

A logical interpretation of Java-style exceptions

Jeffrey A. Vaughan

Harvard University

Abstract. This paper presents a novel, type-based exception analysis for EC (an Exception Calculus)—a higher-order, typed programming language with first-class, Java-style exceptions. EC exceptions are predeclared and subject to a nominal subtyping relation. Every exception has its own type, and generic exception handling is performed by handling a supertype of all applicable exceptions.

Typing is based on *result contexts*, which are analogous to the multiple-conclusion succedents in Gentzen’s LK. This paper proves the syntactic soundness of EC, and presents a shallow embedding of EC types and judgments in LK. The embedding gives rise to a pleasant interpretation of subtyping as logical consequence.

1 Introduction

Most programming languages allow expressions to produce side effects, such as mutating the heap or raising exceptions. Such effects make programs hard for programmers to reason about and compilers to optimize. This problem may be mitigated by type-based static analysis. Well known examples include IO monads as in Haskell, type and effect systems [18], and checked exceptions as in Java.

Haskell-style monadic reasoning is inspired by category theory, but many other effects analyses are ad-hoc. One notable exception is a line of work started by Nakano [13, 14] that performs exception tracking using intuitionistic variants of Gentzen’s [4] system LK. Additionally, languages such as $\lambda\mu$ and the dual calculus provide precise characterizations of some control effects in exotic reduction systems.

This paper presents a novel, type-based exception analysis for EC (an Exception Calculus)—a higher-order, typed programming language with Java-style exceptions. As in Java, EC exceptions are predeclared and subject to a nominal subtyping relation. Every exception has its own type, and generic exception handling is performed by handling a supertype of all applicable exceptions. Exceptions are first class and raised after being passed to a function.

This region of the design space differs from ML, which also features first class exceptions. ML exceptions can be created generatively at runtime [12], and ML programmers implement generic exception handlers using wildcard pattern matching, not subtyping. Also, all ML exceptions inhabit the single type `exn`.

EC’s exception analysis synthesizes ideas from Nakano [14] and from type-and-effects languages [18]. An expression is associated with both a return type

and the set of exceptions that it might raise. These are presented uniformly in a *result context*. Consider

$$e \equiv \text{if } b \text{ then } 3 \text{ else (raise Fail "oops").}$$

The typing judgment (modulo subtyping) assigns e the following result context.

$$e: (\text{ret}: \text{int}, \text{exn Fail}: \text{string})$$

Here we see that e has two *result channels*. It may return an `int`, or it may raise a `string`-carrying `Fail` exception. Abstracting over e internalizes the entire result context, not just the return channel.

$$\lambda b: \text{bool}. e: (\text{ret}: \text{bool} \rightarrow (\text{ret}: \text{int}, \text{exn Fail}: \text{string}))$$

This type captures e 's exceptions as *latent effects* [7]. Intuitively, latent effects are suspended exceptions which may be raised during a function call; here they correspond to the `throws` component of a Java method signature. Section 2 details EC's semantics and theory.

Because EC's type system assigns multiple-channel result contexts to expressions, it resembles sequent systems with multiple conclusions. Gentzen's LK [4] is the prototypical example. LK judgments have the form $\Phi \vdash \Theta$ where Φ and Θ are lists of propositions. The judgment is intended to hold when the conjunction of the Φ s implies the disjunction of the Θ s.

The relation between EC and LK extends beyond syntactic similarity. EC typing derivations have form $e: \Sigma; \Gamma \vdash \Delta$ where Γ assigns types to variables, and Δ is a result context. (Σ defines the exception hierarchy.) Intuitively, such a derivation may be translated into an LK derivation ending with $[\Gamma] \vdash [\Delta]$. Furthermore, EC subtyping derivations ending in $\Sigma \vdash \tau_1 <: \tau_2$ correspond to LK derivations ending in $[\tau_1] \vdash [\tau_2]$. Section 3 provides an explicit account of the translation and its properties.

EC's design embraces a tension between type and effect reasoning—which is precise but specialized to call-by-value reduction—and an approach based on a more literal formulation of logical realizers for LK. We intend for EC to represent a practical compromise between these models. By contrast, other logic-based effects systems stay closer to LK. For example, Nakano [13] can only type $\lambda x.e$ when e has no unhandled exceptions. (The programmer must use source-level “tag abstraction” to express e 's exception channels as part of its return type.) Another approach [14] types the entire abstraction using the same, unsuspending, effects as e . Such techniques appear to yield general analyses that are sound for full reduction, but lack EC's simplicity and precision for call-by-value evaluation.

Above all, the design of EC is motivated by a desire to make rigorous the sort of effects analysis performed in a language like Java, and to show that this analysis has a clear, proofs-as-programs interpretation. Toward this end, we keep EC's feature-set minimal and ignore interesting but extraneous issues such as modules, separate compilation, and effects polymorphism.

This paper's main contributions follow.

- We define system EC, a language with a sequent-style effects system. It treats latent effects properly for call-by-value evaluation.
- We prove syntactic soundness of EC.
- We present a shallow embedding of EC types and judgments in LK. This both refines Nakano’s treatment of effects and provides a pleasing interpretation of subtyping as logical consequence.

2 System EC

In this section we define system EC. EC is a call-by-value lambda calculus extended with exceptions. Its syntax is defined in Figure 1. The syntax of types is non-standard. In particular, arrows have form $\tau \rightarrow \Delta$, where Δ is an entire result context (see Section 2.1).

Exceptions are declared in signatures, Σ . A signature entry $E \sim \tau$ introduces a new exception named E carrying data of type τ . (Signature names are drawn from a distinguished set of identifiers.) Additionally, signatures may introduce subtyping constraints of form $E_1 <: E_2$. During evaluation, exception handlers for E_2 also handle E_1 exceptions. Type preservation requires that exception subtyping be consistent: E_1 ’s data must be a subtype of E_2 ’s data. This constraint is enforced by the well-formed signature judgment, $\Sigma \vdash \diamond$.

Exceptions are first class, and exception values are created using application syntax: $E e$. Expressions of the form $\text{raise } e$ raise an exception, and dynamically propagating exceptions are also represented using raise . The final exception-specific construct is $e_1 \text{ handle } E x \Rightarrow e_2$. This expression evaluates e_1 . If e_1 raises exception E_1 where $E_1 <: E$, then e_2 is evaluated with E_1 ’s data substituted for x . Otherwise the handle expression evaluates to the result of e_1 .

The formalization of EC’s dynamic semantics (Figure 2) is straightforward. Reduction is based on shallow evaluation contexts called *frames*. In addition to call-by-value lambda calculus reductions, EC has evaluation rules for exceptions. Rules E-HANDLE and E-PROP handle and propagate exceptions as discussed above. Choosing which rule to apply depends on dynamic knowledge of the subtyping hierarchy, and the evaluation relation is parameterized by a signature. The signature is also used by judgment $\Sigma \vdash E \curvearrowright F$, pronounced “E hops F.” This holds when frame F *does not* handle exception E . Note that for the special case of exception subtyping, a simple graph traversal suffices to decide $\Sigma \vdash \text{Exn } E_1 <: \text{Exn } E_2$ —there is no reason to invoke the typechecker at runtime.

2.1 Static Semantics

The full static semantics for EC are given by Figures 3 and 4. There we use $x \nmid \Delta$ to denote that a variable x is fresh with respect to Δ .

The typing relation is written $e : \Sigma; \Gamma \vdash \Delta$. The judgment’s unusual shape foreshadows the embedding we define in Section 3.1. *Result context* Δ may have multiple components, for instance $\text{ret} : \tau_1, \text{exn } E_2 : \tau_2, \text{exn } E_3 : \tau_3$. This can be read in two ways. The programmer’s reading is, “ e is a program that either

Result channels		Signatures	
$\varepsilon ::= \text{ret}$	Return channel	$\Sigma ::= \cdot$	Empty signature
$\text{exn } E$	Exception channel	$\Sigma, E \sim \tau$	Exception declaration
		$\Sigma, E_1 <: E_2$	Subtype declaration
Types		Expressions	
$\tau ::= A$	Type variable	$e ::= x \mid e_1 e_2 \mid \lambda x : \tau. e$	Lambda calculus
$\tau \rightarrow \Delta$	Arrow	top	Top/unit value
\top	Top	$E e$	Exception expression
$\text{Exn } E$	Exception type	raise e	Throw exception
		$e_1 \text{ handle } E x \Rightarrow e_2$	Exception handler
Result contexts		Values	
$\Delta ::= \cdot \mid \Delta, \varepsilon : \tau$		$v ::= \lambda x : \tau. e \mid E v \mid \text{top}$	
Variable contexts		Frames	
$\Gamma ::= \cdot \mid \Gamma, x : \tau$		$F ::= [] e \mid v []$	CVB application
		raise $[] \mid E []$	Exception construction
		$[] \text{ handle } E x \Rightarrow e$	Exception handler

Fig. 1. The syntax of EC.

returns a τ_1 or raises one of exception E_2 or E_3 .” In contrast, the logician reads, “ e is a proof that demonstrates the validity of τ_1 , or τ_2 , or τ_3 .” As we will see in Section 3.2, both are right!

The Typing Relation We will explain EC’s typing rules from the programmer’s perspective. Σ declares exceptions as discussed above, and Γ maps variables to their types. Δ is a list of *result channels*, each corresponding to either the result of evaluating e or a potentially uncaught exception. Note that Δ is not allowed to contain multiple occurrences of the same channel. (This is enforced by the well-formedness judgments).

T-VAR and T-TOP simply translate standard typing rules into our setting. Because such expressions always return, each only has a **ret** channel.

Rule T-EXN types expressions of form $E e$, which introduce first-class exception values of type $\text{Exn } E$. As e may raise exceptions during evaluation, its exception channels must also appear in $E e$ ’s result context.

Rule T-RAISE takes an exception of type $\text{Exn } E$ and throws it. There is no way for **raise** e to simply return a value; T-RAISE reflects this, as its conclusion does not contain a **ret** channel. The typing rule’s premise assumes that e may throw E —this is not a problem because subtyping will allow us to add such an exception channel when one is not otherwise present.

As usual, T-ABS types $\lambda x : \tau. e$ by augmenting the environment with $x : \tau$ and recursively checking e . However the function is assigned type $\tau \rightarrow \Delta$. Thus the arrow type constructor internalizes the *entire* result context. The logician may find this unexpected; perhaps he was anticipating something more like LK-

$\Sigma \vdash E \rightsquigarrow F$			
E-F-APP1 $\frac{}{\Sigma \vdash E \rightsquigarrow [] e}$	E-F-APP2 $\frac{}{\Sigma \vdash E \rightsquigarrow v []}$	E-F-EXN $\frac{}{\Sigma \vdash E_1 \rightsquigarrow E_2 []}$	E-F-RAISE $\frac{}{\Sigma \vdash E \rightsquigarrow \text{raise} []}$
E-F-HANDLE $\frac{\Sigma \not\prec \text{Exn } E_1 <: \text{Exn } E_2}{\Sigma \vdash E_1 \rightsquigarrow [] \text{handle } E_2 \Rightarrow e}$			
$\Sigma \vdash e \rightarrow e'$			
E-BETA $\frac{}{\Sigma \vdash (\lambda x: \tau. e) v \rightarrow e\{v/x\}}$	E-HANDLE $\frac{\Sigma \vdash \text{Exn } E_1 <: \text{Exn } E_2}{\Sigma \vdash (\text{raise } E_1 v) \text{handle } E_2 x \Rightarrow e \rightarrow e\{v/x\}}$		
E-PROPAGATE $\frac{\Sigma \vdash E \rightsquigarrow F}{\Sigma \vdash F[\text{raise } E v] \rightarrow \text{raise } E v}$	E-ENDHANDLER $\frac{}{\Sigma \vdash v \text{handle } E \Rightarrow e \rightarrow v}$	E-FRAME $\frac{\Sigma \vdash e \rightarrow e'}{\Sigma \vdash F[e] \rightarrow F[e']}$	

Fig. 2. Evaluation relation for EC.

IMPR from Figure 5. The programmer may be less surprised. The exception channels on the right of an arrow are *latent effects*, which are delayed until the corresponding function is applied [7].

Application behaves mostly as expected. T-APP’s premises mention three result contexts which must contain identical exception channels. As with T-RAISE, subtyping can coerce premises that would not otherwise match.

Rule T-HANDLE checks $e_1 \text{handle } E x \Rightarrow e_2$ as follows. Expression e_1 ’s result context is partitioned into Δ_1, Δ_2 where Δ_1 is made of exceptions beneath E in the subtype relation. That is, Δ_1 is a set of caught exceptions. Checking e_2 (in a suitably augmented environment) must then yield Δ_2 . As both e_1 and e_2 produce Δ_2 s, the two expressions must agree on *ret* channels (if any). Exception channels beneath E may be split between Δ_1 and Δ_2 ; this is useful if e_2 re-raises exceptions.

The final typing rule, T-SUB casts expressions to supertypes. As we have seen, the availability of subtyping simplifies other rules, such as T-APP.

Subtyping and Well-formedness EC contains three subtyping relations. These relate pairs of types, pairs of result channels, and pairs of result contexts.

The relation on types contains several standard rules familiar from system $F_{<}: S\text{-REFL}$, $S\text{-TRANS}$, and $S\text{-TOP}$. The arrow rule is contravariant on arguments and covariant on results. But this is not completely standard as the results are contexts, not basic types.

Exception subtyping is nominal and is governed by signature declarations. Signature well-formedness ensures that all declared subtypes respect covariance.

EC does not require that there be a top exception—a supertype of all other exceptions like Java’s `throwable`. Of course, such an exception may be defined.

Result context subtyping is more interesting. Intuitively, $\Sigma \vdash \Delta_1 <: \Delta_2$ when program contexts that handle all channels in Δ_2 also handle all channels in Δ_1 . Each of Δ_1 ’s exception channels must be a subtype of an exception channel in Δ_2 , and Δ_1 ’s return channel must be a subtype of Δ_2 ’s. For example, consider the following well-formed signature.

$$\begin{aligned} \Sigma &= \text{Disk} \sim \text{string}, \text{Sound} \sim \top, \text{IO} \sim \top, \\ &\text{Disk} <: \text{IO}, \text{Sound} <: \text{IO}. \end{aligned}$$

Then the following judgments hold.

$$\Sigma \vdash \text{exn Disk} : \text{string}, \text{exn Sound} : \top <: \text{exn IO} : \top$$

$$\Sigma \vdash \text{exn Sound} : \top <: \text{exn Sound} : \top, \text{exn Disk} : \text{string}$$

But these cannot be derived.

$$\Sigma \not\vdash \text{exn Disk} : \text{string}, \text{exn Sound} : \top <: \text{exn Disk} : \text{string}$$

$$\Sigma \not\vdash \text{ret} : \text{string} <: \text{exn IO} : \top$$

The programmer recognizes that result context subtyping lets him “over-annotate” an expression’s effects—a familiar operation in type and effect systems [19]. The logician views this as a form of hypothetical reasoning: $\Delta_1 <: \Delta_2$ when the disjunction of Δ_1 ’s formulas implies the disjunction of Δ_2 ’s. The logician’s view is made rigorous by Theorem 1.

Result channel subtyping ($\Sigma \vdash \varepsilon_1 : \tau_1 <: \varepsilon_2 : \tau_2$) is simple: return channels admit subtyping, but exception channels do not. It appears the latter restriction could be relaxed, but doing so does not provide an obvious advantage. Furthermore, the restriction helps maintain a useful invariant: each exception channel is annotated by the exception’s declared type (and not a subtype).

EC also includes several well-formedness judgments. These are needed to maintain three invariants. First, signatures, environments, and result contexts do not bind the same element twice. Second, exception references (e.g., in types) only name previously defined exceptions. Third, as mentioned above, subtype declarations are consistent with structural subtyping.

2.2 Metatheory

This section provides a brief overview of the main properties of EC.

The type system is set up such that the well-formedness of signatures and environments is checked at the leaves of typing derivations. Consequently EC enjoys a variety of regularity lemmas like the following.

Lemma 1. *Suppose $\mathcal{D} :: \Sigma \vdash \tau_1 <: \tau_2$. Then a subderivation of \mathcal{D} shows $\Sigma \vdash \diamond$.*

$$\boxed{e: \Sigma; \Gamma \vdash \Delta}$$

$$\frac{\text{T-SUB}}{e: \Sigma; \Gamma \vdash \Delta} \quad \frac{\Sigma \vdash \Delta_0 <: \Delta}{e: \Sigma; \Gamma \vdash \Delta}$$

$$\frac{\text{T-VAR}}{x: \Sigma; \Gamma \vdash \text{ret}: \tau} \quad \frac{(x: \tau) \in \Gamma \quad \Sigma \vdash \Gamma}{x: \Sigma; \Gamma \vdash \text{ret}: \tau}$$

$$\frac{\text{T-TOP}}{\text{top}: \Sigma; \Gamma \vdash \text{ret}: \top} \quad \frac{\Sigma \vdash \Gamma}{\text{top}: \Sigma; \Gamma \vdash \text{ret}: \top}$$

$$\frac{\text{T-ABS}}{\lambda x: \tau. e: \Sigma; \Gamma \vdash \text{ret}: \tau \rightarrow \Delta} \quad \frac{e: \Sigma; \Gamma, x: \tau \vdash \Delta}{\lambda x: \tau. e: \Sigma; \Gamma \vdash \text{ret}: \tau \rightarrow \Delta}$$

$$\frac{\text{T-APP}}{e_1 e_2: \Sigma; \Gamma \vdash \text{ret}: \tau, \Delta} \quad \frac{e_1: \Sigma; \Gamma \vdash \text{ret}: \tau_2 \rightarrow (\text{ret}: \tau, \Delta), \Delta \quad e_2: \Sigma; \Gamma \vdash \text{ret}: \tau_2, \Delta}{e_1 e_2: \Sigma; \Gamma \vdash \text{ret}: \tau, \Delta}$$

$$\frac{\text{T-EXN}}{E e: \Sigma; \Gamma \vdash \text{ret}: \text{Exn } E, \Delta} \quad \frac{e: \Sigma; \Gamma \vdash \text{ret}: \tau, \Delta \quad E \sim \tau \in \Sigma}{E e: \Sigma; \Gamma \vdash \text{ret}: \text{Exn } E, \Delta}$$

$$\frac{\text{T-RAISE}}{\text{raise } e: \Sigma; \Gamma \vdash \Delta, \text{exn } E: \tau} \quad \frac{e: \Sigma; \Gamma \vdash \text{ret}: \text{Exn } E, \Delta, \text{exn } E: \tau}{\text{raise } e: \Sigma; \Gamma \vdash \Delta, \text{exn } E: \tau}$$

$$\frac{\text{T-HANDLE}}{e_1 \text{ handle } E x \Rightarrow e_2: \Sigma; \Gamma \vdash \Delta_2} \quad \frac{\Sigma \vdash \Delta_1 <: \cdot, \text{exn } E: \tau \quad e_1: \Sigma; \Gamma \vdash \Delta_1, \Delta_2 \quad e_2: \Sigma; \Gamma, x: \tau \vdash \Delta_2}{e_1 \text{ handle } E x \Rightarrow e_2: \Sigma; \Gamma \vdash \Delta_2}$$

$$\boxed{\Sigma \vdash \tau_1 <: \tau_2}$$

$$\frac{\text{S-REFL}}{\Sigma \vdash \tau <: \tau} \quad \frac{\Sigma \vdash \tau}{\Sigma \vdash \tau <: \tau}$$

$$\frac{\text{S-TRANS}}{\Sigma \vdash \tau_1 <: \tau_3} \quad \frac{\Sigma \vdash \tau_1 <: \tau_2 \quad \Sigma \vdash \tau_2 <: \tau_3}{\Sigma \vdash \tau_1 <: \tau_3}$$

$$\frac{\text{S-TOP}}{\Sigma \vdash \tau <: \top} \quad \frac{\Sigma \vdash \tau}{\Sigma \vdash \tau <: \top}$$

$$\frac{\text{S-EXNT}}{\Sigma \vdash \text{Exn } E_1 <: \text{Exn } E_2} \quad \frac{\Sigma \vdash \diamond \quad E_1 <: E_2 \in \Sigma}{\Sigma \vdash \text{Exn } E_1 <: \text{Exn } E_2}$$

$$\frac{\text{S-ARR}}{\Sigma \vdash \tau_1 \rightarrow \Delta_1 <: \tau_2 \rightarrow \Delta_2} \quad \frac{\Sigma \vdash \tau_2 <: \tau_1 \quad \Sigma \vdash \Delta_1 <: \Delta_2}{\Sigma \vdash \tau_1 \rightarrow \Delta_1 <: \tau_2 \rightarrow \Delta_2}$$

$$\boxed{\Sigma \vdash \Delta_1 <: \Delta_2}$$

$$\frac{\text{S-NIL}}{\Sigma \vdash \cdot <: \Delta} \quad \frac{\Sigma \vdash \Delta}{\Sigma \vdash \cdot <: \Delta}$$

$$\frac{\text{S-CONS}}{\Sigma \vdash \Delta_1, \varepsilon_1: \tau_1 <: \Delta_2} \quad \frac{\Sigma \vdash \Delta_1 <: \Delta_2 \quad \varepsilon_2: \tau_2 \in \Delta_2 \quad \Sigma \vdash \varepsilon_1: \tau_1 <: \varepsilon_2: \tau_2}{\Sigma \vdash \Delta_1, \varepsilon_1: \tau_1 <: \Delta_2}$$

$$\boxed{\Sigma \vdash \varepsilon_1: \tau_1 <: \varepsilon_2: \tau_2}$$

$$\frac{\text{S-RET}}{\Sigma \vdash \text{ret}: \tau_1 <: \text{ret}: \tau_2} \quad \frac{\Sigma \vdash \tau_1 <: \tau_2}{\Sigma \vdash \text{ret}: \tau_1 <: \text{ret}: \tau_2}$$

$$\frac{\text{S-EXNC}}{\Sigma \vdash \text{exn } E_1: \tau_1 <: \text{exn } E_2: \tau_2} \quad \frac{\Sigma \vdash \text{Exn } E_1 <: \text{Exn } E_2 \quad \Sigma \vdash \text{exn } E_1: \tau_1 \quad \Sigma \vdash \text{exn } E_2: \tau_2}{\Sigma \vdash \text{exn } E_1: \tau_1 <: \text{exn } E_2: \tau_2}$$

Fig. 3. Static Semantics for EC. (1/2)

$\Sigma \vdash \diamond$

$$\begin{array}{c} \text{WF-S-NIL} \\ \hline \cdot \vdash \diamond \end{array} \qquad \begin{array}{c} \text{WF-S-DECL} \\ \Sigma \vdash \tau \\ \hline \Sigma, E \sim \tau \vdash \diamond \end{array} \text{ where } E \# \Sigma$$

$$\begin{array}{c} \text{WF-S-SUB} \\ E_1 \sim \tau_1 \in \Sigma \quad E_2 \sim \tau_2 \in \Sigma \quad \Sigma \vdash \tau_1 <: \tau_2 \\ \hline \Sigma, E_1 <: E_2 \vdash \diamond \end{array}$$

 $\Sigma \vdash \tau$

$$\begin{array}{c} \text{WF-T-VAR} \\ \Sigma \vdash \diamond \\ \hline \Sigma \vdash A \end{array} \qquad \begin{array}{c} \text{WF-T-TOP} \\ \Sigma \vdash \diamond \\ \hline \Sigma \vdash \top \end{array} \qquad \begin{array}{c} \text{WF-T-ARR} \\ \Sigma \vdash \tau \quad \Sigma \vdash \Delta \\ \hline \Sigma \vdash \tau \rightarrow \Delta \end{array} \qquad \begin{array}{c} \text{WF-T-EXN} \\ \Sigma \vdash \diamond \quad E \sim \tau \in \Sigma \\ \hline \Sigma \vdash \text{Exn } E \end{array}$$

 $\Sigma \vdash \Delta$

$$\begin{array}{c} \text{WF-C-NIL} \\ \Sigma \vdash \diamond \\ \hline \Sigma \vdash \cdot \end{array} \qquad \begin{array}{c} \text{WF-C-CONS} \\ \Sigma \vdash \Delta \quad \Sigma \vdash \varepsilon: \tau \\ \hline \Sigma \vdash \Delta, \varepsilon: \tau \end{array} \text{ where } \varepsilon \# \Delta$$

 $\Sigma \vdash \varepsilon: \tau$

$$\begin{array}{c} \text{WF-B-RET} \\ \Sigma \vdash \tau \\ \hline \Sigma \vdash \text{ret}: \tau \end{array} \qquad \begin{array}{c} \text{WF-B-EXN} \\ \Sigma \vdash \diamond \quad E \sim \tau \in \Sigma \\ \hline \Sigma \vdash \text{exn } E: \tau \end{array}$$

 $\Sigma \vdash \Gamma$

$$\begin{array}{c} \text{WF-E-NIL} \\ \Sigma \vdash \diamond \\ \hline \Sigma \vdash \cdot \end{array} \qquad \begin{array}{c} \text{WF-E-CONS} \\ \Sigma \vdash \Gamma \quad \Sigma \vdash \tau \\ \hline \Sigma \vdash \Gamma, x: \tau \end{array} \text{ where } x \# \Gamma$$

Fig. 4. Static Semantics for EC. (2/2)

Proof (Proof Sketch). Proof is by simultaneous induction with the following proposition: $\mathcal{D} :: \Sigma \vdash \tau$ implies that a subderivation of \mathcal{D} shows $\Sigma \vdash \diamond$.

Lemma 1 mentions subderivations to facilitate proofs by structural-induction.

Because EC includes subtyping, inverting typing judgments provides little traction for reasoning. We must prove a variety of canonical forms and inversion properties. One such lemma follows.

Lemma 2 (Canonical Form—Abs). *Assume*

$$v: \Sigma; \Gamma \vdash \text{ret}: \tau_2 \rightarrow (\text{ret}: \tau_1, \Delta_1), \Delta$$

then $v = \lambda x: \tau_3. e$ where $\Sigma \vdash \tau_2 <: \tau_3$.

Proof. By induction on the typing derivation with an inner induction for T-SUB.

The proof of progress is straightforward, but preservation requires several auxiliary lemmas. Lemma 5 is particularly interesting. It states that values can always be typed in a result context containing only a `ret` channel. Intuitively this holds because values do not reduce and therefore cannot raise exceptions. Lemma 6 establishes a useful invariant of exception and frame typing.

Lemma 3 (Progress). *Assume $e : \Sigma; \cdot \vdash \Delta$. Then either e is value, e has form $\text{raise } E v$ (i.e., is an uncaught exception), or $\Sigma \vdash e \rightarrow e'$.*

Proof. Proof by induction on the typing derivation. Case T-APP uses Lemma 2.

Lemma 4 (Substitution). *Suppose $e : \Sigma; \Gamma, x : \tau, \Gamma' \vdash \Delta$ and $v : \Sigma; \Gamma, \Gamma' \vdash \text{ret} : \tau$, then $e\{v/x\} : \Sigma; \Gamma, \Gamma' \vdash \Delta$.*

Lemma 5 (Value Strengthening). *Suppose $v : \Sigma; \Gamma \vdash \text{ret} : \tau, \Delta$, then $v : \Sigma; \Gamma \vdash \text{ret} : \tau$*

Proof. Lemmas 4 and 5 follow from straightforward induction.

Lemma 6 (Unhandled Exception Frame Typing). *Suppose $\Sigma \vdash E \curvearrowright F$ and $F[\text{raise } E v] : \Sigma; \Gamma \vdash \Delta$. Then there exist E' and τ' such that $\text{exn } E' : \tau' \in \Delta$ where $E <: E' \in \Sigma$.*

Proof. Direct using elided inversion principles.

Lemma 7 (Preservation). *Assume $\Sigma \vdash e \rightarrow e'$. For all Δ such that $e : \Sigma; \cdot \vdash \Delta$, it follows that $e' : \Sigma; \cdot \vdash \Delta$.*

Proof. By induction on the reduction relation.

3 The logical content of EC programs

Thus far we have discussed EC from the perspective of a programmer, as a language with a sound, type-based exception analysis. This section shows that EC has a logical reading, namely typing derivations in EC can be translated to proof trees in Gentzen's system LK [4].

To begin, we review LK. LK is a logic which operates on propositions drawn from the following grammar.

$$\begin{aligned} &\text{LK Propositions} \\ P, Q ::= A_{\text{LK}} \mid \top \mid \perp \mid P \wedge Q \mid P \vee Q \mid P \supset Q \mid \neg P \end{aligned}$$

Atomic propositions, A_{LK} are subscripted to avoid confusion with EC types. LK judgments have the form $\Phi \vdash \Theta$, where Φ —the *antecedent*—and Θ —the *succedent*—are lists of propositions. Intuitively, $\Phi \vdash \Theta$ holds when the conjunction of the Φ s implies the disjunction of the Θ s. The rules of deduction for propositional LK are given by Figure 5.

$\Phi \vdash \Theta$				
LK-ID $\frac{}{P \vdash P}$	LK-THINL $\frac{\Phi \vdash \Theta}{P, \Phi \vdash \Theta}$	LK-THINR $\frac{\Phi \vdash \Theta}{\Phi \vdash \Theta, P}$	LK-CONTRACTL $\frac{P, P, \Phi \vdash \Theta}{P, \Phi \vdash \Theta}$	LK-CONTRACTR $\frac{\Phi \vdash \Theta, P}{\Phi \vdash \Theta, P, P}$
LK-INTERL $\frac{\Phi_1, Q, P, \Phi_2 \vdash \Theta}{\Phi_1, P, Q, \Phi_2 \vdash \Theta}$		LK-INTERR $\frac{\Phi \vdash \Theta_1, Q, P, \Theta_2}{\Phi \vdash \Theta_1, P, Q, \Theta_2}$		LK-CUT $\frac{\Phi_1 \vdash \Theta_1, P \quad P, \Phi_2 \vdash \Theta_2}{\Phi_1, \Phi_2 \vdash \Theta_1, \Theta_2}$
LK-ANDL1 $\frac{P, \Phi \vdash \Theta}{P \wedge Q, \Phi \vdash \Theta}$		LK-ANDL2 $\frac{Q, \Phi \vdash \Theta}{P \wedge Q, \Phi \vdash \Theta}$		LK-ANDR $\frac{\Phi \vdash \Theta, P \quad \Phi \vdash \Theta, Q}{\Phi \vdash \Theta, P \wedge Q}$
LK-ORL $\frac{P, \Phi \vdash \Theta \quad Q, \Phi \vdash \Theta}{P \vee Q, \Phi \vdash \Theta}$		LK-ORR1 $\frac{\Phi \vdash \Theta, P}{\Phi \vdash \Theta, P \vee Q}$	LK-ORR2 $\frac{\Phi \vdash \Theta, Q}{\Phi \vdash \Theta, P \vee Q}$	
LK-IMPL $\frac{\Phi_1 \vdash \Theta_1, P \quad Q, \Phi_2 \vdash \Theta_2}{P \supset Q, \Phi_1, \Phi_2 \vdash \Theta_1, \Theta_2}$			LK-IMPR $\frac{P, \Phi \vdash \Theta, Q}{\Phi \vdash \Theta, P \supset Q}$	
LK-NOTL $\frac{\Phi \vdash \Theta, P}{\neg P, \Phi \vdash \Theta}$	LK-NOTR $\frac{P, \Phi \vdash \Theta}{\Phi \vdash \Theta, \neg P}$	LK-FALSEL $\frac{}{\perp \vdash}$	LK-TRUER $\frac{}{\vdash \top}$	

Fig. 5. Validity judgement for LK.

3.1 An embedding of EC types in LK

We show that every EC typing judgment can be translated into a valid LK judgment. The translation is defined by four functions. $\llbracket \tau \rrbracket_{\text{typ}}^{\Sigma}$ maps types to propositions, $\llbracket \Delta \rrbracket_{\text{ctx}}^{\Sigma}$ maps EC result contexts into LK succedents, and $\llbracket \Gamma \rrbracket_{\text{env}}^{\Sigma}$ maps EC environments into LK antecedents. Lastly $\bigvee \Theta$, pronounced “lift Θ ,” internalizes Θ ’s commas as disjunctions.

The type-to-proposition translation is defined as follows.

$$\begin{aligned}
\llbracket A \rrbracket_{\text{typ}}^{\Sigma} &= A_{\text{LK}} \\
\llbracket \tau \rightarrow \Delta \rrbracket_{\text{typ}}^{\Sigma} &= \llbracket \tau \rrbracket_{\text{typ}}^{\Sigma} \supset \bigvee \llbracket \Delta \rrbracket_{\text{ctx}}^{\Sigma} \\
\llbracket \top \rrbracket_{\text{typ}}^{\Sigma} &= \top \\
\llbracket \text{Exn } E \rrbracket_{\text{typ}}^{\Sigma} &= \llbracket \tau \rrbracket_{\text{typ}}^{\Sigma} \quad \text{where } E \sim \tau \in \Sigma
\end{aligned}$$

We assume that there is an atomic LK proposition for each EC type variable. The interesting cases are for arrows and exceptions. The translation of $\tau \rightarrow \Delta$ can’t simply return $\llbracket \tau \rrbracket_{\text{typ}}^{\Sigma} \supset \llbracket \Delta \rrbracket_{\text{ctx}}^{\Sigma}$ —it’s not well-formed. Instead we take the

disjunction of all elements in $\llbracket \Delta \rrbracket_{\text{ctx}}^\Sigma$. This follows the intuition that succedents correspond to lists of disjunctions. The translation of an exception is the translation of the type it carries.

The environment and result context translations are simple.

$$\begin{aligned} \llbracket \cdot \rrbracket_{\text{env}}^\Sigma &= \cdot \\ \llbracket \Gamma, x : \tau \rrbracket_{\text{env}}^\Sigma &= \llbracket \Gamma \rrbracket_{\text{env}}^\Sigma, \llbracket \tau \rrbracket_{\text{typ}}^\Sigma \end{aligned}$$

$$\begin{aligned} \llbracket \cdot \rrbracket_{\text{ctx}}^\Sigma &= \cdot \\ \llbracket \Delta, \varepsilon : \tau \rrbracket_{\text{ctx}}^\Sigma &= \llbracket \Delta \rrbracket_{\text{ctx}}^\Sigma, \llbracket \tau \rrbracket_{\text{typ}}^\Sigma \end{aligned}$$

When lifting LK antecedents we do a little extra work to avoid unnecessary occurrences of \perp ; this leads to an aesthetically cleaner translation.

$$\begin{aligned} \bigvee \cdot &= \perp \\ \bigvee \cdot, P &= P \\ \bigvee \Theta, P &= \bigvee \Theta \vee P \quad \text{where } \Theta \neq \cdot \end{aligned}$$

The \bigvee function internalizes result context commas as \vee s. While this is a syntactically well defined operation, it's not a priori clear that it preserves logical validity. The following lemma shows that it does.

Lemma 8. $\Phi \vdash \Theta$ iff $\Phi \vdash \bigvee \Theta$.

Proof. Proof by induction on the length of Θ .

3.2 Main Results

This section presents our main results. Lemma 9 shows that EC's subtyping relations corresponds to logical consequence in LK. And Theorem 1 shows this for EC typing judgments.

Lemma 9 (Subtyping Lemma). *Suppose \mathcal{D} a subtyping derivation in EC.*

- If $\mathcal{D} :: \Sigma \vdash \Delta_1 <: \Delta_2$ then $\bigvee \llbracket \Delta_1 \rrbracket_{\text{ctx}}^\Sigma \vdash \llbracket \Delta_2 \rrbracket_{\text{ctx}}^\Sigma$.
- If $\mathcal{D} :: \Sigma \vdash \tau_1 <: \tau_2$ then $\llbracket \tau_1 \rrbracket_{\text{typ}}^\Sigma \vdash \llbracket \tau_2 \rrbracket_{\text{typ}}^\Sigma$.
- If $\mathcal{D} :: \Sigma \vdash \diamond$ then $\llbracket \tau_1 \rrbracket_{\text{typ}}^\Sigma \vdash \llbracket \tau_2 \rrbracket_{\text{typ}}^\Sigma$ where $E_1 <: E_2$, $E_1 \sim \tau_1$, $E_2 \sim \tau_2 \in \Sigma$.

Proof. Proof is by structural induction on \mathcal{D} . We will show only the most interesting cases. Begin by examining the cases for $\Sigma \vdash \tau_1 <: \tau_2$.

- Case S-TRANS: Immediate from the induction hypothesis and LK-CUT.
- Case S-ARR: The subtyping derivation ends in

$$\frac{\Sigma \vdash \tau_2 <: \tau_1 \quad \Sigma \vdash \Delta_1 <: \Delta_2}{\Sigma \vdash \tau_1 \rightarrow \Delta_1 <: \tau_2 \rightarrow \Delta_2}$$

We want to show

$$([\tau_1]_{\text{typ}}^\Sigma \supset \bigvee [\Delta_1]_{\text{ctx}}^\Sigma) \vdash ([\tau_2]_{\text{typ}}^\Sigma \supset \bigvee [\Delta_2]_{\text{ctx}}^\Sigma).$$

The induction hypothesis gives $[\tau_2]_{\text{typ}}^\Sigma \vdash [\tau_1]_{\text{typ}}^\Sigma$ and $\bigvee [\Delta_1]_{\text{ctx}}^\Sigma \vdash \bigvee [\Delta_2]_{\text{ctx}}^\Sigma$. Lemma 8 transforms the latter to $\bigvee [\Delta_1]_{\text{ctx}}^\Sigma \vdash \bigvee [\Delta_2]_{\text{ctx}}^\Sigma$. Conclude with an easy LK derivation.

Cases for $\Sigma \vdash \diamond$.

- Case WF-S-SUB: Suppose Σ ends with $E_3 <: E_4$. Conclude with the induction hypothesis, using Lemma 1 if $E_1 \neq E_3$ or $E_2 \neq E_4$.

Cases for $\Sigma \vdash \Delta_1 <: \Delta_2$

- Case S-CONS. The subcontext derivation ends in

$$\frac{\Sigma \vdash \Delta_1 <: \Delta_2 \quad \varepsilon_2 : \tau_2 \in \Delta_2 \quad \Sigma \vdash \varepsilon_1 : \tau_1 <: \varepsilon_2 : \tau_2}{\Sigma \vdash \Delta_1, \varepsilon_1 : \tau_1 <: \Delta_2}$$

First establish $[\tau_1]_{\text{typ}}^\Sigma \vdash [\tau_2]_{\text{typ}}^\Sigma$ and $\bigvee [\Delta_1]_{\text{ctx}}^\Sigma \vdash \bigvee [\Delta_2]_{\text{ctx}}^\Sigma$ via inversion and appeal to the induction hypotheses. Next build this LK derivation:

$$\frac{\bigvee [\Delta_1]_{\text{ctx}}^\Sigma \vdash \bigvee [\Delta_2]_{\text{ctx}}^\Sigma \quad \frac{[\tau_1]_{\text{typ}}^\Sigma \vdash [\tau_2]_{\text{typ}}^\Sigma}{\text{several uses of LK-THINR and LK-INTERR}}}{\bigvee [\Delta_1]_{\text{ctx}}^\Sigma \vee [\tau_1]_{\text{typ}}^\Sigma \vdash \bigvee [\Delta_2]_{\text{ctx}}^\Sigma}$$

When Δ_1 is non-empty, we're done. Otherwise, we conclude with a simple LK derivation.

Theorem 1 (Main Theorem). *Suppose $e : \Sigma; \Gamma \vdash \Delta$. Then $[\Gamma]_{\text{env}}^\Sigma \vdash [\Delta]_{\text{ctx}}^\Sigma$.*

Proof. Proof is by induction on the typing derivation. Cases T-SUB and T-HANDLE use Lemmas 8 and 9.

4 Discussion

On call-by-value evaluation Because EC evaluation follows a call-by-value evaluation strategy the arrow typing rules describe latent effects as suspended. This is advantageous from a programming perspective—it enables precise reasoning about the sequencing of effects. However, this evaluation strategy cannot correspond to cut elimination in LK; functions may contain suspended applications that translate to LK cuts. In systems studied by Nakano [13], Parigot [16], and others, reduction does correspond to cut elimination, but these languages lack EC's natural and precise reasoning for call-by-value exceptions.

EC’s call-by-value reduction makes type-and-effect-inspired analysis far more appealing than a monadic effect system. Consider the following program.

$$(\lambda x: \top. \text{top}) (\text{raise Fail "disaster!"})$$

EC types this with $\text{ret}: \top, \text{exn Fail}: \text{string}$. This is safe and (reasonably) precise for a strict language; evaluation will raise the reported exception.

Call-by-name throws away the `raise` expression and no exception is thrown. Haskell’s monadic effect discipline accounts for this in a precise way. The `raise` expression can be assigned type $\text{M } b$, where M is the exception monad, and the function gets type $\text{forall } a, a \rightarrow \top$. (We have to squint a bit because EC lacks polymorphism and Haskell lacks subtyping.) The entire expression has type \top —a sound result given call-by-name evaluation.

However, type \top is unsound for call-by-value evaluation. What goes wrong is that unused function parameters are relevant to evaluation, but not to the monadic effects system. It appears possible to cook the monadic analysis so that it’s sound for call-by-value. One way to do this is by imposing side conditions on the abstraction or application typing rules. However such an invasive change is unsatisfying. It negates one of the key advantages of monad-based effect analysis in Haskell: independence of effect monads from the rest of the type system.

Open questions There are several more questions we can ask about EC itself and its relation to LK: What is the right way to add polymorphism to EC? Is EC strongly normalizing? How can EC extensions deal with world effects like memory access?

What is the logical content of EC proofs? EC is weaker than classical logic, perhaps even strictly weaker. It appears obvious that EC is as powerful as intuitionistic logic. Can the shallow embedding in this paper be extended to an isomorphism? Should such an isomorphism be with full LK, an intuitionistic restriction of LK [4, 9, 17], or something else?

5 Related Work

Logics with control Griffin [5] first established that proofs-as-program correspondence between lambda calculus and intuitionistic logic could be extended to control effects and classical logic. Subsequently, several authors have investigated the relation between control effects, classical logic, and sequent calculi.

Nakano [14] observed a correlation between exception analysis and LK, a classical logic. His system $L_{c/t}$ uses lexically scoped variables, or “tags,” as identifiers for exception tracking. However the tags-as-variables style requires many administrative tag abstractions and instantiations to avoid problems stemming from variable confusion. Additionally, Nakano’s static semantics fails to directly account for latent effects. His notion of “compatible contexts” hints at subtyping, but does explore it substantially.

Ong and Stewart [15] extend Parigot’s $\lambda\mu$ calculus in their μPCF system and propose an encoding of ML-like exceptions similar to EC’s. Unlike EC, μPCF identifies exception names with covariables and lacks exception subtyping.

De Groote [2] examines an exception analysis based on classical natural deduction. Exceptions that carry type τ have type $\tau \rightarrow \perp$, and classical reasoning principles may be encoded in his system. The system uses lexically scoped variables as exception names and requires a nonstandard reduction scheme to prevent these names from escaping their scope. The analysis only types programs that do not raise uncaught exceptions and does not account for latent effects.

Several authors have encoded exceptions and exception analysis with control operators. As illustrated by Lillibridge [11], the exact relation between these features is highly sensitive to the surrounding language. We choose to define exceptions directly as definition by elaboration fails to provide programmers with an intuitive language specification.

Intuitionistic sequent calculi This paper embeds EC into classical LK, making extensive use of multiple-conclusion succedents. However, as mentioned in Section 4, intuitionistic variants might give a tight isomorphism with EC.

There are several well-studied multiple-conclusion intuitionistic logics which might serve as translation targets. For example, Kleene [9] cuts-down LK by placing cardinality restrictions on the succedents in LK-NOTR and LK-THINR, and system LJ' [17] does so for LK-NOTR and LK-IMPR. More recently de Paiva and Pereira [3] show how to avoid cardinality constraints by labeling conclusions with unique indices describing relevant hypotheses.

The type-and-effect discipline for exceptions Type-and-effects systems track side effects in (usually) call-by-value languages [18, 7] Typically this uses a judgment of the form $\Gamma \vdash e : \tau ! \ell$. Label ℓ describes the set of effects which may be triggered when evaluating e . Additionally, arrow types are annotated with latent effects.

Leroy and Pessaux [10] use an effects system to detect uncaught exceptions in ML programs. (Research by Guzmán and Suárez [6] foreshadows this work.) This analysis differs from ours as it considers ML-style exceptions with no subtyping, and attempts to track values carried by exceptions. Blume et al. [1] use a refinement of Leroy's techniques to analyze exceptions in a language with extensible cases, showing similarity between the exception and case mechanisms. Kennedy [8], working with a CPS-style intermediate language, and shows type assignment for the exception continuation corresponds to an effect analysis.

6 Conclusion

The paper introduces EC, a language with a sound, typed-based analysis for Java-style exceptions. We show that EC's types and static judgments can be embedded in Gentzen's LK. This clarifies the relation between effects analysis, an important enabling technology for compilation and verification, and classical sequent calculi, an interesting and well-understand family of logics.

Acknowledgments I thank Steve Zdancewic for his helpful feedback on drafts of this paper, and Hugo Herbelin and Guillaume Munch-Maccagnoni for their discussions about classical logic.

Bibliography

- [1] M. Blume, U. Acar, and W. Chae. Exception handlers as extensible cases. In *Proc. APLAS '08*, 2008.
- [2] P. de Groote. A simple calculus of exception handling. In *Proc. TLCA '95*. Springer-Verlag, 1995.
- [3] Valeria de Paiva and Luiz Carlos Pereira. A short note on intuitionistic propositional logic with multiple conclusions. *Manuscrito—Rev. Int. Fil., Campinas*, 28(2):317–329, July 2005.
- [4] G. Gentzen. Investigations into logical deduction. In M. E. Szabo, editor, *The Collected Papers of Gerhard Gentzen*. North-Holland, 1935. Collection 1969.
- [5] T. Griffin. A formulae-as-type notion of control. In *Proc. POPL '90*, 1990.
- [6] J. C. Guzmán and A. Suárez. A type system for exceptions. In *Proc. ML Workshop*, 1994.
- [7] P. Jouvelot and D. Gifford. Reasoning about continuations with control effects. *SIGPLAN Not.*, 24(7):218–226, 1989.
- [8] A. Kennedy. Compiling with continuations, continued. In *Proc. ICFP '07*, 2007.
- [9] Stephen Kleene. *Introduction to Metamathematics*. D. Van Nostrand Company, Inc., New York, 1952.
- [10] X. Leroy and F. Pessaux. Type-based analysis of uncaught exceptions. *ACM Trans. Program. Lang. Syst.*, 22(2):340–377, 2000.
- [11] M. Lillibridge. Exceptions are strictly more powerful than call/cc. Technical Report CMU-CS-95-178, Carnegie Mellon University, July 1995.
- [12] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT, 1997.
- [13] H. Nakano. A constructive formalization of the catch and throw mechanism. In *LICS '92*, pages 82–89, 1992.
- [14] H. Nakano. The non-deterministic catch and throw mechanism and its subject reduction property. In *Logic, Language and Computation, Festschrift in Honor of Satoru Takasu*, pages 61–72, 1994.
- [15] L. Ong and C. Stewart. A Curry-Howard foundation for functional computation with control. In *POPL '97*, pages 215–227, 1997.
- [16] Michel Parigot. $\lambda\mu$ -calculus: An algorithmic interpretation of classical natural deduction. In *Proc. LPAR*, pages 190–201. 1992.
- [17] Gaisi Takeuti. *Proof Theory*. North-Holland, Amsterdam, 1975.
- [18] J.-P. Talpin and P. Jouvelot. The type and effect discipline. In *Proc. LICS '92*, 1992.
- [19] G. Washburn. Cause and effect: type systems for effects and dependencies. Technical Report MS-CIS-05-05, University of Pennsylvania, July 2005.