

Exceptional Syntax

Nick Benton

Andrew Kennedy

Microsoft Research Cambridge, UK
 {nick,akenn}@microsoft.com

Abstract

From the points of view of programming pragmatics, rewriting and operational semantics, the syntactic construct used for exception handling in ML-like programming languages, and in much theoretical work on exceptions, has subtly undesirable features. We propose and discuss a more well-behaved construct.

1 Introduction

Most modern programming languages (e.g. SML, Java, Ada, C++, Common LISP) include *exceptions* to provide a structured, but non-local, way of signalling and recovering from error conditions. Programmers often also use exceptions as convenient, and sometimes more efficient, way of varying control flow in code which has nothing to do with what most people would consider error-handling (for example, the parser combinators and unification code in [16]).

The basic idea of exceptions is simple and familiar: the evaluation of an expression may, instead of completing normally by returning a value or diverging, terminate abnormally by *raising* a named exception. The evaluation of any expression may be wrapped in an *exception handler*, which provides an alternative expression to be evaluated in the case that the wrapped expression raises a particular exception. The way in which a raised exception unwinds the evaluation stack until the closest matching handler is found is syntactically implicit, so the handler may be dynamically far from the point at which the exception is raised without the intervening calls having explicitly to test for, and propagate, an error value.

There are many subtle differences between exception mechanisms in different programming languages, but for the purposes of this paper we shall take a simplified version of the constructs provided in Standard ML [14] as paradigmatic. To the usual simply-typed lambda calculus we add a set \mathbb{E} of exception names, a new base type `exn`, and new constructs with the typing rules

$$\frac{}{\Gamma \vdash E : \text{exn}} \quad E \in \mathbb{E} \qquad \frac{\Gamma \vdash M : \text{exn}}{\Gamma \vdash \text{raise } M : A}$$

$$\frac{\Gamma \vdash M : A \quad \{\Gamma \vdash N_i : A\}_{i=1\dots n}}{\Gamma \vdash M \text{ handle } E_1 \Rightarrow N_1 \mid \dots \mid E_n \Rightarrow N_n : A} \quad \{E_i\}_{i=1\dots n} \subseteq \mathbb{E}$$

where, in the last rule, the E_i are required to be distinct. We take as basic a form of *handle* in which multiple handlers may cover the evaluation of a single expression, as this

is strictly more expressive than the simpler form in which only one named exception may be caught at once. We will sometimes use an abbreviated notation, using H to range over finite sets $\{E_i \Rightarrow N_i\}$ of handlers, and writing $E \in H$ for $\exists N. E \Rightarrow N \in H$ and $H(E)$ for the (unique) N such that $E \Rightarrow N \in H$ if that exists. Write $\Gamma \vdash H : A$ for $H = \{E_i \Rightarrow N_i\}$ and $\forall i. \Gamma \vdash N_i : A$.

One way of explaining the intended behavior of these constructs is to give a big-step operational semantics in which there are two (mutually inductive) forms of judgement: $M \Downarrow V$ means that the closed expression M evaluates to the value V , whereas $M \Uparrow E$ means that the expression M raises the exception E . The rules for deriving these judgements comprise the usual evaluation rules for a call by value lambda calculus¹ together with at least the following:

$$\frac{}{E \Downarrow E} \quad E \in \mathbb{E} \qquad \frac{M \Downarrow E}{(\text{raise } M) \Uparrow E}$$

$$\frac{M \Downarrow V}{(M \text{ handle } H) \Downarrow V} \qquad \frac{M \Uparrow E}{(M \text{ handle } H) \Uparrow E} \quad E \notin H$$

$$\frac{M \Uparrow E \quad N \Downarrow V}{(M \text{ handle } H) \Downarrow V} \quad H(E) = N$$

$$\frac{M \Uparrow E \quad N \Uparrow E'}{(M \text{ handle } H) \Uparrow E'} \quad H(E) = N$$

$$\frac{M \Uparrow E}{(M N) \Uparrow E} \qquad \frac{M \Downarrow \lambda x. M' \quad N \Uparrow E}{(M N) \Uparrow E}$$

There will be further rules, similar to the last two above, which express the way in which thrown exceptions propagate through whatever other constructs we choose to add to our language.

We became aware of essentially the same shortcoming of the *handle* construct in three different ways whilst working on our Standard ML compiler, MLj [3]. Firstly, when coding in SML to implement both the compiler itself and its libraries, we occasionally came across situations in which exception-handling behaviour could only be expressed clumsily. Secondly, when performing rewriting on the compiler

¹We restrict attention to call by value, as the naïve addition of exceptions to a language with call by name semantics wrecks the equational theory to the extent that the resulting language is essentially unusable. There is, however, an ingenious proposal for adding exceptions to Haskell, which sidesteps some of the problems; see [12] for details.

intermediate language, we found that some rewrites were inexpressible if the intermediate language contained the usual exception handling construct. Thirdly, when formalising the intermediate language in order to prove some theorems about the validity of optimising transformations [2], we found that the alternative syntax we had chosen (for the previous reason) allowed a neat and tractable presentation of the operational semantics in terms of a structurally inductive termination predicate, which would not otherwise have been possible.

2 The New Construct

Since the fix for the problems we observed is actually rather simple, and to avoid building unnecessary suspense in the reader, we will reverse the usual order of presentation by giving our solution straight away and then going into the more technical explanations of the problems it solves.

We replace the ML-style *handle* construct with a new one, which builds in a continuation to be applied only in the case that no exception is raised:

$$\text{try } x \leftarrow M \text{ catch } E_1 \Rightarrow N_1 \mid \dots \mid E_n \Rightarrow N_n \text{ in } P$$

This first evaluates M and, if it returns a value, binds that to x and evaluates P . If M raises the exception E_i , however, N_i is evaluated instead. If M raises an exception distinct from all the E_i , then so does the whole expression.

More formally, Figure 1 presents the (simple, non-computational) typing rule for *try* along with its natural semantics rules. Note that we find it convenient to allow empty handlers in this construct and that the type of the expressions N_i in a handler is the same as that of the continuation P , *not* the same as that of the expression M being covered, as is the case with the traditional *handle*.

3 So What Was Wrong With *handle*?

We now describe the problem with the traditional *handle* construct in each of the three contexts in which we observed it. To avoid dragging in too much extraneous material concerning, for example, our compiler intermediate language, we will often gloss over the non-exceptional details of the various languages mentioned: this should not (we hope!) obscure our main point.

3.1 The Programming Problem

Suppose one has a library of ML functions to open, read and close files, all of which raise the `Io` exception if something goes wrong. The problem is to write a function which runs down a list of filenames, performing some processing on the contents of the first one that can be opened successfully. Ignoring the possibilities that the list is empty or that *none* of the files exists/is openable, one's first thought might be that the following will suffice:

```
fun lookup (name :: names) =
  let val f = openIn name (* try to open the file *)
      in readIt f          (* do some processing *)
      end handle Io =>
        lookup names      (* if failed, move on *)
```

However, this doesn't quite do what we want, as the function `readIt` might also raise the `Io` exception: when that

happens then we want the exception to be passed up to the caller of `lookup`, but the above code will handle the exception and move on to the next name in the list irrespective of whether the error occurred in `openIn` or `readIt`.

There are, of course, various straightforward ways of programming around this problem. For example, we might use a datatype:

```
datatype 'a option = NONE | SOME of 'a

fun lookup1 (name :: names) =
  case SOME(openIn name) handle Io => NONE
  of NONE => lookup1 names
   | SOME(f) => readIt f
```

Or use abstraction to delay the call to `readIt` so that the handler doesn't cover it:

```
fun lookup2 (name :: names) =
  (let val f = openIn name
   in fn () => readIt f
   end handle Io => fn () => lookup2 names) ()
```

Or use another exception:

```
fun lookup3 (name :: names) =
  let val f = openIn name handle Io => raise NotFound
      in readIt f
      end handle NotFound => lookup3 names
```

But none of these seems entirely satisfactory as they all introduce a new value (sum, closure or exception) only to eliminate it straight away – it's just there to express some control flow which the `handle` construct is too weak to express directly.

The fix: Programming with *try*

The *try-catch-in* syntax nicely solves our programming problem:

```
fun lookup (name :: names) =
  try f = openIn name
  catch Io => lookup names
  in readIt f
  end
```

and also generalises both *let* and *handle*:

$$\begin{aligned} \text{let } x \leftarrow M \text{ in } N &= \text{try } x \leftarrow M \text{ catch } \{\} \text{ in } N \\ M \text{ handle } H &= \text{try } x \leftarrow M \text{ catch } H \text{ in } x \end{aligned}$$

3.2 The Transformation Problem

Like many compilers for functional languages, MLj performs fairly extensive rewriting in order to optimise programs. The design of MLj's intermediate language, MIL, and its rewrites is motivated by a somewhat informal belief in 'taking the proof theory seriously'. One instance of this prejudice is that the compiler transforms programs into a 'cc-normal form', in which all of the *commuting conversions* have been applied.

In natural deduction presentations of logics (and hence, via the Curry-Howard correspondence, in typed lambda calculi), commuting conversions occur when logical rules (usually eliminations) have what Girard calls a 'parasitic formula' [9], a typical case being that of the sum. The elimination rule for sums is

$$\frac{\Gamma \vdash M : A + B \quad \Gamma, x_1 : A \vdash N_1 : C \quad \Gamma, x_2 : B \vdash N_2 : C}{\Gamma \vdash \text{case } M \text{ of } \text{in}_1 x_1. N_1 \mid \text{in}_2 x_2. N_2 : C}$$

$$\begin{array}{c}
\frac{\Gamma \vdash M : A \quad \Gamma, x : A \vdash P : B \quad \{\Gamma \vdash N_i : B\}_{i=1..n \geq 0} \quad \{E_i\}_{i=1..n} \subseteq \mathbb{E}}{\text{try } x \leftarrow M \text{ catch } E_1 \Rightarrow N_1 \mid \dots \mid E_n \Rightarrow N_n \text{ in } P : B} \\
\\
\frac{M \Downarrow V \quad P[V/x] \Downarrow V'}{\text{try } x \leftarrow M \text{ catch } H \text{ in } P \Downarrow V'} \qquad \frac{M \Downarrow V \quad P[V/x] \Uparrow E}{\text{try } x \leftarrow M \text{ catch } H \text{ in } P \Uparrow E} \\
\\
\frac{M \Uparrow E \quad N \Downarrow V}{\text{try } x \leftarrow M \text{ catch } H \text{ in } P \Downarrow V} H(E) = N \qquad \frac{M \Uparrow E \quad N \Uparrow E'}{\text{try } x \leftarrow M \text{ catch } H \text{ in } P \Uparrow E'} H(E) = N \\
\\
\frac{M \Uparrow E}{\text{try } x \leftarrow M \text{ catch } H \text{ in } P \Uparrow E} E \notin H
\end{array}$$

Figure 1: Typing rule and natural semantics for *try*

in which the formula/type C has no connection with that being eliminated. The presence of such rules introduces undesirable distinctions between proofs and also, for example, causes the subformula property of normal deductions to fail. These problems are addressed by adding commuting conversions to the more familiar β and η rules. Commuting conversions typically have the general form

$$\begin{array}{c}
\begin{array}{c} \vdots \\ A \end{array} \quad \begin{array}{c} \vdots \\ C \end{array} \quad \dots \quad \begin{array}{c} \vdots \\ C \end{array} \quad \mathcal{E}_1 \\
\hline
C \quad \dots \quad \mathcal{E}_2 \\
\hline
D
\end{array} \\
\text{rewrites to} \\
\begin{array}{c}
\begin{array}{c} \vdots \\ C \end{array} \quad \dots \quad \begin{array}{c} \vdots \\ C \end{array} \quad \dots \quad \mathcal{E}_2 \\
\hline
A \quad \dots \quad \mathcal{E}_1 \\
\hline
D
\end{array}
\end{array}$$

where \mathcal{E}_1 is the ‘bad’ elimination rule for the top-level connective in A , with parasitic formula C (which may occur one or more times as a premiss, according to the connective being eliminated), and \mathcal{E}_2 is the elimination rule for the top-level connective in C . For example, if \mathcal{E}_1 is \vee -elimination and \mathcal{E}_2 is \rightarrow -elimination, we get the following commuting conversion on terms:²

$$\begin{array}{c}
(\text{case } M \text{ of } \text{in}_1 x_1. N_1 \mid \text{in}_2 x_2. N_2) P \\
\rightsquigarrow \\
\text{case } M \text{ of } \text{in}_1 x_1. (N_1 P) \mid \text{in}_2 x_2. (N_2 P).
\end{array}$$

(Here and elsewhere, we adopt the ‘variable convention’: sufficient α -conversion to avoid unwanted variable capture is assumed. In the above, this implies that neither x_1 nor x_2 is free in P .) Commuting conversions often enable further

²Applied naively, of course, the duplication of terms in conversions like this one could lead to an unacceptable blowup in code size. In MLj, we avoid this by selective use of a special abstraction construct which compiles to a block of code accessed by jumps, but we will not discuss this subtlety further here.

reductions which would otherwise be blocked, as in

$$\begin{array}{c}
(\text{case } M \text{ of } \text{in}_1 x_1. \lambda y. y + x_1 \mid \text{in}_2 x_2. \lambda y. y) 2 \\
\rightsquigarrow \text{case } M \text{ of } \text{in}_1 x_1. ((\lambda y. y + x_1) 2) \mid \text{in}_2 x_2. ((\lambda y. y) 2) \\
\rightsquigarrow \text{case } M \text{ of } \text{in}_1 x_1. (2 + x_1) \mid \text{in}_2 x_2. 2
\end{array}$$

and we also find generating code from cc-normal forms considerably more straightforward than for arbitrary terms. Other compilers perform similar rewrites (for example, the *case-of-case* and *let-floating* transformations in [13]), though we are unusually dogmatic in recognising them as instances of a common pattern and performing *all* of them.

Interestingly, cc-normal form for our intermediate language, which is based on Moggi’s computational metalanguage [15], turns out to be almost the same thing as Sabry and Felleisen’s *A-normal form* [20, 8], which was derived from an analysis of CPS-based compilation. A nice discussion of the connection between CPS and Moggi’s metalanguage may be found in [11].

For most of the type constructors of our intermediate language, MIL, we have well-behaved introduction and elimination rules for which it is clear how to derive the commuting conversions. For the exception-related constructs, the situation is messier (since part of the point of exceptions is that they are not explicitly visible in source-language types) but it is nevertheless obvious that there are some cc-like rewrites which we would like to perform. For example

$$(M \text{ handle } E \Rightarrow N) P$$

looks as though it should convert to something like

$$(M P) \text{ handle } E \Rightarrow (N P)$$

so that if, for example, N is a λ -abstraction, we get to perform a compile-time β -reduction. But this transformation is not generally sound if either P or the application of the value of M to the value of P might raise the exception E . Furthermore, there isn’t a correct transformation which we can use instead. It should be remarked at this point that the limited expressibility of an intermediate language based on a λ -calculus with *handle* is not shared by lower-level target languages. Using Java bytecodes, for example, a code sequence corresponding to a correct version of the above transformation is easily written:

```

11: Code to evaluate M
12: Code to evaluate P
    invokevirtual <Method resulttype apply(argtype)>
13: Code for rest of computation

14: pop \\ throw away the actual exception object
    Code to evaluate N
    Code to evaluate P
    invokevirtual <Method resulttype apply(argtype)>
    jmp 13

```

```

Exception table:
  from  to  target type
  11    12  14    <Class E>

```

and the same is true of target languages in which exception handlers are explicitly pushed onto and popped from a stack.

In fact, because of the separation of computations from values in MIL, we would have to express the first term above as

$$\text{let } f \Leftarrow (M \text{ handle } E \Rightarrow N) \text{ in let } v \Leftarrow P \text{ in } (f v)$$

but the essential point remains unchanged: there's simply no correct way to write the transformation which we feel we should be able to perform.

Of course, one could simply accept the inexpressibility of such transformations and generate slightly lower quality code. Alternatively, one can observe that the commuting conversions are not in themselves generally optimisations; they are reorganisations of the code which enable more computationally significant β redexes to be exposed. Hence the same optimisations might well be obtained by using non-local rewrites which look for larger patterns in the term. This would, however, significantly increase the complexity of the rewriting function and, we believe, would make it less efficient (despite the fact that the non-local steps would combine the effect of more than one local rewrite).

The fix: Rewriting with *try*

The *try-catch-in* syntax comes with unsurprising β -like reductions, similar to those for *handle* and *let*

$$\begin{aligned} \text{try } x \Leftarrow \text{raise } E \text{ catch } H \text{ in } P &\rightsquigarrow N && (N = H(E)) \\ \text{try } x \Leftarrow \text{raise } E \text{ catch } H \text{ in } P &\rightsquigarrow \text{raise } E && (N \notin H) \\ \text{try } x \Leftarrow V \text{ catch } H \text{ in } P &\rightsquigarrow P[V/x] && (V \text{ a value}) \end{aligned}$$

but, unlike *handle*, also has well-behaved commuting conversions, which allow us to express useful compiler transformations. We present in Figure 2 a general list of conversions for *try-catch-in* against itself and the eliminations for sums, products and functions. Although these look complex, it should be noted that in a language like MIL (which separates values from computations at both the type and term levels) or Pitts's language from [17] (which has term-level restrictions on the places where non-values may occur), most of these cases either do not occur or only occur in a simplified form. In MIL, for example, only *try-case* and *try-try* are well typed, because projection, application and *case* can only be applied to values, whereas a *try* is always a computation. Furthermore, the restriction that *M* in the *try-case* rewrite be a value simplifies it to

(*try-case*):

$$\begin{aligned} \text{try } x \Leftarrow (\text{case } V \text{ of } \text{in}_1 y_1. N_1 \mid \text{in}_2 y_2. N_2) \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \\ \rightsquigarrow \text{case } V \text{ of } \text{in}_1 y_1. \text{try } x \Leftarrow N_1 \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \mid \\ \text{in}_2 y_2. \text{try } x \Leftarrow N_2 \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \end{aligned}$$

The *try-catch-in* construct is the one which we use in MIL, and the MLj compiler actually does perform the *try-try* and *try-case*' rewrites.

As an interesting example of MIL rewriting, showing the *try* construct working with our monadic effect analysis [2], consider the following ML function for summing all the elements of an array:

```

fun sumarray a =
  let fun s(n,sofar) = let val v = Array.sub(a,n)
                      in s(n+1, sofar+v)
                      end handle Subscript => sofar

      in s(0,0)
      end

```

Because the SML source language doesn't have *try*, the programmer has made the handler cover both the array access and the recursive call to the inner function *s*. But this would prevent a naïve compiler from recognising that call as tail-recursive. In MLj, the intermediate code for *s* looks like (in MLish, rather than MIL, syntax):

```

fun s(n,sofar) =
  try val x = try val v = Array.sub(a,n)
              catch {}
              in s(n+1, sofar+v)
              end
  catch Subscript => sofar
  in x
  end

```

The *try-try* rewrite turns this into

```

fun s(n,sofar) = try val v = Array.sub(a,n)
                  catch Subscript => sofar
                  in try val x = s(n+1, sofar+v)
                    catch Subscript => sofar
                    in x
                    end
                  end

```

(The two identical handlers are actually abstracted as a shared local block.) The effect analysis detects that the recursive call to *s* cannot, in fact, ever throw the *Subscript* exception, so the function is rewritten again to

```

fun s(n,sofar) = try val v = Array.sub(a,n)
                  catch Subscript => sofar
                  in s(n+1, sofar+v)
                  end

```

which *is* tail recursive, and so gets compiled as a loop in the final code for *sumarray*.

3.3 The Semantics Problem

There are several different styles in which one can specify the operational semantics of ML-like languages. We have already seen (in Section 1) a big-step, natural semantics presentation, but this is not always the most convenient formulation with which to work when proving results about observational equivalences. A popular alternative is to use a small-step semantics presented using Felleisen's notion of *evaluation context* [7]. In this style, one first defines axioms for the primitive transitions $R \rightarrow M$, saying that redex *R* reduces to term *M*, and then gives an inductive definition of

$\begin{aligned} & \pi_j(\text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow N_i\} \text{ in } P) \\ & \rightsquigarrow \text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow \pi_j(N_i)\} \text{ in } \pi_j(P) \end{aligned}$	<i>proj-try</i>
$\begin{aligned} & (\text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow N_i\} \text{ in } P) Q \\ & \rightsquigarrow \text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow (N_i Q)\} \text{ in } (P Q) \end{aligned}$	<i>app-try</i>
$\begin{aligned} & \text{case } (\text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow N_i\} \text{ in } P) \text{ of } \text{in}_1 y_1. Q_1 \mid \text{in}_2 y_2. Q_2 \\ & \rightsquigarrow \text{try } x \leftarrow M \text{ catch } \{E_i \Rightarrow \text{case } N_i \text{ of } \text{in}_1 y_1. Q_1 \mid \text{in}_2 y_2. Q_2\} \\ & \quad \text{in case } P \text{ of } \text{in}_1 y_1. Q_1 \mid \text{in}_2 y_2. Q_2 \end{aligned}$	<i>case-try</i>
$\begin{aligned} & \text{try } x \leftarrow (\text{case } M \text{ of } \text{in}_1 y_1. N_1 \mid \text{in}_2 y_2. N_2) \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \\ & \rightsquigarrow \text{try } z \leftarrow M \text{ catch } \{E_i \Rightarrow P_i\} \text{ in} \\ & \quad \text{case } z \text{ of } \text{in}_1 y_1. \text{try } x \leftarrow N_1 \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \mid \\ & \quad \quad \text{in}_2 y_2. \text{try } x \leftarrow N_2 \text{ catch } \{E_i \Rightarrow P_i\} \text{ in } Q \end{aligned}$	<i>try-case</i>
$\begin{aligned} & \text{try } x \leftarrow (\text{try } y \leftarrow M \text{ catch } \{E_i \Rightarrow N_i\}_{i \in I} \text{ in } P) \\ & \quad \text{catch } \{E'_j \Rightarrow N'_j\}_{j \in J} \text{ in } Q \\ & \rightsquigarrow \text{try } y \leftarrow M \text{ catch } \{E_i \Rightarrow \text{try } x \leftarrow N_i \text{ catch } \{E'_j \Rightarrow N'_j\}_{j \in J} \text{ in } Q\}_{i \in I} \\ & \quad \cup \{E'_j \Rightarrow N'_j\}_{E'_j \notin \{E_i\}_{i \in I}} \\ & \quad \text{in } \text{try } x \leftarrow P \text{ catch } \{E'_j \Rightarrow N'_j\}_{j \in J} \text{ in } Q \end{aligned}$	<i>try-try</i>

Figure 2: Conversions

evaluation contexts as terms $E[\cdot]$ containing a single ‘hole’ in the place where the next reduction will take place. A simple lemma that every non-value is uniquely of the form $E[\bar{R}]$ then allows the one-step transition relation to be defined as $E[\bar{R}] \rightarrow E[\bar{M}]$ for every evaluation context $E[\cdot]$ and primitive transition $R \rightarrow M$ (and the evaluation relation to be defined in terms of the reflexive transitive closure of the transition relation). Wright and Felleisen [21] give an evaluation context semantics for ML with exceptions which uses a second kind of context for propagating exceptions.

Pitts has argued [17] that for reasoning about contextual equivalences it is convenient to reify the notion of evaluation context and give a small-step operational semantics in which a configuration is a pair of a term and an explicit context (continuation). The advantages of this approach include the fact that the right-hand sides of transitions are all defined by structural induction over the left-hand side and that there is a Galois connection between relations on terms and relations on contexts which has proved useful in reasoning about, for example, equivalence of polymorphic functions. This style of presentation is also particularly natural if the language includes first-class continuations, in the style of Scheme or SML/NJ (see, for example, [10]).

Pitts formalises contexts by introducing new syntactic categories for defining *continuation stacks*: a configuration looks like

$$\langle (x_1).N_1 \circ \dots \circ (x_n).N_n, M \rangle$$

where M is the term being evaluated (in a λ -calculus with a strict **let** construct and a restriction that only values and variables may occur in eliminations) and $(x_1).N_1 \circ \dots \circ (x_n).N_n$ is a sequence of (closed) abstractions representing the context in which the evaluation takes place. The rules defining the transition relation include

$$\begin{aligned} \langle K \circ (x).N, V \rangle & \rightarrow \langle K, N[V/x] \rangle \\ \langle K, \text{let } x \leftarrow M \text{ in } N \rangle & \rightarrow \langle K \circ (x).N, M \rangle \\ \langle K, (\lambda x.M) V \rangle & \rightarrow \langle K, M[V/x] \rangle \end{aligned}$$

which, it should be apparent, amounts to defining a kind of abstract machine.³ This style of semantics has been applied in [18, 19, 5], and the relational operators it induces are further discussed in [1]. Coincidentally, the current implementation of MLj uses essentially the same representation internally for efficient rewriting of terms in context.

Pitts gives the relationship between the stack-based semantics and a natural semantics using the following lemma: For all appropriately-typed, closed K, M and V

$$\langle K, M \rangle \rightarrow^* \langle \cdot, V \rangle \iff K @ M \Downarrow V$$

where \cdot is the empty continuation stack and the ‘unwinding’ operator $@$ is defined by

$$\begin{aligned} \cdot @ M & = M \\ (K \circ (x).N) @ M & = K @ (\text{let } x \leftarrow M \text{ in } N). \end{aligned}$$

Note how the place where the action (reduction) happens is at the root of the syntax tree of a stack configuration but buried deep in that of its unwinding, as

$$\begin{aligned} & ((x_1).N_1 \circ \dots \circ (x_n).N_n) @ M \\ & \quad \text{let } x_1 \leftarrow (\\ & \quad \quad \text{let } x_2 \leftarrow \\ & \quad \quad \quad \dots (\text{let } x_n \leftarrow M \text{ in } N) \dots \\ & \quad \quad \quad \text{in } N_2) \\ & \quad \text{in } N_1 \end{aligned}$$

It is straightforward to extend Pitts’s semantics to a language with exceptions: one simply allows (closed) handlers

³Actually, since Pitts is interested in which configurations lead to termination, for reasoning about contextual equivalence, the one-step transitions are implicit in inference rules defining the termination predicate \searrow directly, such as

$$\frac{\langle K, N[V/x] \rangle \searrow}{\langle K \circ (x).N, V \rangle \searrow}$$

H (which we previously introduced as an abbreviation for part of the syntax of the *handle* construct and are now making slightly more first-class) to appear as a new kind of element in continuation stacks, with the new transitions

$$\begin{aligned} \langle K \circ H, V \rangle &\rightarrow \langle K, V \rangle \\ \langle K \circ H, \text{raise } E \rangle &\rightarrow \langle K, N \rangle \quad \text{if } H(E) = N \\ \langle K \circ H, \text{raise } E \rangle &\rightarrow \langle K, \text{raise } E \rangle \quad \text{if } E \notin H \\ \langle K \circ (x).N, \text{raise } E \rangle &\rightarrow \langle K, \text{raise } E \rangle \\ \langle K, M \text{ handle } H \rangle &\rightarrow \langle K \circ H, M \rangle \end{aligned}$$

The connection with the natural semantics extends to

$$\langle K, M \rangle \rightarrow^* \langle \cdot, \text{raise } E \rangle \iff K @ M \uparrow E$$

where the definition of $@$ is extended by

$$(K \circ H) @ M = K @ (M \text{ handle } H)$$

and this is the formulation we initially used when working on the equational theory of MIL. However, there is a certain amount of clutter involved in using stacks (extra syntax, type rules, etc.), and we noticed that if one's syntax is sufficiently well-behaved then it is possible to obtain an equally tractable presentation of the transition relation just using terms of the original language. For Pitts's language *without* exceptions, the idea is to axiomatise directly transitions between terms of the form $\text{let } x \Leftarrow M \text{ in } N$ by using commuting conversion transitions to 'bubble up' the next redex in M until it is at the top (and its surrounding context within M has been pushed into N). For example:

$$\begin{aligned} \text{let } x \Leftarrow V \text{ in } N &\rightarrow \text{let } y \Leftarrow N[V/x] \text{ in } y \quad (N \neq x) \\ \text{let } x \Leftarrow (\text{let } y \Leftarrow M \text{ in } N) \text{ in } P &\rightarrow \text{let } y \Leftarrow M \text{ in let } x \Leftarrow N \text{ in } P \\ \text{let } x \Leftarrow (\lambda y.M) V \text{ in } N &\rightarrow \text{let } x \Leftarrow M[V/y] \text{ in } N \end{aligned}$$

Using this style of presentation, the relationship between the big-step and small-step semantics becomes

$$(\text{let } x \Leftarrow M \text{ in } x) \rightarrow^* (\text{let } x \Leftarrow V \text{ in } x) \iff M \Downarrow V.$$

Intuitively, the stack-free transition relation is defined directly on a variant of Pitts's 'unwound' terms, in which the *lets* associate the other way around from the original definition:

$$(K \circ (x).N) @ M = \text{let } x \Leftarrow M \text{ in } (K @ N).$$

The equivalence of the two definitions of $@$ depends on the validity of the associativity of *let* (which, if one separates computations from values in the type system as in Moggi's computational metalanguage, is a commuting conversion in the associated logic [4]).

However, if we add exceptions and the *handle* construct, the definition of the stack-free transition relation fails to extend. Once again, the problem is the lack of commuting conversions which would allow an exception handler to be pushed into a surrounding context so that the evaluation of the expression covered by the handler 'bubbles' to the top. More concretely, consider the following putative transition:

$$\text{let } x \Leftarrow (M \text{ handle } E \Rightarrow N) \text{ in } P \rightarrow \Gamma$$

We'd like to put something on the right-hand side in which the evaluation of M is at the top of the syntax tree, but there's no rewrite to anything of the form $\text{let } x \Leftarrow M \text{ in } \dots$. Nor can we extend the collection of top-level forms to include *handle* as well as *let* constructs: there's no rewrite to something of the form $M \text{ handle } E \Rightarrow \dots$ either.

The fix: Operational semantics with *try*

If our language includes *try-catch-in*, then there is no difficulty in giving a stack-free presentation of a structurally inductive transition semantics. Figure 3 presents transitions between terms of the form $\text{try } x \Leftarrow M \text{ catch } H \text{ in } P$ (recall that *try-catch-in* generalises *let*). The syntax ($H \text{ catch } H' \text{ in } x.Q$) is an abbreviation for the covering of one handler by the other handler and continuation used in the *try-try* conversion.

The connection between the transition semantics and the big-step semantics is then expressed by

$$\begin{aligned} M \Downarrow V &\iff \text{try } x \Leftarrow M \text{ catch } \{ \} \text{ in } x \\ &\rightarrow^* \text{try } x \Leftarrow V \text{ catch } H \text{ in } x \\ M \uparrow E &\iff \text{try } x \Leftarrow M \text{ catch } \{ \} \text{ in } x \\ &\rightarrow^* \text{try } x \Leftarrow \text{raise } E \text{ catch } H \text{ in } P \quad (E \notin H) \end{aligned}$$

This formulation of the transition semantics is the one which we have used when reasoning about observational congruence for MIL in order to validate effect-based transformations [2].⁴

4 Conclusions

Although the point is undeniably a small one, we hope we have convinced the reader that the *try-catch-in* syntax for exception handling really is more well-behaved than the traditional *handle* construct. It is also probably worth noting that if one translates a language with exceptions into one without them, by using sums to encode the exceptions monad (if the set of exceptions is infinite then this requires either infinite syntax or defaults in pattern matching), then the derived elimination construct for computations is essentially *try-catch-in*. (The difference is that *all* exceptions are always caught, though all but a finite number are then rethrown.)

As far as we know, MIL is the first language to use *try-catch-in*, though we are not the only people to have spotted that it might be a useful programming construct – whilst we were writing this Judicael Courant suggested the same thing on the CAML mailing list [6]. As far as programming language design goes, we would advocate retaining the *handle* construct as well as allowing *try-catch-in*, as it is simpler and suffices in many cases. In a compiler intermediate language, or for theoretical work, *try-catch-in* is a better choice.

From a methodological perspective, we feel that this is another small piece of evidence for the benefits of taking insights from proof-theory seriously when doing language design. Although the solution seems obvious in retrospect, and other people might have reached it by a different route, we personally would not have recognised that there was an identifiable problem in the first place (as opposed to a sequence of disconnected ugly bits of code and messy proofs) had we not been thinking in terms of proof-theoretic normal forms.

References

- [1] M. Abadi. $\top\top$ -closed relations and admissibility. *Mathematical Structures in Computer Science*, 10(1), 2000.

⁴Though, embarrassingly, the HOTS paper gives an incorrect shorthand for one handler covering another in the operational semantics.

$$\begin{aligned}
\text{try } x \Leftarrow V \text{ catch } H \text{ in } P &\rightarrow \text{try } y \Leftarrow P[V/x] \text{ catch } \{\} \text{ in } y & (P \neq x) \\
\text{try } x \Leftarrow \text{raise } E \text{ catch } H \text{ in } P &\rightarrow \text{try } y \Leftarrow H(E) \text{ catch } \{\} \text{ in } y \\
\text{try } x \Leftarrow (\lambda y. M V) \text{ catch } H \text{ in } P &\rightarrow \text{try } x \Leftarrow M[V/y] \text{ catch } H \text{ in } P \\
\text{try } x \Leftarrow (\text{try } y \Leftarrow M \text{ catch } H \text{ in } P) \text{ catch } H' \text{ in } Q & \\
&\rightarrow \text{try } y \Leftarrow M \text{ catch } (H \text{ catch } H' \text{ in } Q) \text{ in } \text{try } x \Leftarrow P \text{ catch } H' \text{ in } Q
\end{aligned}$$

$$\begin{aligned}
\{E_i \Rightarrow N_i\} \text{ catch } \{E'_j \Rightarrow N'_j\} \text{ in } x.Q &\stackrel{\text{def}}{=} \{E_i \Rightarrow \text{try } x \Leftarrow N_i \text{ catch } \{E'_j \Rightarrow N'_j\} \text{ in } Q\} \\
&\cup \{E'_j \Rightarrow N'_j \mid \exists i. E_i = E'_j\}
\end{aligned}$$

Figure 3: Transition semantics

- [2] N. Benton and A. Kennedy. Monads, effects and transformations. In *Third International Workshop on Higher Order Operational Techniques in Semantics (HOOTS), Paris*, volume 26 of *Electronic Notes in Theoretical Computer Science*. Elsevier, September 1999.
- [3] N. Benton, A. Kennedy, and G. Russell. Compiling Standard ML to Java bytecodes. In *3rd ACM SIGPLAN Conference on Functional Programming*, September 1998.
- [4] P. N. Benton, G. M. Bierman, and V. C. V. de Paiva. Computational types from a logical perspective. *Journal of Functional Programming*, 8(2):177–193, March 1998. Preliminary version appeared as Technical Report 365, University of Cambridge Computer Laboratory, May 1995.
- [5] G.M. Bierman. A computational interpretation of the $\lambda\mu$ -calculus. In L. Brim, J. Gruska, and J. Zlatuška, editors, *Proceedings of Symposium on Mathematical Foundations of Computer Science*, volume 1450 of *Lecture Notes in Computer Science*, pages 336–345, 1998.
- [6] J. Courant. A common use of try...with. Message to the CAML mailing list 16 December 1999 <http://pauillac.inria.fr/caml/caml-list/2121.html>, 1999.
- [7] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103:235–271, 1992.
- [8] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *Proceedings of the 1993 Conference on Programming Language Design and Implementation*. ACM, 1993.
- [9] J.-Y. Girard, Y. Lafont, and P. Taylor. *Proofs and Types*. Number 7 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1989.
- [10] R. Harper, B. Duba, and D. MacQueen. Typing first-class continuations in ML. *Journal of Functional Programming*, 4(3):465–484, October 1993.
- [11] J. Hatcliff and O. Danvy. A generic account of continuation-passing styles. In *Proceedings of the 21st Annual Symposium on Principles of Programming Languages*. ACM, January 1994.
- [12] S. Peyton Jones, A. Reid, T. Hoare, S. Marlow, and F. Henderson. A semantics for imprecise exceptions. In *Proceedings of PLDI'99, Atlanta*, 1999.
- [13] S. Peyton Jones and A. Santos. A transformation-based optimiser for Haskell. *Science of Computer Programming*, 32(1-3):3–47, September 1998.
- [14] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. The MIT Press, 1997.
- [15] E. Moggi. Notions of computation and monads. *Information and Computation*, 93:55–92, 1991.
- [16] L. C. Paulson. *ML for the Working Programmer*. Cambridge University Press, 1991.
- [17] A. M. Pitts. Operational semantics for program equivalence. Invited talk at MFPS XIII, CMU, Pittsburgh. See <http://www.cl.cam.ac.uk/users/amp12/talks/>, 1997.
- [18] A. M. Pitts. Parametric polymorphism and operational equivalence. *Mathematical Structures in computer Science*, 10:1–39, 2000. A preliminary version appeared in *Proceedings, Second Workshop on Higher Order Operational Techniques in Semantics (HOOTS II), Stanford CA, December 1997*, Electronic Notes in Theoretical Computer Science 10, 1998.
- [19] A. M. Pitts and I. D. B. Stark. Operational reasoning for functions with local state. In A. D. Gordon and A. M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute, pages 227–273. Cambridge University Press, 1998.
- [20] A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. *Lisp and Symbolic Computation*, 6(3/4):289–360, 1993.
- [21] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 15 November 1994.