declaration at runtime. One challenge that this entails is how to define the meaning of rows, in general, and of \top , in particular. Unlike Leroy and Pessaux [40], we cannot define the row \top as the set of all exception names that statically appear in the program: we should be careful about the scope of each exception.

Extensible variant types. Standard ML [46, §2.6] and OCaml [35] provide the *dynamically extensible variant type* **exn** to treat exceptions as first-class data. Each time a new exception is declared, it is added as a new variant constructor of **exn**. The **raise** construct takes any value of type **exn**. Using **try** *e*₁ **with** *x* \Rightarrow *e*₂ one can catch *any* exception (the bound variable *x* has type **exn**) and one can **match** on values of type **exn** to determine which exception is at hand.

Similar to Leroy and Pessaux [40], we annotate the **exn** _{ρ} type with a row ρ . This extension allows us to assign types to functions which use exceptions as first-class data. For example, the **Lazy** [38] and **Domain** [36] libraries in OCaml allow one to suspend computations or to execute computations in a new thread (called a *domain* in OCaml). In both cases, if the computation *f* raises *any* exception, this exception is stored in a reference and re-raised each time the computation is forced/joined. Using the **exn** _{ρ} type for the reference allows us to assign these libraries types that are polymorphic in the exceptions raised by *f*. Our general types for these libraries, in turn, can be subtyped to the ordinary OCaml types without exception tracking.

Safety and abstraction theorems. Using the Rocq prover and the Iris framework [25–30] we prove that our type system enjoys type and exception safety and representation independence. Since our type system is an extension of an ordinary ML-style type system, our proofs extend to ordinary ML with local exceptions and dynamically extensible variant types. Representation independence of ML with neither of these features has been considered in prior work.

Let us consider some examples to see how representation independence is relevant. The **map**_{**exn**} function we previously discussed should be *contextually equivalent* to a version **map**_{**opt**} (with the

same type) which internally uses the option monad instead of exceptions:

$$\mathbf{mapM}_{\text{opt}} f \, xs \triangleq \mathbf{match} \, xs \, \mathbf{with} \, [] \Rightarrow \mathbf{Some} \, [] \mid x :: xs' \Rightarrow f \, x \gg= \lambda y. \mathbf{mapM}_{\text{opt}} f \, xs' \gg= \lambda ys. \mathbf{Some} \, (y :: ys)$$

Contextual equivalence $e_1 =_{\text{ctx}} e_2$ means that e_1 can be replaced with e_2 in any typed program e without affecting e 's behavior. The equivalence between $\mathbf{mapM}_{\text{exn}}$ and $\mathbf{mapM}_{\text{opt}}$ relies on the exception *Return* (used in **with_return**) being private and, by consequence, on the function argument f not being able to throw *Return*. If *Return* were not private (that is, if it were declared as a global exception) or if the operational semantics did not guarantee freshness of exception labels, then the equivalence would fail: in $\mathbf{mapM}_{\text{exn}}$, a *Return* exception thrown by f would be handled accidentally by the **try** in **with_return**, whereas in $\mathbf{mapM}_{\text{opt}}$ such an exception would propagate.

There is a close relationship between local references and local exceptions. It is known that keeping local references private is crucial for representation independence [2, 54, 63]. Consider:

$$\mathbf{let} \, r = \mathbf{ref} \, 10 \, \mathbf{in} \, f \, (); !r =_{\text{ctx}} f \, (); 10 \quad \mathbf{let} \, \mathbf{exn} \, E \, \mathbf{in} \, \mathbf{try} \, f \, () \, \mathbf{with} \, E \, x \Rightarrow e =_{\text{ctx}} f \, ()$$

The first (well-known) equivalence holds because the reference r is local and the unknown function f is thus unable to modify it. Similarly, we observe that the second equivalence holds because the exception E is local and the unknown function f cannot raise it. The more complicated equivalence between $\mathbf{mapM}_{\text{exn}}$ and $\mathbf{mapM}_{\text{opt}}$ crucially relies on equivalences of the second form.

We develop a binary logical relations model from which safety and representation independence follows. We show that our model is rich enough to prove non-trivial contextual equivalences and refinements, even involving concurrency. We prove equivalences between $\mathbf{mapM}_{\text{exn}}$ and $\mathbf{mapM}_{\text{opt}}$ and that a version of the **Lazy** module from the OCaml standard library is a refinement of the **Domain** module. Finally, we prove various equivalences involving interesting implementations of extensible variant types introduced by Filinski in his PhD thesis [16, §4.5, Remark 4.30].

Our relational separation logic. We define our binary logical relations in a novel relational concurrent separation logic for exceptions. Although the approach of using a relational logic for defining binary logical relations is not new [7, 15, 63, 68], our logic is the first that supports (local) exceptions and extensible variant types. Our program logic has a number of novel ingredients.

First, inspired by Simuliris [19], we generalize the postcondition to range over expressions instead of values. This allows us to obtain simple reasoning rules and state simple specifications to relate code that raises exceptions to code that does not—a crucial ability for modular verification. In contrast to previous relational logics for effects [12] based on *biorthogonality* [53], which have a similar goal, our logic keeps a simple refinement judgment (without an extra parameter).

Second, we show that the generalization of the postcondition allows us to state a single invariant rule. In contrast, prior relational logics with support for invariants such as ReLoC [17, 18], have separate invariant rules for each atomic instruction (such as, load, store, and **cas**).

Third, based on Iris's higher-order ghost state [25] and on Timany et al. [64]'s *monotone partial bijections*, we develop a theory of *exception-signature assertions* to model rows and dynamically extensible variant types. Our theory supports variant constructors that refer to the variant type itself, even in negative position. We show how it enables reasoning about freshness of exception labels and how it extends to Filinski [16]'s implementations of extensible variant types.

Contributions. In sum, our contributions are as follows:

- We define ExceptionLang, a language with higher-order references, polymorphism, concurrency, and a type system based on rows for (local) exceptions (§2). We provide a *top row* \top and a rich notion of subtyping to ensure backwards compatibility with ML-style types.

$$\begin{aligned}
v &::= () \mid \mathbf{rec} \, f \, x = e \mid C \, v \mid \ell \mid l \, v \mid \dots & N &::= E \mid l \\
e &::= v \mid x \mid e \, e \mid C \, e \mid \overline{\mathbf{match} \, e \, \mathbf{with} \, C \, x \Rightarrow e} \mid \mathbf{ref} \, e \mid !e \mid e \leftarrow e \mid \mathbf{cas} \, e \, e \, e \mid \mathbf{fork} \{e\} \\
&\quad \mid \mathbf{let} \, \mathbf{exn} \, E \, \mathbf{in} \, e \mid N \, e \mid \mathbf{raise} \, e \mid \mathbf{try} \, e \, \mathbf{with} \, x \Rightarrow e \mid \mathbf{match} \, e \, \mathbf{with} \, N \, x \Rightarrow e \mid y \Rightarrow e \mid \dots \\
K &::= [] \mid e \, K \mid K \, v \mid l \, K \mid \mathbf{raise} \, K \mid \mathbf{try} \, K \, \mathbf{with} \, x \Rightarrow e \mid \mathbf{match} \, K \, \mathbf{with} \, l \, x \Rightarrow e \mid y \Rightarrow e \mid \dots
\end{aligned}$$

(a) Syntax.

$$\begin{aligned}
&\mathbf{try} \, v \, \mathbf{with} \, x \Rightarrow h \rightarrow_p v && (\mathbf{let} \, \mathbf{exn} \, E \, \mathbf{in} \, e, s) \rightarrow_t (e[l/E], \{l\} \uplus s, \epsilon) \\
&\mathbf{try} \, (\mathbf{raise} \, v) \, \mathbf{with} \, x \Rightarrow h \rightarrow_p h[v/x] && (\mathbf{fork} \{e\}, s) \rightarrow_t ((), s, [e]) \\
&\mathbf{match} \, (l \, v) \, \mathbf{with} \, l \, x \Rightarrow h \mid y \Rightarrow h' \rightarrow_p h[v/x] && (\mathbf{ref} \, v, s) \rightarrow_t (\ell, \{\ell \mapsto v\} \uplus s, \epsilon) \\
&\mathbf{match} \, (l' \, v) \, \mathbf{with} \, l \, x \Rightarrow h \mid y \Rightarrow h' \rightarrow_p h'[l' \, v/y] \text{ (if } l \neq l') \\
&\quad K[\mathbf{raise} \, v] \rightarrow_p \mathbf{raise} \, v \text{ (if neutral } K)
\end{aligned}$$

$\frac{\text{PURE-STEP} \quad e_1 \rightarrow_p e_2}{(e_1, s) \rightarrow_t (e_2, s, \epsilon)}$	$\frac{\text{BIND-STEP} \quad (e_1, s_1) \rightarrow_t (e_2, s_2, \vec{e}_f)}{(K[e_1], s_1) \rightarrow_t (K[e_2], s_2, \vec{e}_f)}$	$\frac{\text{TP-STEP} \quad (e_1, s_1) \rightarrow_t (e_2, s_2, \vec{e}_f)}{(\vec{e} \uplus e_1 :: \vec{e}', s_1) \rightarrow_{tp} (\vec{e} \uplus e_2 :: \vec{e}' \uplus \vec{e}_f, s_2)}$
--	--	---

(b) Operational semantics.

Fig. 1. Syntax and operational semantics of ExceptionLang.

- We develop a relational concurrent separation logic for compositional verification of ExceptionLang programs. The logic employs a generalized form of postcondition suitable to relate code that raises an exception to code that returns normally (§3). The generalized postcondition additionally enables us to state a single proof rule for Iris-style invariants.
- We develop a logical relations model for the ExceptionLang type system on top of our program logic (§4). The key ingredient is a higher-order ghost theory to concisely model rows and the extensible variant type **exn**. Using our logical relations model, we prove type and exception safety, as well as representation independence (§5).
- We mechanize all our results in the Rocq prover using Iris [69].

2 Language and Type System

We present the syntax and semantics of ExceptionLang (§2.1 and 2.2), its type system (§2.3), our approach to backwards compatibility (§2.4 and 2.5), and the safety and abstraction theorems (§2.6).

2.1 Syntax

Figure 1a shows the syntax of values, expressions, and evaluation contexts in ExceptionLang. The syntax of evaluation contexts reveals a call-by-value semantics with right-to-left evaluation order. Most of the value and expression constructs are standard; they closely follow the syntax of ML. In the following paragraphs, we discuss the unusual elements of the syntax.

Variants. The syntax $C \, e$ is used to construct values using one of the constructors C of a user-defined variant type.¹ Such values can be exhaustively pattern matched using the syntax $\mathbf{match} \, e \, \mathbf{with} \, \overline{C \, x \Rightarrow e}$, where the overline indicates a list of branches. Variant types are orthogonal to the dynamically extensible variant type **exn**.

¹The constructor C is a tag that, statically, indicates the type of its argument, and that, at runtime, indicates the pattern in a **match**. Tags could be represented by any set that fits the constructors used in the program, in Rocq we use strings.

Function declarations. Function declarations $\mathbf{rec} f x = e$ have one formal argument x and include a binder f denoting the function itself to allow recursive declarations. We obtain a standard syntax via syntactic sugar: non-recursive functions $\lambda x. e$ are defined as $\mathbf{rec} _ x = e$ and functions with multiple formal parameters $\lambda x. \dots \lambda y. e$ are defined as $\lambda x. \dots \lambda y. e$. Let bindings $\mathbf{let} x = e_1 \mathbf{in} e_2$ are defined as $(\lambda x. e_2) e_1$, and sequencing $e_1; e_2$ is sugar for $\mathbf{let} _ = e_1 \mathbf{in} e_2$.

Exceptions. The exception handler $\mathbf{try} e \mathbf{with} x \Rightarrow h$ handles any exception raised by e . If e raises an exception with $\mathbf{raise} v$, then control is transferred to the *handler branch* h with x bound to v . Any expression e can be used as the argument of \mathbf{raise} , however, as we will see (§ 2.3), the type system allows only *exception expressions*. Exception expressions are introduced as $E e$, where E is an *exception name*. At runtime, an exception name is bound to an *exception label* l . Exception labels are not part of the surface language, as indicated by the different color. The *local-exception declaration* $\mathbf{let} \mathbf{exn} E \mathbf{in} e$ introduces an exception name E whose scope is e . At runtime, this instruction allocates a fresh exception label l to which E is bound. The *exception-matching construct* $\mathbf{match} e \mathbf{with} E' x \Rightarrow h \mid y \Rightarrow h'$ expects e to reduce to an expression of the form $E v$. At runtime, if E and E' are bound to the same exception label, then the branch h is selected with x bound to v , otherwise the branch h' is selected with y bound to $E v$. Our syntax allows the same exception name E to be introduced in different parts of a program with different purposes. It is therefore important that the comparison be at the level of exception labels and that each occurrence of E be bound to a fresh label. Finally, we note that it is easy to define a construct for catching a specific exception as syntactic sugar (similar to Pierce [52, §14.3, Ex. 4]):

$$\mathbf{try} e \mathbf{with} E x \Rightarrow h \triangleq \mathbf{try} e \mathbf{with} x \Rightarrow (\mathbf{match} x \mathbf{with} E x \Rightarrow h \mid y \Rightarrow \mathbf{raise} y) \quad (1)$$

Absence of common constructs. For brevity, the presented syntax omits standard constructs such as integers, Booleans, pairs, arithmetic operations, and \mathbf{if} statements. All of these are included in our Rocq mechanization [69].

2.2 Operational Semantics

Figure 1b shows the semantics of ExceptionLang. The semantics is captured by a small-step *thread-pool reduction* \rightarrow_{tp} , a relation between two pairs (of the same type) consisting of a *thread pool* \vec{e}_1 , formalized as a list of expressions, and a *heap* s , formalized as a finite map where keys can be either *memory locations* ℓ that map to values, or exception labels l that map to a reserved symbol \ddagger . Intuitively, the assertion $(\vec{e}_1, s_1) \rightarrow_{\text{tp}} (\vec{e}_2, s_2)$ means that one of the expressions in \vec{e}_1 takes one step changing the state from s_1 to s_2 . This is captured by Rule TP-STEP, which stratifies the construction of \rightarrow_{tp} into the construction of the *thread-local reduction* \rightarrow_{t} , a relation formalizing the reduction of a single expression: formally, it is a relation between (1) a pair of an expression and a heap and (2) a tuple of an expression, a heap, and a list of expressions. Intuitively, the assertion $(e_1, s_1) \rightarrow_{\text{t}} (e_2, s_2, \vec{e})$ means that e_1 reduces to e_2 in one step, updating the state from s_1 to s_2 , and spawning new threads initialized with each of the expressions in \vec{e} . The construction of such assertions is stratified by Rule BIND-STEP, which allows one to focus on a subexpression e under an evaluation context K , or by Rule PURE-STEP, which allows one to ignore the state s in case of a *pure step* \rightarrow_{p} . Some base cases of \rightarrow_{p} and \rightarrow_{t} are shown in Figure 1b. They rephrase, mathematically, our informal discussion from § 2.1 about the constructs of ExceptionLang. The condition *neutral* K means that no exception handler ($\mathbf{try} K' \mathbf{with} x \Rightarrow e$) occurs in K , and ensures that exceptions are always caught by the nearest handler. The notation $\{l\} \cup s$ corresponds to the insertion of the key-value binding (l, \ddagger) in s plus the condition that l is not a key of s , and similar for $\{\ell \mapsto v\} \cup s$.

We omit various standard thread-local and pure reduction rules for brevity.

$$\tau ::= \mathbf{unit} \mid \tau \rightarrow_{\rho} \sigma \mid \overline{C \text{ of } \tau} \mid \forall \theta. \tau \mid \mathbf{exn}_{\rho} \mid \dots \quad \rho ::= \langle \rangle \mid \theta \mid \top \mid E \mid \rho \cdot \rho$$

(a) Types and rows.

$$\begin{array}{c}
\text{FORALL-ROW-INTRO} \\
\frac{\Gamma \vdash e : \langle \rangle : \tau \quad \mathit{value} \ e}{\Gamma \vdash e : \langle \rangle : \forall \theta. \tau} \\
\\
\text{LET-EXN-TYPED} \\
\frac{E \text{ fresh in } \Gamma, \rho, \tau \quad (E : \sigma), \Gamma \vdash e : \rho : \tau}{\Gamma \vdash \mathbf{let} \ \mathbf{exn} \ E \ \mathbf{in} \ e : \rho : \tau} \\
\\
\text{MATCH-EXN-TYPED} \\
\frac{(E : \sigma) \in \Gamma \quad \Gamma \vdash e : \rho : \mathbf{exn}_{E, \rho'} \quad (x : \sigma), \Gamma \vdash h : \rho : \tau \quad (y : \mathbf{exn}_{\rho'}), \Gamma \vdash h' : \rho : \tau}{\Gamma \vdash (\mathbf{match} \ e \ \mathbf{with} \ E \ x \Rightarrow h \mid y \Rightarrow h') : \rho : \tau} \\
\\
\text{TRY-TYPED} \\
\frac{\Gamma \vdash e : \rho : \tau \quad (x : \mathbf{exn}_{\rho}), \Gamma \vdash h : \rho' : \tau}{\Gamma \vdash \mathbf{try} \ e \ \mathbf{with} \ x \Rightarrow h : \rho' : \tau} \\
\\
\text{SUB-TYPED} \\
\frac{\Gamma \vdash e : \rho : \tau \quad \rho <: \rho' \quad \tau <: \tau'}{\Gamma \vdash e : \rho' : \tau'}
\end{array}$$

(b) Typing rules.

$$\begin{array}{c}
\text{ROW-SUB} \\
\frac{\mathit{Exns} \ \rho \subseteq \mathit{Exns} \ \rho' \quad \mathit{RVars} \ \rho \subseteq \mathit{RVars} \ \rho'}{\rho <: \rho'} \\
\\
\text{EXN-SUB} \\
\frac{\rho <: \rho'}{\mathbf{exn}_{\rho} <: \mathbf{exn}_{\rho'}} \\
\\
\text{FORALL-ROW-ELIM-SUB} \\
\forall \theta. \tau <: \tau[\rho/\theta] \\
\\
\text{ARROW-SUB} \\
\frac{\tau' <: \tau \quad \rho <: \rho' \quad \sigma <: \sigma'}{\tau \rightarrow_{\rho} \sigma <: \tau' \rightarrow_{\rho'} \sigma'} \\
\\
\text{FORALL-ROW-INTRO-SUB} \\
\frac{\tau <: \sigma \quad \theta \notin \tau}{\tau <: \forall \theta. \sigma}
\end{array}$$

(c) Subtyping rules.

Fig. 2. Type system of ExceptionLang.

2.3 Type System

The formal syntax of types and exception rows is defined in Figure 2a.² Recall that a row ρ intuitively denotes a set of exceptions and that, syntactically, it can either be empty $\langle \rangle$; a singleton E ; a row variable θ ; the top row \top ; or the union $\rho_1 \cdot \rho_2$ of ρ_1 and ρ_2 . An alternative (equivalent) syntactic view of a row ρ is as a pair of a set of exception names $\mathit{Exns} \ \rho$ and a set of row variables $\mathit{RVars} \ \rho$. Such a view can be obtained by *flattening* unions in ρ . This alternative view of rows is used in the subtyping rules (Figure 2c), which we discuss soon.

A selection of typing rules is shown in Figure 2b. We selected rules where rows play a non-trivial role. A typing judgment has the shape $\Gamma \vdash e : \rho : \tau$, where Γ is a *typing context* that maps variables x and exception names E to types. Its informal meaning is that e either diverges, or returns an output of type τ , or raises an exception included in ρ .

Rule **FORALL-ROW-INTRO** introduces a row-polymorphic type $\forall \theta. \tau$. The rule applies only to values. This restriction is known as the *value restriction*, which is necessary in the presence of both polymorphism and mutable references [72].³ Because values do not raise exceptions, the row is $\langle \rangle$.

²See our Rocq mechanization [69] for omitted types such as integers, references, value-polymorphism and recursive types.

³In Rocq, we formalize the value restriction as a separate typing judgment.

Rule **REC-TYPED** shows that a function $\mathbf{rec} f x = e$ has type $\tau \rightarrow_{\rho} \sigma$ if its body e has type σ and row ρ in an extended typing context where x has type τ and f has type $\tau \rightarrow_{\rho} \sigma$.

Rule **LET-EXN-TYPED** extends the typing context with the exception name E at an arbitrary type σ . We call σ the *type of E* . This type is used to constrain the type of e' in an exception value $E e'$. The rule requires E to be *fresh* to avoid clashes of exception names.

Rule **EXN-TYPED** places two conditions to the typing of an exception value $E e$: the exception name E must be present in the typing context and e must have the type σ to which E is bound. The concluding type \mathbf{exn}_E is the most precise, but, as we will see, it can be weakened by subtyping.

Rule **RAISE-TYPED** is relatively simple: $\mathbf{raise} e$ can be typed at any type τ , provided e has type \mathbf{exn}_{ρ} . Thanks to this type, only the constructors in ρ could have been used to build the exception value e . It is therefore sound to assign the row ρ to $\mathbf{raise} e$.

Rule **MATCH-EXN-TYPED** shows the typing of an exception-matching construct. It assumes the scrutinee e is an exception value of type $\mathbf{exn}_{E, \rho'}$ and that the branch h can be typed in an extended context where x has type σ (the type of E). Crucially, the rule also assumes that h' can be typed in an extended context where y has type $\mathbf{exn}_{\rho'}$. Intuitively, the typing of h' can exploit the fact that, if h' is selected, then the exception value returned by e could not have been constructed by E .

Rule **TRY-TYPED** shows that the type τ of an exception handler is the same as the type of its handlee e and the type of its handler branch h . Because every exception in e is captured by the exception handler, the handler raises exceptions only during the execution of h . This is why the concluding judgment has row ρ' . The typing of h can rely on a typing context extended with $(x : \mathbf{exn}_{\rho})$, because every exception raised by e uses one of the constructors in ρ .

Rule **SUB-TYPED** adds the possibility to weaken the type τ and the row ρ of a derived typing judgment to a type σ and row ρ' according to the subtyping relations $\rho <: \rho'$ and $\tau <: \sigma$. A selection of the subtyping rules appear in [Figure 2c](#). Again, we select those in which rows are important. Rule **ROW-SUB** rephrases the row subtyping relation $\rho <: \rho'$ as two set inclusions: $\mathit{Exns} \rho \subseteq \mathit{Exns} \rho'$ and $\mathit{RVars} \rho \subseteq \mathit{RVars} \rho'$. If \top occurs in ρ' , these inclusions are vacuously true. If \top occurs in ρ , then these inclusions demand that \top occur in ρ' as well. Rule **ARROW-SUB** adapts the standard subtyping of arrow types to include rows. Rule **EXN-SUB** allows one to loosen the information of the constructors used to build an exception value. Finally, Rules **FORALL-ROW-ELIM-SUB** and **FORALL-ROW-INTRO-SUB** are the elimination and introduction subtyping rules for polymorphic types.

The following typing rule for the construct defined in [Equation \(1\)](#) can be derived:

$$\text{TRY-EXN-TYPED} \frac{(E : \sigma) \in \Gamma \quad \Gamma \vdash e : E \cdot \rho : \tau \quad (x : \sigma), \Gamma \vdash h : \rho : \tau}{\Gamma \vdash \mathbf{try} e \mathbf{with} E x \Rightarrow h : \rho : \tau}$$

2.4 Extension of a Type System Without Exception Tracking

THEOREM 2.1 (EXTENSION). *For each ML-style typing derivation $\Gamma \vdash e : \tau$ without exception tracking, we can derive $\Gamma \vdash e : \top : \tau$ in *ExceptionLang*.*

We implicitly translate ordinary ML-style types to our type system by mapping the exception type \mathbf{exn} to \mathbf{exn}_{\top} and the function type $\tau \rightarrow \sigma$ to $\tau \rightarrow_{\top} \sigma$. This theorem is proved by induction on the typing derivation. In most cases, the typing judgments all have the same row, so the derivation of the ML typing rule with \top is direct. Let us consider a case where subsumption is also needed:

$$\frac{(E : \sigma) \in \Gamma \quad \Gamma \vdash e : \top : \mathbf{exn}_{\top} \quad (x : \sigma), \Gamma \vdash h : \top : \tau \quad (y : \mathbf{exn}_{\top}), \Gamma \vdash h' : \top : \tau}{\Gamma \vdash \mathbf{match} e \mathbf{with} E x \Rightarrow h \mid y \Rightarrow h' : \top : \tau}$$

To derive this case from **MATCH-EXN-TYPED**, we use **SUB-TYPED** on $\Gamma \vdash e : \top : \mathbf{exn}_{\top}$ and have to show that $\top <: E \cdot \top$. This subtyping holds vacuously because the row on the right includes \top .

2.5 Graceful Generalization via Subtyping

We say that our type system satisfies *graceful generalization*, which allows a function assigned a generalized type that tracks exceptions to be subtyped to its original type without exception tracking. For this property to be meaningful, the type system must be expressive enough to enable the type generalization of interesting libraries and functions, as we demonstrate in this section.

Recall the definition and signature of a simplified version of Jane Street’s `with_return` [23] (§1):

$$\begin{aligned} \text{with_return} &: \forall \theta, \alpha. (\forall \theta'. (\alpha \rightarrow_{\theta'} \perp) \rightarrow_{\theta, \theta'} \alpha) \rightarrow_{\theta} \alpha \\ \text{with_return } f &\triangleq \text{let exn } \text{Return} \text{ in try } f (\lambda x. \text{raise } (\text{Return } x)) \text{ with } \text{Return } x \Rightarrow x \end{aligned}$$

A call to `with_return` f , executes f applied to a *return* function. Consequently, f can terminate (1) with a *return* x call, (2) normally, or (3) by raising an exception. Raised exceptions by f are not caught by `with_return` as it only handles the local *Return* exception, which is used to implement the *return* call. The rank-2 row quantifier θ' in the type of `with_return` hides the use of a local exception and allows f to raise it only via a call to the *return* function.

To derive `with_return`’s ordinary type $\forall \alpha. ((\alpha \rightarrow \perp) \rightarrow \alpha) \rightarrow \alpha$, we use `SUB-TYPED` and should show the subtyping $(\forall \theta. \forall \alpha. (\forall \theta'. (\alpha \rightarrow_{\theta'} \perp) \rightarrow_{\theta, \theta'} \alpha) \rightarrow_{\theta} \alpha) <: (\forall \alpha. ((\alpha \rightarrow \perp) \rightarrow \alpha) \rightarrow \alpha)$. We first apply `FORALL-ROW-ELIM-SUB` to instantiate θ with \top , followed by a congruence rule for type polymorphism and `ARROW-SUB`. Next, we show $((\alpha \rightarrow \perp) \rightarrow \alpha) <: (\forall \theta'. (\alpha \rightarrow_{\theta'} \perp) \rightarrow_{\top, \theta'} \alpha)$, where the direction has swapped due to the contravariant nature of `ARROW-SUB`. We now apply `FORALL-ROW-INTRO-SUB` to introduce the universal quantifier on the right. The rest of the derivation reduces to showing $\theta' <: \top$ and $\top <: \top \cdot \theta'$, which hold vacuously because \top appears on the right.

We implement a version of OCaml’s `Domain` library [36] as follows:

$$\begin{aligned} \text{Domain.spawn} &: \forall \theta. \forall \alpha. (\text{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\langle \rangle} (\text{unit} \rightarrow_{\theta} \alpha) \\ \text{Domain.spawn } f &\triangleq \text{let } r = \text{ref Wait} \text{ in fork } \{ \text{try } r \leftarrow \text{Val } (f \ ()) \text{ with } ex \Rightarrow r \leftarrow \text{Exn } ex \}; \\ &\quad \lambda _ . \text{Domain.join } r \\ \text{Domain.join} &: \forall \theta. \forall \alpha. \text{ref } \{ \text{Wait} \mid \text{Val of } \alpha \mid \text{Exn of } \text{exn}_{\theta} \} \rightarrow_{\theta} \alpha \\ \text{Domain.join } r &\triangleq \text{match } !r \text{ with Wait} \Rightarrow \text{Domain.join } r \mid \text{Val } res \Rightarrow res \mid \text{Exn } ex \Rightarrow \text{raise } ex \end{aligned}$$

The function `Domain.spawn` takes a closure, runs it in a new thread, and returns a *join handler* h . Calling $h \ ()$ joins the thread and returns its result or re-raises any exception raised by $f \ ()$. This behavior is evident from its type: the returned join handler is annotated with the same row θ as the given closure. The derivation of the ordinary ML type $\forall \alpha. (\text{unit} \rightarrow \alpha) \rightarrow (\text{unit} \rightarrow \alpha)$ follows by instantiating θ with \top via `FORALL-ROW-ELIM-SUB` and `ARROW-SUB`.

Internally, `Domain.spawn` uses the native `fork` of `ExceptionLang` to run $f \ ()$ in a new thread. The reference r is used to track whether the computation is pending (`Wait`), has terminated normally (`Val`), or with an exception (`Exn`). The returned join handler uses the internal function `Domain.join` to perform a busy loop until the computation finishes (that is, it is no longer `Wait`). To type check this implementation, it is crucial that the exception type exn_{θ} is annotated with the row θ .

We implement a version of OCaml’s `Lazy` library [38] as follows:

$$\begin{aligned} \text{Lazy.from_fun} &: \forall \theta. \forall \alpha. (\text{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\langle \rangle} (\text{unit} \rightarrow_{\theta} \alpha) \\ \text{Lazy.from_fun } f &\triangleq \text{let } r = \text{ref Wait} \text{ in } \lambda _ . \text{Lazy.force } f r \\ \text{Lazy.force} &: \forall \theta. \forall \alpha. (\text{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\langle \rangle} \text{ref } \{ \text{Wait} \mid \text{Lock} \mid \text{Val of } \alpha \mid \text{Exn of } \text{exn}_{\theta} \} \rightarrow_{\theta} \alpha \\ \text{Lazy.force } f r &\triangleq \text{match } !r \text{ with} \\ &\quad \mid \text{Wait} \Rightarrow \text{if not (cas } r \text{ Wait Lock) then Lazy.force } f r \\ &\quad \quad \text{else try let } res = f \ () \text{ in } r \leftarrow \text{Val } res; res \\ &\quad \quad \text{with } ex \Rightarrow r \leftarrow \text{Exn } ex; \text{raise } ex \\ &\quad \mid \text{Lock} \Rightarrow \text{Lazy.force } f r \mid \text{Exn } ex \Rightarrow \text{raise } ex \mid \text{Val } res \Rightarrow res \end{aligned}$$

The function **Lazy.from_fun** has the same type as **Domain.spawn** and a similar behavior. The key difference is that the computation $f ()$ is not performed in a new thread, but upon the first call to the returned *suspension* h . Also, similar to **Domain.spawn**, if $f ()$ raises an exception, then every call to $h ()$ will raise that same exception, as evident from the row θ in the signature. The implementation uses an internal reference r to track the computation's state. To ensure that the computation $f ()$ is performed at most once, we add a **Lock** state. We use the atomic compare-and-set operation **cas** to ensure that only a single thread can move from the **Wait** to the **Lock** state, and perform $f ()$. Other threads perform a busy loop until the computation finishes (that is, it is no longer **Wait** or **Lock**).

We implement a version of OCaml's **Fun.protect** function [37] as follows:

$$\text{Fun.protect} : \forall \theta. \forall \alpha. (\text{unit} \rightarrow_{\langle \rangle} \text{unit}) \rightarrow_{\langle \rangle} (\text{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\theta} \alpha$$

Fun.protect *finally work* \triangleq **let** $y = (\text{try } \text{work} () \text{ with } x \Rightarrow \text{finally} (); \text{raise } x) \text{ in } \text{finally} (); y$

A call to **Fun.protect** *finally work* first performs *work* (), and then performs *finally* () regardless of whether *work* () returned normally or raised an exception. OCaml considers exceptions raised by *finally* a programming error, and performs a runtime check to handle these in a special way. We do not implement such a runtime check and instead rule out these cases statically by restricting the row of *finally* to $\langle \rangle$. To see **Fun.protect** in action, we implement a spin-lock:

$$\text{new_lock} : \text{unit} \rightarrow_{\langle \rangle} \forall \theta. \forall \alpha. (\text{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\theta} \alpha$$

new_lock () \triangleq **let** $r = \text{ref } \text{false}$ **in** $\lambda f. \text{acquire } r; \text{Fun.protect } (\lambda _ . \text{release } r) f$

A call to **new_lock** () returns a lock lk . Calling $lk f$ executes $f ()$ inside a critical section. The lock is implemented using a reference r that is **false** (unlocked) or **true** (locked). The internal function **acquire** performs a busy loop until it can update r from **false** to **true**, and **release** resets r to **false**. We wrap the call to f inside **Fun.protect** to prevent deadlocks when f raises an exception.

Finally, we consider three implementations of a module for dynamically extensible variant types:

$$\text{new_dyn}_{\text{exn}}, \text{new_dyn}_{\text{ref}} : \text{unit} \rightarrow_{\langle \rangle} \exists \alpha. \forall \beta. \text{unit} \rightarrow_{\langle \rangle} (\beta \rightarrow_{\langle \rangle} \alpha) \times (\alpha \rightarrow_{\langle \rangle} \text{option } \beta)$$

new_dyn_{exn} () \triangleq **pack** $\langle \text{exn}_{\top} \rangle (\lambda _ . \text{let } \text{exn } E \text{ in } ((\lambda x. E x), (\lambda d. \text{match } d \text{ with } E x \Rightarrow \text{Some } x \mid _ \Rightarrow \text{None})))$

new_dyn_{ref} () \triangleq **let** $lk = \text{new_lock} ()$ **in** **pack** $\langle \text{unit} \rightarrow_{\top} \text{unit} \rangle (\lambda _ . \text{let } r = \text{ref } \text{None}$ **in** $((\lambda x. \lambda _ . r \leftarrow \text{Some } x), (\lambda f. lk (\lambda _ . r \leftarrow \text{None}; f (); !r))))$

new_dyn_{ctr} () \triangleq **let** $c = \text{ref } 0$ **in** **pack** $\langle \text{int} \times \text{Obj.t} \rangle (\lambda _ . \text{let } i = \text{FAA } c \ 1 \ \text{in } ((\lambda x. (i, x)), (\lambda (j, x). \text{if } i = j \text{ then } \text{Some } x \text{ else } \text{None})))$

The implementations **new_dyn**_{exn} and **new_dyn**_{ref} are taken from Filinski [16, §4.5, Remark 4.30], while **new_dyn**_{ctr} is an *unsafe* implementation we wrote ourselves. All of these three functions generate a new dynamically extensible variant module, modeled as a closure that can be used to install new *constructor-destructor* pairs (c, d) . The destructor d is akin to a **match**. That is, a call to $d y$ returns **Some** x if y was obtained through the corresponding constructor $c x$, and returns **None** when the y has been obtained via a different constructor. For example:

let $\text{dyn} = \text{new_dyn}_{\text{exn}} ()$ **in** **let** $(c_1, d_1) = \text{dyn} \langle \text{int} \rangle ()$ **in** **let** $(c_2, d_2) = \text{dyn} \langle \text{bool} \rangle ()$ **in** $d_1 (c_2 \ \text{true})$ (* returns **None** *); $d_1 (c_1 \ 37)$ (* returns **Some** 37 *)

The implementation **new_dyn**_{exn} uses the **exn** type to represent variant constructors, and implements the destructor using **match**. The implementation **new_dyn**_{ref} uses a local reference r for each variant constructor. The constructor c returns a closure that sets the reference r to the argument value, and the destructor calls that closure and reads the reference r . If the reference contains **None** it means the constructor of a different variant was used. Compared to Filinski [16] we consider a concurrent language and need to wrap the destructor in a lock to ensure thread safety. Finally, **new_dyn**_{ctr}

represents a variant constructor as a pair of an integer tag (which we atomically increment for each new constructor) and the constructor argument. This implementation is not statically typeable in our language and requires the unsafe `Obj` module in OCaml [39]. Even though `new_dyncntr` is not typeable, one can reason about it semantically using logical relations [24, 63], as shown in §5.

2.6 Safety and Abstraction Guarantees

Safety. A program is *safe* if it either diverges or terminates with a result that is either a value or an unhandled exception (that is, a `raise` expression). A program is *exception safe*, if, on top of being type safe, it does not raise unhandled exceptions. A well-typed program is always type safe, and, if it type checks at an empty row, it is also exception safe:

THEOREM 2.2 (SAFETY). *If $\vdash e : \rho : \tau$, then e is safe, and, if $\rho = \langle \rangle$, then e is also exception safe.*

Abstraction. A key abstraction property enjoyed by well-typed programs is representation independence: changing the representation of a type does not affect programs that use this type abstractly. The main ingredient in formulating such abstraction properties is *contextual equivalence*. Informally, a program e_1 is *contextually equivalent* to e_2 , if, for any program e , replacing every occurrence of e_2 with e_1 in e does not change the behavior of e . Formally, contextual equivalence and refinement are defined as follows:

$$\begin{aligned} \Gamma \vDash e_1 \leq_{\text{ctx}} e_2 : \tau : \rho &\triangleq \forall (C : (\Gamma : \rho : \tau) \rightsquigarrow (\emptyset : \langle \rangle : \text{unit})), s. (C[e_1], s) \downarrow \Rightarrow (C[e_2], s) \downarrow \\ \Gamma \vDash e_1 =_{\text{ctx}} e_2 : \rho : \tau &\triangleq \Gamma \vDash e_1 \leq_{\text{ctx}} e_2 : \rho : \tau \wedge \Gamma \vDash e_2 \leq_{\text{ctx}} e_1 : \rho : \tau \end{aligned}$$

In words, an expression e_1 contextually refines e_2 if, for any *typed program context* C and heap s , the termination of $(C[e_1], s)$ implies the termination of $(C[e_2], s)$. Termination is captured by the assertion $(e_1, s) \downarrow$, stating that the main thread e_1 terminates to either a value or a `raise` expression. The relation $C : (\Gamma : \rho : \tau) \rightsquigarrow (\Gamma' : \rho' : \sigma)$ expresses that, for any e , if the judgment $\Gamma \vdash e : \rho : \tau$ is derivable, then $\Gamma' \vdash C[e] : \rho' : \sigma$ is derivable. Program contexts C are more general than evaluation contexts K , because the hole $[]$ in C does not need to appear in evaluation position.

In §5 we show that our type system (and by backwards compatibility, an ordinary ML-style language without exception tracking) satisfies many interesting equivalences and refinements. We first show simple equations indicating that local exceptions are private, before proving refinements between the examples in this section. Due to the quantification over *all* program contexts C , contextual refinement is, however, very hard to prove. As is standard in the literature, we prove refinements indirectly via logical relations (§4), which we define by means of our relational concurrent separation logic for exceptions (§3).

3 Relational Concurrent Separation Logic for Exceptions

We present the refinement judgment and adequacy theorem (§3.1) and proof rules (§3.2) of our relational logic for exceptions. We conclude with the support for concurrency (§3.3).

3.1 Refinement Judgment and Adequacy

The *refinement judgment* $e_1 \lesssim e_2 \{ \Phi \}$ is the mechanism through which the behaviors of programs e_1 and e_2 can be compared. In its most general form, the *postcondition* Φ is a binary relation on expressions. The judgment can then be intuitively read as asserting that either (1) the expression e_1 diverges or (2) both e_1 and e_2 respectively reduce to expressions e'_1 and e'_2 such that $\Phi e'_1 e'_2$ holds.

Our general type of postconditions allows a kind of *asymmetric context-local reasoning*, which we explain during the discussion of the *bind rule* (§3.2). However, it is often the case that, when stating the refinement between two programs, we would like the postcondition to express a relation between the *results* of these programs rather than arbitrary expressions. In a language

$$\begin{array}{c}
\text{STEP-L} \frac{e_1 \rightarrow_p e'_1 \quad \triangleright e'_1 \lesssim e_2 \{\Phi\}}{e_1 \lesssim e_2 \{\Phi\}}^* \quad \text{STEP-R} \frac{e_2 \rightarrow_p e'_2 \quad e_1 \lesssim e'_2 \{\Phi\}}{e_1 \lesssim e_2 \{\Phi\}}^* \quad \text{BASE} \frac{\Phi \ e_1 \ e_2}{e_1 \lesssim e_2 \{\Phi\}}^* \\
\text{WAND} \frac{e_1 \lesssim e_2 \{\Phi\} \quad \forall e'_1, e'_2. \Phi \ e'_1 \ e'_2 \ * \ \Phi' \ e'_1 \ e'_2}{e_1 \lesssim e_2 \{\Phi'\}}^* \quad \text{BIND} \frac{e_1 \lesssim e_2 \{e'_1, e'_2. K_1[e'_1] \lesssim K_2[e'_2] \{\Phi\}\}}{K_1[e_1] \lesssim K_2[e_2] \{\Phi\}}^* \\
\text{LET-EXN-L} \frac{\forall l_1. \text{exnLbl}_i \ l_1 \ * \triangleright e_1[l_1/E] \lesssim e_2 \{\Phi\}}{\text{let exn } E \text{ in } e_1 \lesssim e_2 \{\Phi\}}^* \quad \text{LET-EXN-R} \frac{\forall l_2. \text{exnLbl}_s \ l_2 \ * \ e_1 \lesssim e_2[l_2/E] \{\Phi\}}{e_1 \lesssim \text{let exn } E \text{ in } e_2 \{\Phi\}}^*
\end{array}$$

Fig. 3. Reasoning rules.

with exceptions, the result of a program is either a value or a **raise** expression. Precisely because the type of our postconditions is so general, it is easy to define specialized forms of the judgment:

$$\begin{aligned}
e_1 \lesssim e_2 \{r_1, r_2. R\} &\triangleq e_1 \lesssim e_2 \{e'_1, e'_2. e'_1, e'_2 \in \text{Result} \wedge R[e'_1/r_1, e'_2/r_2]\} \\
e_1 \lesssim e_2 \{\Phi_E \mid \Phi_V\} &\triangleq e_1 \lesssim e_2 \left\{ r_1, r_2. \vee \left\{ \begin{array}{l} \exists v_1, v_2. r_1 = v_1 \wedge r_2 = v_2 \wedge \Phi_V \ v_1 \ v_2 \\ \exists l_1, l_2, v_1, v_2. \wedge \left\{ \begin{array}{l} r_1 = \mathbf{raise} \ (l_1 \ v_1) \\ r_2 = \mathbf{raise} \ (l_2 \ v_2) \\ \Phi_E \ l_1 \ l_2 \ v_1 \ v_2 \end{array} \right. \right. \end{array} \right.
\end{aligned}$$

The specialized form $e_1 \lesssim e_2 \{r_1, r_2. R\}$ restricts the postcondition to a binary relation on expressions of type $\text{Result} \triangleq \text{Val} \cup \{\mathbf{raise} \ v \mid v \in \text{Val}\}$. The binders r_1 and r_2 , which we use exclusively to denote results, disambiguate between the general refinement judgment and this specialized form.

The specialized form $e_1 \lesssim e_2 \{\Phi_E \mid \Phi_V\}$ intuitively states that either (1) e_1 diverges or (2) both e_1 and e_2 return values related by Φ_V or (3) both e_1 and e_2 reduce to **raise** expressions related by Φ_E .

Model. The definition of the refinement judgment $e_1 \lesssim e_2 \{\Phi\}$ is inspired by the definition of *observational refinement* from ReLoC [17, 18]. The main differences are twofold: we employ an *extended weakest precondition* that supports exceptions and has an expression postcondition, and adopt a slightly different *invariant mask* semantics. For concision, we omit the formal definition and refer the interested reader to the Rocq mechanization [69] and §6 for these differences. Using the model, we obtain the adequacy theorem, showing that the logic is sound:

THEOREM 3.1 (ADEQUACY). *If $\vdash e_1 \lesssim e_2 \{\Phi_E \mid \Phi_V\}$, then (1) e_1 is safe, (2) if $(e_1, s) \downarrow$ then $(e_2, s) \downarrow$, and (3) if Φ_E is False, then e_1 is exception safe.*

3.2 Reasoning Rules

Figure 3 shows the reasoning rules of the logic. We omit the rules for reasoning about memory accesses (such as, load, store, and **cas**), which are standard in relational separation logic. Our rules for reasoning about concurrency constructs are less standard. We present them in detail in §3.3.

Rules **STEP-L** and **STEP-R** enable the expressions in either side of the refinement to be partially executed, provided these executions are pure.⁴

⁴To reason about a recursive function f , Iris employs a reasoning principle (called *Löb induction*) whereby f 's own specification can be introduced as an assumption to reason about recursive uses of f . To avoid cyclic proofs, where the assumed specification of f is immediately used to prove itself, Iris guards the assumed specification by a *later modality* \triangleright [7, 48]. To use the assumed specification, the modality must be eliminated. The \triangleright in front of the refinement judgment in **STEP-L** is one way in which \triangleright can be eliminated from assumptions in the proof context. Intuitively, the presence of \triangleright in **STEP-L** (and its absence in **STEP-R**) corresponds to the informal reading of the judgment, which allows e_1 to diverge regardless of e_2 .

Rule **BASE** lets one terminate a proof of refinement whenever the expressions e_1 and e_2 satisfy Φ . It is used to reason about expressions that return a value or raise an exception. For example, consider the following *derived* rules $\Phi_V v_1 v_2 \vdash v_1 \lesssim v_2 \{\Phi_E \mid \Phi_V\}$ and $\Phi_E l_1 l_2 v_1 v_2 \vdash \mathbf{raise}(l_1 v_1) \lesssim \mathbf{raise}(l_2 v_2) \{\Phi_E \mid \Phi_V\}$. Both rules are straightforward applications of **BASE**.

Reading from top to bottom, Rule **WAND** allows the postcondition Φ in a refinement judgment to be weakened to Φ' . The weakening condition is expressed as $\forall e'_1, e'_2. \Phi e'_1 e'_2 \multimap \Phi' e'_1 e'_2$. Intuitively, Rule **WAND** captures the principle that resources available before the execution of a program are still available afterwards provided the program does not interfere with these resources. This principle is more clearly visible in the *frame rule*, $R * e_1 \lesssim e_2 \{\Phi\} \vdash e_1 \lesssim e_2 \{e'_1, e'_2. R * \Phi e'_1 e'_2\}$, which can be stated as a rule derived by a straightforward application of **WAND**.

Rule **BIND** enforces *context-local reasoning*: it is sound to reason about the expressions e_1 and e_2 independently of the contexts in which they occur. There are two interesting aspects about **BIND**.

First, because the postcondition is not restricted to values (and, in fact, not even to results), when applying **BIND**, it is not needed that e_1 and e_2 terminate synchronously. This generalization allows a form of *asymmetric context-local reasoning*, whereby one can focus on a subexpression in the implementation side regardless of the specification side (and vice-versa). The following rule (and its specification-side counterpart) can be derived: $e_1 \lesssim e_2 \{e'_1, e'_2. K[e'_1] \lesssim e'_2 \{\Phi\}\} \vdash K[e_1] \lesssim e_2 \{\Phi\}$.

Second, unlike previous relational logics for non-local control effects where the bind rule is restricted to contexts K that do not contain handlers for the effects the subexpressions might perform, Rule **BIND** applies to any context, even to contexts that include **try-with** constructs. Indeed, the following rule (and its specification-side counterpart) can be derived:

$$\text{TRY-WITH-L} \frac{e_1 \lesssim e_2 \{e'_1, e'_2. (e'_1 \in \text{Val} \wedge e'_1 \lesssim e'_2 \{\Phi\}) \vee (\exists v. e'_1 = \mathbf{raise} v \wedge h_1[v/x] \lesssim e'_2 \{\Phi\})\}}{\mathbf{try} e_1 \text{ with } x \Rightarrow h_1 \lesssim e_2 \{\Phi\}}^*$$

Finally, Rule **LET-EXN-L** can be used to reason about the allocation of exceptions, which corresponds in the logic to the introduction of a resource $\text{exnLbl}_i l$. Intuitively, $\text{exnLbl}_i l$ asserts that l is fresh: $\text{exnLbl}_i l * \text{exnLbl}_i l' \vdash l \neq l'$. By default, $\text{exnLbl}_i l$ is non-duplicable, but it can be updated to a persistent assertion $\text{exnLbl}_i^\square l$, which can still be used for establishing freshness of exception labels: $\text{exnLbl}_i l * \text{exnLbl}_i^\square l' \vdash l \neq l'$.⁵ The reading of Rule **LET-EXN-R** and exnLbl_s is analogous.

Example. We illustrate how our reasoning rules can be used to derive the following rule:

$$\text{WITH-RETURN-L} \frac{\forall \text{return}, l. \left(\text{exnLbl}_i l * (\Box \forall v, \Phi', e'. \mathbf{raise}(l v) \lesssim e' \{\Phi'\} \multimap \text{return} v \lesssim e' \{\Phi'\}) \multimap \right.}{\mathbf{with_return} \text{ main} \lesssim e \{r_1, r_2. \Phi r_1 r_2\}}^* \left. \begin{array}{l} \text{main return} \lesssim e \left\{ r_1, r_2. \left\{ \begin{array}{ll} \Phi v r_2 & \text{if } r_1 = \mathbf{raise}(l v) \\ \Phi r_1 r_2 & \text{otherwise} \end{array} \right\} \right\} \end{array} \right)$$

This rule allows one to reason about occurrences of **with_return** on the left-hand side of the judgment.⁶ Specifically, to reason about the application of **with_return** to a client function main , it suffices to reason about the application of main to a function return , which stands for the function in **with_return** that raises the locally declared exception Return . Rule **WITH-RETURN-L** exposes the label l to which Return is bound at runtime to assert its freshness using $\text{exnLbl}_i l$. The postcondition for the refinement between main return and e reflects the behavior of the exception handler in the implementation of **with_return**: the case of a **raise** expression with label l unwraps the payload v before asserting the results to be related by Φ . In addition to $\text{exnLbl}_i l$, Rule **WITH-RETURN-L** also provides a specification of return that, in a sense, inlines the underlying definition of return : the

⁵Under the hood, a exnLbl predicate is defined as a points-to predicate with a discardable fractional permission [70].

⁶An analogous (derived) rule for occurrences of **with_return** on the right-hand side of the judgment exists.

$$\begin{array}{c}
\text{FORK-R} \\
\frac{\forall j. j \Rightarrow e_2 \text{ * } e_1 \lesssim () \{ \Phi \}}{e_1 \lesssim \text{fork} \{ e_2 \} \{ \Phi \}} \text{ *}
\end{array}
\qquad
\begin{array}{c}
\text{INV-OPEN} \\
\frac{\text{atomic}(e_1) \quad \boxed{P} \quad \triangleright P \text{ * } e_1 \lesssim e_2 \{ e'_1, e'_2. e'_1 \in \text{Result} \text{ * } \triangleright P \text{ * } \Phi e'_1 e'_2 \}}{e_1 \lesssim e_2 \{ \Phi \}} \text{ *}
\end{array}$$

Fig. 4. Reasoning rules for concurrency.

application *return v* refines any expression e' at any postcondition Φ' provided that **raise** (lv) refines e' at Φ' . Iris's persistence modality \square allows this specification to be used as many times as *return* occurs in the body of *main*.

The main rules used in the derivation of **WITH-RETURN-L** are Rule **LET-EXN-L**, to introduce l and *exnLbl*; l ; and Rule **TRY-WITH-L**, to reason about the exception handler installed by **with_return**.

3.3 Concurrency and Iris-Style Invariants

We explain how our relational logic supports concurrency by showing that **Lazy.from_fun** refines **Domain.spawn**, our versions of the OCaml libraries we discussed in §2.5.⁷ To prove this refinement, we use Iris's *invariants* \boxed{P} to assert that P holds at any moment. We show that, by leveraging the expressivity of postconditions ranging over expressions, we can state a single invariant-opening rule that applies to any atomic instruction. In contrast, prior relational logics such as ReLoC [17, 18] include one invariant-opening rule *per* atomic instruction (see §6 for details).

We state the refinement between **Lazy.from_fun** and **Domain.spawn** as follows:

$$\forall \Phi_E, \Phi_V, f_1, f_2. \left(f_1 () \lesssim f_2 () \{ \square \Phi_E \mid \square \Phi_V \} \text{ * } \right. \\
\left. \text{Lazy.from_fun } f_1 \lesssim \text{Domain.spawn } f_2 \{ g_1, g_2. \square g_1 () \lesssim g_2 () \{ \square \Phi_E \mid \square \Phi_V \} \} \right)$$

It is assumed that f_1 and f_2 either both raise an exception or both return a value. The refinement states that the functions g_1 and g_2 satisfy the same relational specification as f_1 and f_2 : from a logical point of view, g_1 and g_2 behave like f_1 and f_2 . This reading corresponds to the intuition that, under the hood, g_1 forces the execution of f_1 and g_2 reads the result of executing f_2 . The occurrences of the \square modalities allow g_1 and g_2 to be called repeatedly and their results to be shared freely.⁸

We now briefly explain the proof. After the allocation of a reference r_2 by **Domain.spawn**, we must reason about a **fork** instruction. To this end, we use Rule **FORK-R**, which lets us introduce a *ghost thread-pool assertion* [67] to represent the newly spawned thread with a fresh id j running the exception handler: $j \Rightarrow (\text{try } r_2 \leftarrow \text{Val } (f_2 ()) \text{ with } ex \Rightarrow r_2 \leftarrow \text{Exn } ex)$.⁹ After the allocation of r_1 by **Lazy.from_fun**, both **Lazy.from_fun** and **Domain.spawn** reach a point in which they return a value. Rule **BASE** can then be used to reduce the proof to:

$$\square (\lambda_. \text{Lazy.force } f_1 r_1) () \lesssim (\lambda_. \text{Domain.join } r_2) () \{ \square \Phi_E \mid \square \Phi_V \}$$

The modality \square prevents us from keeping exclusive ownership of resources (in particular, the points-to assertions and the ghost thread-pool assertion). This is problematic because **Lazy.force**, for example, accesses the location r_1 . Fortunately, in Iris, it is possible to express a relation among

⁷We can prove a refinement in this direction because the execution of the source is angelic: we can execute the spawned source thread from **Domain.spawn** at the moment the suspension is forced in **Lazy.from_fun**. The other direction does not hold because the scheduling of the target is demonic: the execution of the spawned thread could be interleaved arbitrarily or not executed at all, which is not possible for the computation of the suspension.

⁸In our Rocq mechanization we have also proven a *oneshot* refinement, where all persistence modalities are removed. This enforces that the closures are only called once but permits their results to contain exclusive resources.

⁹Note that exposing the ghost thread-pool assertion $j \Rightarrow e_2$ via **FORK-R** does not break abstraction. The only way to interact with such assertions is via de Vilhena et al. [12]'s logical fork rule, as we will see next. Since our logic does not provide the usual stepping rules for ghost thread-pool assertions from ReLoC, we treat them as an abstract resource.

$$\begin{aligned}
\text{Inv} \triangleq \vee \left\{ \begin{array}{l} (1) \quad \left(\begin{array}{l} r_1 \mapsto_i \text{Wait} * r_2 \mapsto_s \text{Wait} * f_1 () \lesssim f_2 () \{ \Box \Phi_E \mid \Box \Phi_V \} * \\ j \Rightarrow (\text{try } r_2 \leftarrow \text{Val } (f_2 ()) \text{ with } ex \Rightarrow r_2 \leftarrow \text{Exn } ex) \end{array} \right) \quad (1) \\ (2) \quad (r_1 \mapsto_i^{1/2} \text{Lock}) \quad (2) \\ (3) \quad (\exists l_1, l_2, v_1, v_2. r_1 \mapsto_i^\Box \text{Exn } (l_1 v_1) * r_2 \mapsto_s^\Box \text{Exn } (l_2 v_2) * \Box \Phi_E l_1 l_2 v_1 v_2) \quad (3) \\ (4) \quad (\exists v_1, v_2. r_1 \mapsto_i^\Box \text{Val } v_1 * r_2 \mapsto_s^\Box \text{Val } v_1 * \Box \Phi_V v_1 v_2) \quad (4) \end{array} \right.
\end{aligned}$$

Fig. 5. The invariant for the refinement between `Lazy.from_fun` and `Domain.spawn`.

the contents of r_1 and r_2 and the state of the running thread j as an *invariant assertion* $\Box P$, an assertion that holds at any moment. To introduce this assertion, the ownership of the resources in P is relinquished. In exchange, because $\Box P$ asserts only the *knowledge* that P holds at any point, it does not need to be exclusively owned: it is in fact persistent and can be freely shared.

The relation among r_1 , r_2 , and j is formalized as the assertion Inv shown in Figure 5. It represents the state system informally discussed in §2.5 as a disjunction of four cases. Case (1) is a snapshot of the current resources, case (2) asserts that someone currently holds the lock and is executing f_1 and f_2 , and cases (3) and (4) state that the computations have terminated with either related exceptions or related values. By relinquishing the resources we own, we can prove the first case, introduce $\Box \text{Inv}$, and keep it even after removing the \Box modality to simplify the goal to:

$$\text{Lazy.force } f_1 r_1 \lesssim \text{Domain.join } r_2 \{ \Box \Phi_E \mid \Box \Phi_V \}$$

To access r_1 , **Lazy.force** must claim its ownership back from Inv . It is possible to reclaim ownership of the resources in Inv (or to *open* the invariant $\Box \text{Inv}$) for the exact duration of a single execution step. This is sound because, as long as Inv is *preserved*, that is, as long as it holds at the beginning and at the end of each step, Inv holds at any point between two steps (as asserted by $\Box \text{Inv}$). Rule **INV-OPEN** captures this intuition precisely. Access to P is only granted for the duration of one step by requiring the expression on the implementation side to be atomic.¹⁰ If this expression is not atomic, the user can use **BIND** to focus on an atomic subexpression before applying **INV-OPEN**. The statement, in a relational setting, of a single general invariant-opening rule applicable to any atomic expression is an original contribution of ours.

Using Rule **INV-OPEN**, the rest of the proof is straightforward. If the **cas** operation succeeds, we can take out the resources of case (1), transition to case (2) to indicate that we currently hold the lock, and execute f_1 and f_2 , using de Vilhena et al. [12]’s logical fork rule, which allows one to temporarily swap the right-hand side of a refinement with an expression being executed by another thread. Finally, we update the invariant to either case (3) or (4), depending on whether f_1 and f_2 terminated with an exception or normally. Subsequent calls to g_1 and g_2 are in case (3) or (4), and can duplicate the contents of the persistently stored postconditions to finish the proof.

4 Logical Relations

To formally justify the safety and abstraction guarantees discussed in §2.6, we develop a *logical interpretation* of the type system through *logical relations*, where types are interpreted as value relations and typing judgments as relational specifications (§4.1). Based on the *logical approach* [63], we define our logical relations in terms of our relational program logic. Consequently, our safety theorem as well as standard properties of logical relations follow from adequacy of our logic (§4.2). Moreover, the abstraction guarantees follow from properties of the interpretation of types (§5). For

¹⁰For soundness, Iris also uses *invariant masks* \mathcal{E} and *namespaces* \mathcal{N} to prevent an invariant being opened twice in a nested fashion. We omit masks and namespaces for brevity.

$$\begin{aligned}
\Gamma \vDash e : \rho : \tau &\triangleq \Gamma \vDash e \leq_{\log} e : \rho : \tau \\
\Gamma \vDash e_1 \leq_{\log} e_2 : \rho : \tau &\triangleq \forall \eta, \gamma_1, \gamma_2. \llbracket \Gamma \rrbracket_{\eta \gamma_1 \gamma_2} * \gamma_1 e_1 \lesssim \gamma_2 e_2 \{ \llbracket \rho \rrbracket_{\eta \gamma_1 \gamma_2} \mid \llbracket \tau \rrbracket_{\eta \gamma_1 \gamma_2} \} \\
\tau <: \tau' &\triangleq \forall \eta, \gamma_1, \gamma_2, v_1, v_2. \llbracket \tau \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 * \llbracket \tau' \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 \\
\rho <: \rho' &\triangleq \forall \eta, \gamma_1, \gamma_2. (\rho)_{\eta \gamma_1 \gamma_2} \subseteq (\rho')_{\eta \gamma_1 \gamma_2}
\end{aligned}$$

(a) Interpretation of typing and subtyping judgments.

$$\llbracket \Gamma \rrbracket_{\eta \gamma_1 \gamma_2} \triangleq (*_{(x:\tau) \in \Gamma} \cdot \llbracket \tau \rrbracket_{\eta \gamma_1 \gamma_2} (\gamma_1 x) (\gamma_2 x)) * (*_{(E:\sigma) \in \Gamma} \cdot (\gamma_1 E, \gamma_2 E) \ltimes \llbracket \sigma \rrbracket_{\eta \gamma_1 \gamma_2})$$

(b) Interpretation of typing contexts.

$$\begin{aligned}
\llbracket \rho \rrbracket_{\eta \gamma_1 \gamma_2} l_1 l_2 v_1 v_2 &\triangleq (l_1, l_2) \in (\rho)_{\eta \gamma_1 \gamma_2} * \exists \Phi. (l_1, l_2) \ltimes \Phi * \square \Phi v_1 v_2 \\
\langle \theta \rangle_{\eta \gamma_1 \gamma_2} &\triangleq \eta \theta & \langle \langle \rangle \rangle_{\eta \gamma_1 \gamma_2} &\triangleq \emptyset & (\rho_1 \cdot \rho_2)_{\eta \gamma_1 \gamma_2} &\triangleq (\rho_1)_{\eta \gamma_1 \gamma_2} \cup (\rho_2)_{\eta \gamma_1 \gamma_2} \\
(E)_{\eta \gamma_1 \gamma_2} &\triangleq \{(\gamma_1 E, \gamma_2 E)\} & (\top)_{\eta \gamma_1 \gamma_2} &\triangleq \top
\end{aligned}$$

(c) Interpretation of rows.

$$\begin{aligned}
\llbracket \mathbf{unit} \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 &\triangleq v_1 = v_2 = () \\
\llbracket \mathbf{C of } \tau \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 &\triangleq \exists (C \text{ of } \sigma) \in \overline{\mathbf{C of } \tau}, w_1, w_2. v_1 = C w_1 * v_2 = C w_2 * \llbracket \sigma \rrbracket_{\eta \gamma_1 \gamma_2} w_1 w_2 \\
\llbracket \tau \rightarrow_{\rho} \sigma \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 &\triangleq \square \forall w_1, w_2. \llbracket \tau \rrbracket_{\eta \gamma_1 \gamma_2} w_1 w_2 * v_1 w_1 \lesssim v_2 w_2 \{ \llbracket \rho \rrbracket_{\eta \gamma_1 \gamma_2} \mid \llbracket \sigma \rrbracket_{\eta \gamma_1 \gamma_2} \} \\
\llbracket \forall \theta. \tau \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 &\triangleq \forall S. \llbracket \tau \rrbracket_{((\theta \mapsto S) \eta) \gamma_1 \gamma_2} v_1 v_2 \\
\llbracket \mathbf{exn}_{\rho} \rrbracket_{\eta \gamma_1 \gamma_2} v_1 v_2 &\triangleq \exists l_1, l_2, w_1, w_2. v_1 = l_1 w_1 * v_2 = l_2 w_2 * \llbracket \rho \rrbracket_{\eta \gamma_1 \gamma_2} l_1 l_2 w_1 w_2
\end{aligned}$$

(d) Interpretation of types.

Fig. 6. Semantic interpretation of the ExceptionLang type system.

now, it suffices to know that the abstraction guarantees are formulated as assertions relating the execution of two programs, so it is key that our logical relations be *binary*. Finally, we note that our interpretation of the extensible variant type **exn** and the interpretation of rows make use of a novel theory of *exception-signature assertions* (§4.3), otherwise we follow the setup of Timany et al. [63].

4.1 Semantic Interpretation of the Type System

We interpret a typing judgment $\Gamma \vdash e : \rho : \tau$ as a *semantic typing judgment* $\Gamma \vDash e : \rho : \tau$ and the subtyping relations $\tau <: \tau'$ and $\rho <: \rho'$ respectively as *semantic subtyping judgments* $\tau <: \tau'$ and $\rho <: \rho'$. These semantic judgments are logical assertions defined in terms of the interpretation of types and rows. A type τ is interpreted as a *semantic type* $\llbracket \tau \rrbracket$, a binary value relation such that the assertion $\llbracket \tau \rrbracket v_1 v_2$ is persistent for every v_1 and v_2 . Intuitively, the relation $\llbracket \tau \rrbracket$ represents the results of two related expressions of type τ . The persistence condition allows values to be freely shared as permitted by the unrestricted nature of the type system. A row ρ is interpreted as a *semantic row* $\llbracket \rho \rrbracket$, a relation on a label pair and a value pair. Intuitively, the relation $\llbracket \rho \rrbracket$ describes the exceptions raised by two related expressions with row ρ .

Figure 6a defines the semantic judgments. We explain the semantic subtyping judgments when discussing the interpretation of types and rows. A semantic typing judgment $\Gamma \vDash e : \rho : \tau$ is defined as a *logical refinement* $\Gamma \vDash e \leq_{\log} e : \rho : \tau$. A logical refinement between two expressions, e_1 and e_2 , is defined as a relational specification. It quantifies over all substitutions γ_1 and γ_2 that map variables to values and exception names to exception labels. The refinement requires that $\gamma_1 e_1$ refine $\gamma_2 e_2$

under all such substitutions. For the interpretation of the type system, the expressions e_1 and e_2 in a logical refinement are always the same. However, in §5, we use logical refinement as a way to establish contextual refinement (via [Theorem 4.1](#), to be explained in §4.2). To this end, it is important for logical refinement to be able to relate two different expressions. The substitutions γ_1 and γ_2 , as well as the map η of sets of label pairs for row variables, are parameters in the interpretation of types and rows. The substitutions γ_1 and γ_2 are needed to interpret exception names and η is needed to interpret row variables. For brevity, we omit them from subscripts in prose.

The postcondition of $\Gamma \vDash e_1 \leq_{\log} e_2 : \rho : \tau$ concludes that $\gamma_1 e_1$ and $\gamma_2 e_2$ either raise $\llbracket \rho \rrbracket$ -related exceptions or return $\llbracket \tau \rrbracket$ -related values. The precondition of $\Gamma \vDash e_1 \leq_{\log} e_2 : \rho : \tau$ requires γ_1 and γ_2 to be related, as defined in [Figure 6b](#): for each variable x bound to τ in Γ , the values $\gamma_1 x$ and $\gamma_2 x$ are related by $\llbracket \tau \rrbracket$ and, for each exception name E bound to σ in Γ , the labels $\gamma_1 E$ and $\gamma_2 E$, as well as the interpretation $\llbracket \sigma \rrbracket$ are related by the connective \bowtie . The *exception-signature* assertion $(l_1, l_2) \bowtie \Phi$ means l_1 and l_2 are valid labels uniquely associated with Φ . The interpretation of rows relies on this unique association to assert that values raised with these labels are related by Φ . The specific predicate Φ is chosen when the labels l_1 and l_2 are allocated. We discuss this further in §4.3 when we introduce the allocation rule for exception-signatures ([EXN-MAP-ALLOC](#)).

[Figure 6c](#) shows the interpretation of rows. The interpretation of a row ρ asserts (l_1, l_2) belongs to $\langle \rho \rangle$, which corresponds to the set of label pairs denoted by ρ . Its recursive definition is straightforward. Moreover, the interpretation of ρ asserts v_1 and v_2 are allowed payloads of the exceptions l_1 and l_2 . This is done by asserting the existence of a predicate Φ for which the exception-signature assertion holds and by which v_1 and v_2 are related. The exception-signature assertion essentially looks up the predicate which was chosen when the labels were allocated.

The semantic subtyping judgment $\rho <: \rho'$ means exceptions raised in the context of row ρ can be safely raised in a context of row ρ' . Such a property holds whenever $\llbracket \rho \rrbracket$ -related exceptions are also related by $\llbracket \rho' \rrbracket$. It suffices to assert $\langle \rho \rangle$ is included in $\langle \rho' \rangle$, as the rest of the interpretation does not depend on rows.

[Figure 6d](#) shows the interpretation of types. The interpretation of **unit** relates values equal to $()$. The interpretation of a variant type \overline{C} of τ relates values built with one of the constructors C of σ of the type using $\llbracket \sigma \rrbracket$ -related arguments. The interpretation of a function type $\tau \rightarrow_{\rho} \sigma$ relates functions v_1 and v_2 that, when applied to $\llbracket \tau \rrbracket$ -related values, either return $\llbracket \sigma \rrbracket$ -related values or raise $\llbracket \rho \rrbracket$ -related exceptions. The modality \square ensures the interpretation is persistent. The interpretation of a row polymorphic type $\forall \theta. \tau$ extends η by binding θ to a universally quantified set of label pairs S , thereby translating a *syntactic* \forall into a *semantic* \forall . The interpretation of \mathbf{exn}_{ρ} relates exception values $l_1 w_1$ and $l_2 w_2$ such that l_1, l_2, w_1 , and w_2 are related by $\llbracket \rho \rrbracket$, which, to recall, means (1) (l_1, l_2) belongs to the set of label pairs denoted by ρ and (2) w_1 and w_2 are values that can respectively be raised with l_1 and l_2 . The interpretation of the omitted types (such as, integers, references, and value-polymorphic types) follows the usual definition of logical relations in Iris [63]. It be found in our Rocq mechanization [69].

The semantic subtyping judgment $\tau <: \tau'$ means values returned in a context of type τ can also be returned in a context of type τ' . Such a property holds whenever $\llbracket \tau \rrbracket$ -related values are also related by $\llbracket \tau' \rrbracket$. This is exactly the definition of $\tau <: \tau'$.

4.2 Soundness and Fundamental Theorem

The following theorem shows logical refinement can be used to derive contextual refinement:

THEOREM 4.1 (SOUND LOGICAL RELATIONS). *If $\Gamma \vDash e_1 \leq_{\log} e_2 : \rho : \tau$, then $\Gamma \vDash e_1 \leq_{\text{ctx}} e_2 : \rho : \tau$.*

Its proof relies on adequacy of our relational logic ([Theorem 3.1](#)) among other theorems. It is used extensively in the next section (§5) to show contextual refinements.

$$\begin{array}{c}
\text{EXN-MAP-ALLOC} \quad \text{EXN-MAP-AGREE} \quad \text{EXN-MAP-DISAGREE} \\
\frac{\text{exnLbl}_i l_1 \quad \text{exnLbl}_s l_2}{\Rightarrow (l_1, l_2) \times \Phi}^* \quad \frac{(l_1, l_2) \times \Phi \quad (l_1, l'_2) \times \Psi}{l_2 = l'_2 * \triangleright \forall v_1, v_2. \Phi v_1 v_2 ** \Psi v_1 v_2}^* \quad \frac{(l_1, l_2) \times \Phi \quad (l'_1, l'_2) \times \Psi}{l_1 \neq l'_1 \Rightarrow l_2 \neq l'_2}^* \\
\text{EXN-MAP-PROJ} \quad (l_1, l_2) \times \Phi \vdash \text{exnLbl}_i^\square l_1 \wedge \text{exnLbl}_s^\square l_2
\end{array}$$

Fig. 7. Properties of exception-signature assertions.

The *Fundamental Theorem* shows that a well-typed program indeed meets the relational specification given by the semantic interpretation of its typing judgment:

THEOREM 4.2 (FUNDAMENTAL THEOREM). *If $\Gamma \vdash e : \rho : \tau$, then $\Gamma \vDash e : \rho : \tau$.*

This theorem creates the link between the type system and the logic and allows us to exploit soundness of the logic ([Theorem 3.1](#)) to show soundness of the type system ([Theorem 2.2](#)). Indeed, if a program e is well-typed at an empty typing context (as assumed by [Theorem 2.2](#)), that is, if $\vdash e : \rho : \tau$, then it follows by [Theorem 4.2](#) that the specification $\vDash e : \rho : \tau$ is derivable. Then, by [Theorem 3.1](#), it follows, in particular, that e is safe, and that, if $\rho = \langle \rangle$, then e is exception safe.

The proof of [Theorem 4.2](#) follows, as standard in the logical approach, by induction over the derivation of the (syntactic) typing judgment $\Gamma \vdash e : \rho : \tau$. This amounts to proving that, for each typing rule in [Figure 2](#), the *semantic version* of the rule, obtained by replacing syntactic judgments with semantic judgments, holds in the logic. Most cases are straightforward, following directly from the reasoning rules of the logic ([Figure 3](#)). The key cases are those related to exceptions. They depend on the properties of the exception-signature assertion \times described next.

4.3 A Theory of Exception-Signature Assertions

What properties should exception-signature assertions enjoy so that the semantic versions of the typing rules in [Figure 2](#) hold? The answer is in [Figure 7](#).

Rule [EXN-MAP-ALLOC](#) is used in the proof that the semantic version of [LET-EXN-TYPED](#) holds to discard the exnLbl assertions introduced through Rules [LET-EXN-L](#) and [LET-EXN-R](#) of the logic in exchange for an exception-signature assertion with $\llbracket \sigma \rrbracket$ as the value relation (where σ is the exception type of the declared exception E). This exception-signature assertion, on its turn, is used to justify the introduction of E at type σ in the typing context.

Rules [EXN-MAP-AGREE](#) and [EXN-MAP-DISAGREE](#) are used in the proof that the semantic version of [MATCH-EXN-TYPED](#) holds. During this proof, exception-signature assertions come from two places: from the interpretation of the typing context, which must contain the exception name E with type σ , and from the interpretation of **exn**, the type of the scrutinee. In short, Rules [EXN-MAP-AGREE](#) and [EXN-MAP-DISAGREE](#) are used to justify that the exception matching in both sides of the refinement in a semantic judgment either both succeed or both fail.

Rule [EXN-MAP-PROJ](#) is used to justify the freshness of newly allocated exception labels during proofs of contextual refinement (§5). Via this rule, because one can deduce persistent exnLbl^\square assertions hold for l_1 and l_2 , it is possible to conclude l_1 and l_2 are distinct from newly allocated ones for which exnLbl holds with full ownership.

The exception-signature assertion is defined using *monotone partial bijections* [64], *authoritative ghost state* [26, §6.3.3], and *saved predicates* [13, §5.1]. We omit the details, but emphasize that these Iris constructions act as building blocks for its model and to prove its properties.

5 Examples of Representation Independence

We prove contextual refinements between various example programs, including those from § 1 and § 2.5. Since contextual refinement is closed under subtyping,¹¹ we also immediately establish these refinements for their ordinary ML-style types without exception tracking.

Abstract exceptions. The following two equations characterize the property that, if a program uses an exception abstractly, then locally declared exceptions cannot be used to catch this exception:

$$(h : \tau) \vDash \lambda f. \mathbf{let\ exn\ } E \mathbf{ in\ try\ } f () \mathbf{ with\ } E _ \Rightarrow h =_{\text{ctx}} \lambda f. f () : \langle \rangle : \forall \theta. (\mathbf{unit} \rightarrow_{\theta} \tau) \rightarrow_{\theta} \tau$$

$$(h_1 : \tau), (h_2 : \tau) \vDash \lambda x. \mathbf{let\ exn\ } E \mathbf{ in\ match\ } x \mathbf{ with\ } E _ \Rightarrow h_1 \mid _ \Rightarrow h_2 =_{\text{ctx}} \lambda _. h_2 : \langle \rangle : \forall \theta. \mathbf{exn}_{\theta} \rightarrow_{\langle \rangle} \tau$$

Intuitively, because, in both statements, the row θ is polymorphic, both programs should treat the set of exceptions denoted by θ abstractly. In particular, the locally declared exception E cannot accidentally reveal the exceptions denoted by θ . The proof of both cases starts with [Theorem 4.1](#), allowing us to switch to a proof of refinement in the logic. Then, using the exclusive ownership of the label l to which E is dynamically bound and using [EXN-MAP-PROJ](#), we conclude that any exception raised by $f ()$ differs from l , so h and h_1 are never called.

Equivalent implementations of `mapM`. We show $\mathbf{mapM}_{\text{exn}}$ is contextually equivalent to $\mathbf{mapM}_{\text{opt}}$:

$$\vDash \mathbf{mapM}_{\text{exn}} =_{\text{ctx}} \mathbf{mapM}_{\text{opt}} : \langle \rangle : \forall \theta, \alpha, \beta. (\alpha \rightarrow_{\theta} \mathbf{option}\ \beta) \rightarrow_{\langle \rangle} (\mathbf{list}\ \alpha \rightarrow_{\theta} \mathbf{option}\ (\mathbf{list}\ \beta))$$

By definition of $=_{\text{ctx}}$, it suffices to prove contextual refinements on both directions. They are analogous, so let us consider the left-to-right direction. By [Theorem 4.1](#), we exchange \leq_{ctx} for \leq_{log} . The crux of the proof is then an application of [WITH-RETURN-L](#), which gives $\mathbf{mapM}_{\text{exn}}$ exclusive ownership of the exception label l used by `with_return`. By [EXN-MAP-PROJ](#), any exception raised by the iterated function must differ from l . We modularize the proof by stating a generic refinement between `List.map` and $\mathbf{mapM}_{\text{opt}}$. The statement of this refinement relies on the ability to relate code that raises an exception to code that terminates normally. We therefore use the $e_1 \lesssim e_2 \{r_1, r_2. R\}$ form instead of the $e_1 \lesssim e_2 \{\Phi_E \mid \Phi_V\}$ form of our refinement judgment.

Lazy and Domain. We prove that `Lazy.from_fun` contextually refines `Domain.spawn`:

$$\vDash \mathbf{Lazy.from_fun} \leq_{\text{ctx}} \mathbf{Domain.spawn} : \langle \rangle : \forall \theta. \forall \alpha. (\mathbf{unit} \rightarrow_{\theta} \alpha) \rightarrow_{\langle \rangle} (\mathbf{unit} \rightarrow_{\theta} \alpha)$$

By [Theorem 4.1](#), we exchange \leq_{ctx} for \leq_{log} . The proof then follows directly from the refinement in § 3.3 by (roughly) instantiating Φ_E with $\llbracket \theta \rrbracket$ and Φ_V with $\llbracket \alpha \rrbracket$.

Equivalent implementations of `new_dyn`. We show that the three implementations of dynamically extensible variant types (§ 2.5) are contextually equivalent, by proving a number of logical refinements between them. The proofs rely on picking the right Iris-style invariant and the right semantic interpretation for the existentially quantified type α . We delegate the interested reader to the Rocq mechanization for details and only highlight the key aspects here. First, we generalize the theory of exception-signatures assertions (§ 4.3). Instead of considering x_1 and x_2 in $(x_1, x_2) \times \Phi$ to be exception labels, we allow them to be of arbitrary types A_1 and A_2 . To prove the desired refinements, we let A_1 and A_2 be combinations of exception labels (for `new_dynexn`), memory locations (for `new_dynref`), and integers (for `new_dynctr`). Second, to prove that one implementation returns `Some` if and only if the other does so too, we generalize the rules [EXN-MAP-AGREE](#) and [EXN-MAP-DISAGREE](#). Finally, we emphasize that, while the implementation of `new_dynctr` is untyped, nothing prevents us from proving an equivalence involving an untyped implementation. The definition of the logical relation is stated purely in terms of the language semantics (encapsulated by our separation logic, which operates on untyped programs), not the type system.

¹¹This follows directly from the fact that typed program contexts themselves are closed under subtyping.

6 Related Work

Type systems and analyzers. Leroy and Pessaux [40] present a program analysis for estimating uncaught exceptions in ML programs. Their approach combines a type system with exception row polymorphism and control-flow analysis. The use of control-flow analysis is essential because, in languages such as OCaml, exception handlers may decide whether to handle an exception based on the values of its arguments. Consequently, the type system must track which values may be thrown. Leroy and Pessaux [40] prove type and exception safety of a core version of their system using progress and preservation, but they neither study representation independence nor consider local exceptions.

With a similar aim, Lermusiaux and Montagu [34] present a static analyzer for detecting uncaught exceptions in a large subset of OCaml, including local exceptions. Their approach differs from ours in that it is based on *abstract interpretation* rather than a type system, and does not support modular analysis of libraries and instead requires analyzing the entire program at once.

Schwerter et al. [57, 58] apply *gradual typing* [61] to address the related challenge of designing a system where *checked* and *unchecked* effects can co-exist and where the migration to *checked* effects is, in our terms, graceful. Their approach differs in terms of scope, flexibility, and semantics. Our scope is a statically typed language where exceptions are locally declared with a fixed signature. In contrast, Schwerter et al. [57] develop a generic effect calculus with a fixed effect theory, demonstrate how their framework can be used for exception handling, and later extend their work to a dynamically typed language with support for both gradual effect and type information [58]. Their approach is also more flexible than ours in its support for migrating unchecked to checked effects: functions can be migrated at any order and with support for both up and down casts (via dynamic checks). In our system, our subtyping rules allow only up casts: information can only be lost. Moreover, the order in which function signatures are generalized must respect the order in which these functions are defined. Their treatment of the *unknown effect* ζ , which is similar to our top row \top , also differs semantically. In particular, their unknown effect type is not absorbing, that is, $\{\zeta, E\}$ is not equivalent to $\{\zeta\}$, whereas in our system we have that $\top \cdot E$ is a subrow of \top , and *vice versa*. New et al. [50] study the different setting of algebraic effects and handlers and consider modules that may assign different signatures to effects. They prove *static* and *dynamic gradual guarantees* showing that the addition of annotations can only make programs harder to type check or to fail because of dynamic checks. Because our system does not rely on dynamic checks, such a guarantee does not reveal interesting properties of our system.

de Vilhena and Pottier [11] develop a type system for an ML-style language with dynamically allocated labels for local effects similar to ML's **let exception** construct. Like our work, they use row-annotated types to ensure (effect) safety. Unlike our work, (1) they do not consider extensible variant types, (2) they do not consider backwards-compatibility issues, and (3) their interpretation of rows requires the sets of labels denoted by two rows ρ_1 and ρ_2 in a row union $\rho_1 \cdot \rho_2$ to be disjoint. Their disjointness requirement is necessary for soundness, because their system allows the effect signature of an effect name to be changed whenever a handler is installed. In contrast, following OCaml's approach to fix the signature of an exception when it is declared, we can forgo this requirement. They prove type and effect safety of their system using a unary logical relations model defined in Iris, but do not study representation independence. As we discuss next, extending a system like TES for general effect handlers to a relational setting likely requires a complex model based on *biorthogonality* [8, 12], whereas we present a simpler model tailored to exceptions.

Program logics and logical relations. There has been much work on program logics and logical relations for language with non-local control flow, most of which in the unary case [1, 6, 10, 31, 41]. A notable exception is Maillard et al. [42], who develop the first relational program logic for exceptions

as part of a generic framework for relational program logics involving arbitrary monadic effects. Similar to our logic, they support postconditions on results that are either values or exceptions (but not expressions in general). However, they do not consider local exceptions, the extensible variant type `exn`, concurrency, nor separation-logic assertions and reasoning principles.

The literature also features various binary logical relation models for languages with (delimited) continuations (and, in particular, effect handlers) [8, 9, 15, 44, 62, 73]. Since continuations are more expressive than exceptions, they all take a different approach than us. They use *biorthogonality* [53] to define the logical relation in two steps: first relating evaluation contexts, and then relating expressions based on their behavior in related contexts. This approach works well in the context of continuations, but our work shows that the easier approach of generalizing the postcondition is sufficient for exceptions. Up to our knowledge, none of these works consider extensible variant types, nor factor the definition of their logical relation through a relational program logic.

Recent work by de Vilhena et al. [12] develops an Iris-based relational logic for effect handlers. They support local effects in the style of de Vilhena and Pottier [11], but do not consider OCaml's extensible variant type `Effect.t`. More broadly, they provide no connection to a type system and do not establish contextual refinement. Their logic enables compositional verification of programs that use different effects via *relational theories*, which is similar to relating code that raises exceptions to code that terminates normally. Despite the similar aim, their approach is different from ours. Their model is a complex recursive definition based on biorthogonality, whereas we use a simpler model inspired by ReLoC [17, 18], which is in turn inspired by Turon et al. [67].

Abstraction guarantees for exceptions. Zhang and Myers [73] prove the *absence of accidental handling* in a type system with *tunneled effects*, which in turn are inspired by *tunneled exceptions* [74]. Tunneled exceptions can be seen as the combination of `let exn` with an exception handler. The absence of a *catch-all* construct prevents them from supporting functions like `Domain.spawn`, `Lazy.from_fun`, and `Fun.protect`, but allows them to prove exceptions cannot accidentally be caught by a row-polymorphic function, strengthening their abstraction guarantees. Their language and type system are also different from ours: the core language is sequential and pure (that is, no mutable references) and the type system has no support for extensible variant types.

Comparison to ReLoC and Simuliris. Our work is inspired by ReLoC [17, 18], an Iris-based relational separation logic for contextual refinements, but without support for exceptions. Aside from exceptions, we perform two improvements that are of independent interest: (1) the refinement judgment in ReLoC limits postconditions to persistent predicates on values, whereas our refinement judgment allows arbitrary postconditions on expressions, and (2) our logic enjoys a novel general invariant-opening rule (`INV-OPEN`), whereas ReLoC includes one such rule per atomic instruction. We can state such a rule because our *invariant mask* semantics differs from the one in ReLoC. When opening an invariant, like Iris and our logic, ReLoC changes the mask in the relational judgment to prevent *reentrancy*: the same invariant opening twice. However, in ReLoC, any non \top mask prevents the left-hand side from reducing, which is needed to support invariants in combination with their persistent postconditions. This mask semantics makes a general invariant-opening rule unusable in ReLoC as one cannot even perform the atomic step on the left-hand side. In contrast, our logic follows the Iris-style mask semantics, allowing the left-hand side to reduce and requiring the invariant to be closed in the postcondition, made possible by the more general postconditions. Additionally, due to the postcondition ranging over expressions, one can prove the postcondition and close the invariant even if the program on the right-hand side has not terminated yet. This is essential in the refinement from § 3.3: we only step the right-hand side if the `cas` operation succeeds, which is only known after opening the invariant. These differences make ReLoC and our

logic incomparable, but we conjecture that every refinement provable in one can also be proven in the other, albeit with slightly different specifications and proof structures.

Simuliris [19] proves termination-preserving refinements in a step-indexing-free setting that uses parameterized coinduction [22] to reason about loops. Their approach requires the generalization to postcondition on expressions, an idea we adopt. Allain et al. [4] extend Simuliris to handle different calling conventions by generalizing the protocols of de Vilhena and Pottier [10]. Simuliris uses monotone partial bijections [64] to reason about freshness of memory locations, whereas we use them to model our exception-signature assertions.

Representation independence for extensible variant types. Our exception-signature assertions are similar to prior approaches to modeling logical relations for languages that support the allocation of fresh type names or include a dynamic type [3, 43, 49, 51, 66], where the latter is akin to our extensible variant type. Most closely related is New et al. [51] who also consider an extensible variant type (**OSum**), but do not track the allocated injections, meaning pattern matching cannot be exhaustive. They construct step-indexed logical relations [2] where the monotone partial bijection of allocated injections are explicitly tracked by the *possible worlds*. We instead model them using an Iris-style ghost theory that enables local reasoning about partial bijections. This makes it convenient to verify challenging examples, such as Filinski’s encoding of extensible variants, while keeping the reasoning about the global bijection hidden.

7 Future Work

Developing an algorithmic type system and investigating type inference is a natural next step. The type system as presented in this paper is an extension of System-F with a subtyping relation, making type checking and subtyping undecidable [65]. However, we do not expect that the full type inference or subtyping problem needs to be solved. For functions requiring complex (Rank-N) types, such as `with_return`, it is reasonable to expect that the programmer provides the type, as is also the case in Haskell [20]. We also expect that instantiation and elimination of rows could be restricted while retaining our applications for backwards compatibility. Specifically, it may suffice to only decide $\tau <: \sigma$ when one of the types has fewer quantifiers than the other. A potential algorithmic approach for this restricted fragment is to instantiate rows on the left with \top , and to subtype introduced rows on the right to \top . We conjecture that such an algorithm handles all examples in the paper, but a proper investigation remains as future work.

While we focused specifically on exceptions in the style of ML, it would be interesting to investigate whether our techniques can be applied to other languages, such as Java, where exceptions are classes, and therefore have a hierarchical structure. Another natural direction is to adapt the type system to support effect handlers by generalizing the exception type `exn` to OCaml’s effect type `'a Effect.t`, likely requiring a treatment similar to that of *generalized algebraic data types* (GADTs) in the logical relation [60], and a more expressive program logic. The logic blaze [12] supports the required language features, but currently has no connection to a type system and does not establish contextual refinement. Finally, it would be interesting to extend our type system with orthogonal features for exceptions, such as value tracking [40] or down casts [50, 58, 58].

It would be interesting to scale our work to a larger fragment of OCaml, for example, by using the recently developed semantics for OCaml by Seassau et al. [59] or Allain and Scherer [5]. We expect that extending our separation logic and logical relation requires one to investigate several orthogonal features, such as supporting OCaml’s underspecified evaluation order of arguments. Like most developments based on Iris, we consider sequentially consistent concurrency, while it would be interesting to provide support for OCaml’s weak memory model [14, 45].

Data Availability

A snapshot of the Rocq formalization is persistently available at [69].

Acknowledgments

We would like to thank François Pottier for suggesting many of the examples in this paper. Furthermore, we would like to thank Éric Tanter and the anonymous paper and artifact reviewers for their insightful suggestions. The first and third authors are supported, in part, by ERC grant COCONUT (grant no. 101171349), funded by the European Union; and grant [Cyclic Structures in Programs and Proofs](#) (file number OCENW.XL.23.089), funded by the Dutch Research Council (NWO). The second author, who, during the course of this work, was a member of the inspiring Veritas research group led by Azalea Raad, is supported by the UKRI Future Leaders Fellowship MR/V024299/1. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] Danel Ahman, Catalin Hritcu, Kenji Maillard, Guido Martínez, Gordon D. Plotkin, Jonathan Protzenko, Aseem Rastogi, and Nikhil Swamy. 2017. Dijkstra monads for free. In *POPL*. 515–529. <https://doi.org/10.1145/3009837.3009878>
- [2] Amal Ahmed, Derek Dreyer, and Andreas Rossberg. 2009. State-dependent representation independence. In *POPL*. 340–353. <https://doi.org/10.1145/1480881.1480925>
- [3] Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. 2011. Blame for all. In *POPL*. 201–214. <https://doi.org/10.1145/1926385.1926409>
- [4] Clément Allain, Frédéric Bour, Basile Clément, François Pottier, and Gabriel Scherer. 2025. Tail Modulo Cons, OCaml, and Relational Separation Logic. *PACMPL* 9, POPL (2025), 2337–2363. <https://doi.org/10.1145/3704915>
- [5] Clément Allain and Gabriel Scherer. 2026. Zoo: A framework for the verification of concurrent OCaml 5 programs using separation logic. *PACMPL* (2026). <https://clef-men.github.io/publications/allain-scherer-26.pdf>
- [6] Andrew W. Appel and Sandrine Blazy. 2007. Separation Logic for Small-Step Cminor. In *TPHOLs (LNCS, Vol. 4732)*. 5–21. https://doi.org/10.1007/978-3-540-74591-4_3
- [7] Andrew W. Appel, Paul-André Mellies, Christopher D. Richards, and Jérôme Vouillon. 2007. A very modal model of a modern, major, general type system. In *POPL*. 109–122. <https://doi.org/10.1145/1190216.1190235>
- [8] Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2018. Handle with care: Relational interpretation of algebraic effects and handlers. *PACMPL* 2, POPL (2018), 8:1–8:30. <https://doi.org/10.1145/3158096>
- [9] Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2020. Binders by day, labels by night: Effect instances via lexically scoped handlers. *PACMPL* 4, POPL (2020), 48:1–48:29. <https://doi.org/10.1145/3371116>
- [10] Paulo Emílio de Vilhena and François Pottier. 2021. A separation logic for effect handlers. *PACMPL* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434314>
- [11] Paulo Emílio de Vilhena and François Pottier. 2023. A Type System for Effect Handlers and Dynamic Labels. In *ESOP (LNCS, Vol. 13990)*. 225–252. https://doi.org/10.1007/978-3-031-30044-8_9
- [12] Paulo Emílio de Vilhena, Simcha van Collem, Ines Wright, and Robbert Krebbers. 2026. A Relational Separation Logic for Effect Handlers. *PACMPL* 10, POPL (2026), 981–1009. <https://doi.org/10.1145/3776676>
- [13] Mike Dodds, Suresh Jagannathan, Matthew J. Parkinson, Kasper Svendsen, and Lars Birkedal. 2016. Verifying Custom Synchronization Constructs Using Higher-Order Separation Logic. *TOPLAS* 38, 2 (2016), 4:1–4:72. <https://doi.org/10.1145/2818638>
- [14] Stephen Dolan, K. C. Sivaramakrishnan, and Anil Madhavapeddy. 2018. Bounding Data Races in Space and Time. In *PLDI*. 242–255. <https://doi.org/10.1145/3192366.3192421>
- [15] Derek Dreyer, Amal Ahmed, and Lars Birkedal. 2009. Logical Step-Indexed Logical Relations. In *LICS*. 71–80. <https://doi.org/10.1109/LICS.2009.34>
- [16] Andrzej Filinski. 1996. *Controlling Effects*. Ph. D. Dissertation. Carnegie Mellon University.
- [17] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *LICS*. 442–451. <https://doi.org/10.1145/3209108.3209174>
- [18] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2021. ReLoC Reloaded: A Mechanized Relational Logic for Fine-Grained Concurrency and Logical Atomicity. *LMCS* 17, 3 (2021). [https://doi.org/10.46298/LMCS-17\(3:9\)2021](https://doi.org/10.46298/LMCS-17(3:9)2021)
- [19] Lennard Gäher, Michael Sammler, Simon Spies, Ralf Jung, Hoang-Hai Dang, Robbert Krebbers, Jeehoon Kang, and Derek Dreyer. 2022. Simuliris: A separation logic framework for verifying concurrent program optimizations. *PACMPL*

- 6, *POPL* (2022), 1–31. <https://doi.org/10.1145/3498689>
- [20] GHC Team. 2026. Arbitrary-rank polymorphism. https://downloads.haskell.org/ghc/9.14.1/docs/users_guide/exts/rank_polymorphism.html
- [21] Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *TyDe@ICFP*. 15–27. <https://doi.org/10.1145/2976022.2976033>
- [22] Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. 2013. The power of parameterization in coinductive proof. In *POPL*. 193–206. <https://doi.org/10.1145/2429069.2429093>
- [23] Jane Street. 2025. Module Base.With_return. https://ocaml.janestreet.com/ocaml-core/v0.14/doc/base/Base/With_return/index.html
- [24] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: securing the foundations of the rust programming language. *PACMPL* 2, *POPL* (2018), 66:1–66:34. <https://doi.org/10.1145/3158154>
- [25] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *ICFP*. 256–269. <https://doi.org/10.1145/2951913.2951943>
- [26] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *JFP* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- [27] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. 637–650. <https://doi.org/10.1145/2676726.2676980>
- [28] Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A general, extensible modal framework for interactive proofs in separation logic. *PACMPL* 2, *ICFP* (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- [29] Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of Higher-Order Concurrent Separation Logic. In *ESOP (LNCS, Vol. 10201)*. 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- [30] Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive proofs in higher-order concurrent separation logic. In *POPL*. 205–217. <https://doi.org/10.1145/3009837.3009855>
- [31] Robbert Krebbers and Freek Wiedijk. 2013. Separation Logic for Non-local Control Flow and Block Scope Variables. In *FoSSaCS (LNCS, Vol. 7794)*. 257–272. https://doi.org/10.1007/978-3-642-37075-5_17
- [32] Daan Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *MSFP@ETAPS (EPTCS, Vol. 153)*. 100–126. <https://doi.org/10.4204/EPTCS.153.8>
- [33] Daan Leijen. 2017. Type directed compilation of row-typed algebraic effects. In *POPL*. 486–499. <https://doi.org/10.1145/3009837.3009872>
- [34] Pierre Lermusiaux and Benoît Montagu. 2024. Detection of Uncaught Exceptions in Functional Programs by Abstract Interpretation. In *ESOP (LNCS)*. 391–420. https://doi.org/10.1007/978-3-031-57267-8_15
- [35] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2025. OCaml manual (Version 5.4): Extensible variant types. <https://ocaml.org/manual/5.4/extensiblevariants.html>
- [36] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2025. OCaml manual (Version 5.4): Module Domain. <https://ocaml.org/manual/5.4/api/Domain.html>
- [37] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2025. OCaml manual (Version 5.4): Module Fun. <https://ocaml.org/manual/5.4/api/Fun.html>
- [38] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2025. OCaml manual (Version 5.4): Module Lazy. <https://ocaml.org/manual/5.4/api/Lazy.html>
- [39] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2025. OCaml manual (Version 5.4): Module Obj. <https://ocaml.org/manual/5.4/api/Obj.html>
- [40] Xavier Leroy and François Pessaux. 2000. Type-based analysis of uncaught exceptions. *TOPLAS* 22, 2 (2000), 340–377. <https://doi.org/10.1145/349214.349230>
- [41] Kenji Maillard, Danel Ahman, Robert Atkey, Guido Martínez, Catalin Hritcu, Exequiel Rivas, and Éric Tanter. 2019. Dijkstra monads for all. *PACMPL* 3, *ICFP* (2019), 104:1–104:29. <https://doi.org/10.1145/3341708>
- [42] Kenji Maillard, Catalin Hritcu, Exequiel Rivas, and Antoine Van Muylder. 2020. The next 700 relational program logics. *PACMPL* 4, *POPL* (2020), 4:1–4:33. <https://doi.org/10.1145/3371072>
- [43] Jacob Matthews and Amal Ahmed. 2008. Parametric Polymorphism through Run-Time Sealing or, Theorems for Low, Low Prices!. In *ESOP (LNCS, Vol. 4960)*. 16–31. https://doi.org/10.1007/978-3-540-78739-6_2
- [44] Craig McLaughlin. 2020. *Relational reasoning for effects and handlers*. Ph. D. Dissertation. University of Edinburgh, UK. <https://doi.org/10.7488/ERA/537>
- [45] Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2020. Cosmo: A concurrent separation logic for multicore OCaml. *PACMPL* 4, *ICFP* (2020), 96:1–96:29. <https://doi.org/10.1145/3408978>

- [46] Robin Milner and Mads Tofte. 1991. *Commentary on standard ML*. MIT Press.
- [47] John C. Mitchell. 1986. Representation Independence and Data Abstraction. In *POPL*. 263–276. <https://doi.org/10.1145/512644.512669>
- [48] Hiroshi Nakano. 2000. A Modality for Recursion. In *LICS*. 255–266. <https://doi.org/10.1109/LICS.2000.855774>
- [49] Georg Neis, Derek Dreyer, and Andreas Rossberg. 2009. Non-parametric parametricity. In *ICFP*. 135–148. <https://doi.org/10.1145/1596550.1596572>
- [50] Max S. New, Eric Giovannini, and Daniel R. Licata. 2023. Gradual Typing for Effect Handlers. 7, OOPSLA (oct 2023). <https://doi.org/10.1145/3622860>
- [51] Max S. New, Dustin Jamner, and Amal Ahmed. 2020. Graduality and parametricity: together again for the first time. *PACMPL* 4, POPL (2020), 46:1–46:32. <https://doi.org/10.1145/3371114>
- [52] Benjamin C. Pierce. 2002. *Types and programming languages*. MIT Press.
- [53] Andrew M. Pitts. 2005. Typed Operational Reasoning. In *Advanced Topics in Types and Programming Languages*, Benjamin C. Pierce (Ed.). MIT Press, Chapter 7, 245–289.
- [54] Andrew M. Pitts and Ian Stark. 1998. Operational reasoning for functions with local state. In *HOOTS*.
- [55] Didier Rémy. 1989. Typechecking Records and Variants in a Natural Extension of ML. In *POPL*. 77–88. <https://doi.org/10.1145/75277.75284>
- [56] John C. Reynolds. 1974. Towards a theory of type structure. In *Programming Symposium, Proceedings Colloque sur la Programmation (LNCS, Vol. 19)*. 408–423. https://doi.org/10.1007/3-540-06859-7_148
- [57] Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2014. A theory of gradual effect systems. In *IFCP*. 283–295. <https://doi.org/10.1145/2628136.2628149>
- [58] Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2016. Gradual type-and-effect systems. *JFP* 26 (2016), e19. <https://doi.org/10.1017/S0956796816000162>
- [59] Remy Seassau, Irene Yoon, Jean-Marie Madiot, and François Pottier. 2025. Formal Semantics and Program Logics for a Fragment of OCaml. *PACMPL* 9, ICFP (2025), 128–159. <https://doi.org/10.1145/3747509>
- [60] Filip Sieczkowski, Sergei Stepanenko, Jonathan Sterling, and Lars Birkedal. 2024. The Essence of Generalized Algebraic Data Types. *PACMPL* 8, POPL (2024), 695–723. <https://doi.org/10.1145/3632866>
- [61] Jeremy Siek and Walid Taha. 2006. Gradual typing for functional languages. *Scheme and Functional Programming*.
- [62] Amin Timany and Lars Birkedal. 2019. Mechanized relational verification of concurrent programs with continuations. *PACMPL* 3, ICFP (2019), 105:1–105:28. <https://doi.org/10.1145/3341709>
- [63] Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *JACM* 71, 6 (2024), 40:1–40:75. <https://doi.org/10.1145/3676954>
- [64] Amin Timany, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal. 2017. A logical relation for monadic encapsulation of state: proving contextual equivalences in the presence of runST. *PACMPL* 2, POPL (2017). <https://doi.org/10.1145/3158152>
- [65] Jerzy Tiuryn and Pawel Urzyczyn. 2002. The Subtyping Problem for Second-Order Types Is Undecidable. *Inf. Comput.* 179, 1 (2002), 1–18. <https://doi.org/10.1006/INCO.2001.2950>
- [66] Matías Toro, Elizabeth Labrada, and Éric Tanter. 2019. Gradual parametricity, revisited. *PACMPL* 3, POPL (2019), 17:1–17:30. <https://doi.org/10.1145/3290330>
- [67] Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*. 377–390. <https://doi.org/10.1145/2500365.2500600>
- [68] Aaron Joseph Turon, Jacob Thamsborg, Amal Ahmed, Lars Birkedal, and Derek Dreyer. 2013. Logical relations for fine-grained concurrency. In *POPL*. 343–356. <https://doi.org/10.1145/2429069.2429111>
- [69] Simcha van Collem, Paulo Emílio de Vilhena, and Robbert Krebbers. 2026. Rocq mechanization of “Backwards-Compatible Row-Based Exceptions in ML”. <https://doi.org/10.5281/zenodo.19494181>
- [70] Simon Friis Vindum and Lars Birkedal. 2021. Contextual refinement of the Michael-Scott queue (proof pearl). In *CPP*. 76–90. <https://doi.org/10.1145/3437992.3439930>
- [71] Mitchell Wand. 1991. Type Inference for Record Concatenation and Multiple Inheritance. *I&C* 93, 1 (1991), 1–15. [https://doi.org/10.1016/0890-5401\(91\)90050-C](https://doi.org/10.1016/0890-5401(91)90050-C)
- [72] Andrew K. Wright. 1995. Simple Imperative Polymorphism. *LISP Symb. Comput.* 8, 4 (1995), 343–355.
- [73] Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-safe effect handlers via tunneling. *PACMPL* 3, POPL (2019), 5:1–5:29. <https://doi.org/10.1145/3290318>
- [74] Yizhou Zhang, Guido Salvaneschi, Quinn Beightol, Barbara Liskov, and Andrew C. Myers. 2016. Accepting blame for safe tunneled exceptions. In *PLDI*. 281–295. <https://doi.org/10.1145/2908080.2908086>

Received 2025-11-14; accepted 2026-04-03