

# Exceptions

CPSC 509: Programming Language Principles

Ronald Garcia\*

13 November 2011

## Introduction

Abstract machine semantics give you a mechanical process for decomposing a program into an evaluation context and redex. This process has been left implicit in all of our reduction semantics before. In addition, abstract machine semantics have a much closer connection to a real implementation of a programming language. The expressions that make up a program act as an instruction set for the abstract machine, and the evaluation context (in inside-out form) acts as its runtime stack. Many actual language implementations have been developed by expanding upon such abstract machines.

In this section, we use the intuition provided by abstract machines to introduce a common language feature that manipulates program control: exceptions. It turns out that exceptions can also be presented using reduction semantics, but seeing how control really flows through a machine can give a better feel for what may appear too abstract in reduction semantics.

## A Language with Exceptions

We extend the language of Boolean and arithmetic expressions with two mechanisms: one for raising exceptions and one for catching them. The syntax of the language follows.

### Syntax

$$\begin{aligned} x &\in \text{VAR}, & n &\in \mathbb{Z}, & t &\in \text{TERM}, & v &\in \text{VAL}, & r &\in \text{REDEX}, & E &\in \text{ECTXT} \\ t &::= x \mid \text{true} \mid \text{false} \mid \text{if } t \text{ then } t \text{ else } t \mid n \mid t + t \mid \text{raise } t \mid \text{try } t \text{ handle } x.t \\ v &::= n \mid \text{true} \mid \text{false} \\ r &::= v + v \mid \text{if } v \text{ then } t \text{ else } t \mid \text{try } v \text{ handle } x.t \\ E &::= \square \mid E[\square + t] \mid E[v + \square] \mid E[\text{if } \square \text{ then } t \text{ else } t] \mid E[\text{raise } \square] \mid E[\text{try } \square \text{ handle } x.t] \end{aligned}$$

The `raise  $t$`  expression evaluates  $t$  to a value and then raise it as an exception. The first part of that behavior is captured by the evaluation context form  $E[\text{raise } \square]$ . The counterpart expression to `raise` is `try  $t_1$  handle  $x.t_2$` . This expression first evaluates  $t_1$ , as suggested by the evaluation context form  $E[\text{try } \square \text{ handle } x.t_2]$ . If  $t_1$  evaluates to a normal value  $v_1$ , then the result of the expression is  $v_1$ . However, if  $t_1$  raises an exception `raise  $v$` , then the handler  $t_2$  is run, with  $x$  bound to the value  $v$ . The language of terms  $t$  is also extended with variables  $x$  so that exception handlers can refer to the raised value.

---

\*© Ronald Garcia.

## Extending the Machine

To extend the machine with support for these expressions, we first add the basic rules to refocus that point out where evaluation proceeds in the body of **raise** and **try**.

$$\begin{aligned} \langle E, \mathbf{raise} \ t \rangle_{\text{refocus}} &\longrightarrow \langle E[\mathbf{raise} \ \square], t \rangle_{\text{refocus}} \\ \langle E, \mathbf{try} \ t_1 \ \mathbf{handle} \ x.t_2 \rangle_{\text{refocus}} &\longrightarrow \langle E[\mathbf{try} \ \square \ \mathbf{handle} \ x.t_2], t_1 \rangle_{\text{refocus}} \\ \\ \langle E[\square + t], v \rangle_{\text{return}} &\longrightarrow \langle E[v + \square], t \rangle_{\text{refocus}} \\ \langle E[v_1 + \square], v_2 \rangle_{\text{return}} &\longrightarrow \langle E, v_1 + v_2 \rangle_{\text{reduce}} \\ \langle E[\mathbf{if} \ \square \ \mathbf{then} \ t \ \mathbf{else} \ t], v \rangle_{\text{return}} &\longrightarrow \langle E, \mathbf{if} \ v \ \mathbf{then} \ t \ \mathbf{else} \ t \rangle_{\text{reduce}} \end{aligned}$$

We also need to extend the return rules to consider the two new contexts. First, **try** is discarded if a value is returned to it; we'll do that during reduction.

$$\begin{aligned} \langle E[\mathbf{try} \ \square \ \mathbf{handle} \ x.t_2], v_1 \rangle_{\text{return}} &\longrightarrow \langle E, \mathbf{try} \ v_1 \ \mathbf{handle} \ x.t_2 \rangle_{\text{reduce}} \\ \langle E, \mathbf{try} \ v_1 \ \mathbf{handle} \ x.t_2 \rangle_{\text{reduce}} &\longrightarrow \langle E, v_1 \rangle_{\text{refocus}} \end{aligned}$$

If a value is returned to **raise**, though, then we must prepare to raise an exception. To do so, we introduce a new configuration to the machine to raising an exception.

$$\begin{aligned} C &\in \text{CFG} \\ C &::= \langle E, t \rangle_{\text{refocus}} \mid \langle E, r \rangle_{\text{reduce}} \mid \langle E, v \rangle_{\text{return}} \mid \langle E, \mathbf{raise} \ v \rangle_{\text{raise}} \end{aligned}$$

Then when a value is returned to **raise**, an exception is raised:

$$\langle E[\mathbf{raise} \ \square], v \rangle_{\text{return}} \longrightarrow \langle E, \mathbf{raise} \ v \rangle_{\text{raise}}$$

The machine rules for **raise** proceed to pop frames off of the evaluation context as the exception propagates upwards:

$$\begin{aligned} \langle E[\square + t], \mathbf{raise} \ v \rangle_{\text{raise}} &\longrightarrow \langle E, \mathbf{raise} \ v \rangle_{\text{raise}} \\ \langle E[v_1 + \square], \mathbf{raise} \ v \rangle_{\text{raise}} &\longrightarrow \langle E, \mathbf{raise} \ v \rangle_{\text{raise}} \\ \langle E[\mathbf{if} \ \square \ \mathbf{then} \ t \ \mathbf{else} \ t], \mathbf{raise} \ v \rangle_{\text{raise}} &\longrightarrow \langle E, \mathbf{raise} \ v \rangle_{\text{raise}} \\ \langle E[\mathbf{raise} \ \square], \mathbf{raise} \ v \rangle_{\text{raise}} &\longrightarrow \langle E, \mathbf{raise} \ v \rangle_{\text{raise}} \end{aligned}$$

In principle, these reductions could have been worked into the reduce rules, but these separate configurations are more closely related to what happens in a real implementation. However, the **try** contexts catches any raised exception that reaches it and redirects execution to its associated handler.

$$\langle E[\mathbf{try} \ \square \ \mathbf{handle} \ x.t], \mathbf{raise} \ v \rangle_{\text{raise}} \longrightarrow \langle E, [v/x]t \rangle_{\text{refocus}}$$

Here we use substitution to bind the value  $v$  to the variable  $x$  in the body of the handler, and refocus continues computation in the handler.

To get a better sense of how exceptions work, consider the trace of a program that uses exceptions:

$$\begin{aligned}
 & \langle \square, 5 + \text{try (if (7 + raise 4) then true else false) handle } x.x + 2 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[\square + \text{try (if (7 + raise 4) then true else false) handle } x.x + 2], 5 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[\square + \text{try (if (7 + raise 4) then true else false) handle } x.x + 2], 5 \rangle_{\text{return}} \\
 \rightarrow & \langle \square[5 + \square], \text{try (if (7 + raise 4) then true else false) handle } x.x + 2 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2], (\text{if (7 + raise 4) then true else false}) \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}], 7 + \text{raise 4} \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][\square + \text{raise 4}], 7 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][\square + \text{raise 4}], 7 \rangle_{\text{return}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][7 + \square], \text{raise 4} \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][7 + \square][\text{raise } \square], 4 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][7 + \square][\text{raise } \square], 4 \rangle_{\text{return}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}][7 + \square], \text{raise 4} \rangle_{\text{raise}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2][\text{if } \square \text{ then true else false}], \text{raise 4} \rangle_{\text{raise}} \\
 \rightarrow & \langle \square[5 + \square][\text{try } \square \text{ handle } x.x + 2], \text{raise 4} \rangle_{\text{raise}} \\
 \rightarrow & \langle \square[5 + \square], 4 + 2 \rangle_{\text{refocus}} \rightarrow \langle \square[5 + \square][\square + 2], 4 \rangle_{\text{refocus}} \rightarrow \langle \square[5 + \square][\square + 2], 4 \rangle_{\text{return}} \\
 \rightarrow & \langle \square[5 + \square][4 + \square], 2 \rangle_{\text{refocus}} \rightarrow \langle \square[5 + \square], 4 + 2 \rangle_{\text{reduce}} \rightarrow \langle \square[5 + \square], 6 \rangle_{\text{refocus}} \\
 \rightarrow & \langle \square[5 + \square], 6 \rangle_{\text{return}} \rightarrow \langle \square, 5 + 6 \rangle_{\text{reduce}} \rightarrow \langle \square, 11 \rangle_{\text{refocus}} \rightarrow \langle \square, 11 \rangle_{\text{return}} .
 \end{aligned}$$

Phew!

## Reduction Semantics for Exceptions

Now that we've seen how exceptions are represented in an abstract machine, we can consider how they might be represented using reduction semantics. The nice thing about the abstract machine semantics is that the behavior of the machine looks a lot like how a program would really run on a real language implementation. In fact, many language implementations (compilers and interpreters) are little more than an implementation of an abstract machine on a real machine. Here you can really see how raising an exception discards chunks of the stack until it is stopped by an exception handler.

The downside of abstract machine semantics is that they are big and noisy. All of the details about finding a redex, returning results, and transitioning between machine modes help us see how a concrete implementation works, but they add a lot of extra details which can obscure the essence of a language design.

Reduction semantics, on the other hand, are quite economical. They present the core of a language semantics, thereby emphasizing the interesting aspects. As it turns out, we can give a semantics for exceptions in reduction semantics, without resorting to abstract machines. We already have most of the necessary pieces to present the reduction semantics: syntax and evaluation contexts. Also, we have all of the notions of reduction that were introduced in the last lecture, and were ultimately became the `reduce` rules of the abstract machine. What are needed now are notions of reduction that address raising and catching exceptions. Those notions of reduction follow:

$$\begin{aligned}
 \text{raise } v + t & \rightsquigarrow \text{raise } v \\
 v_1 + \text{raise } v & \rightsquigarrow \text{raise } v \\
 \text{if raise } v \text{ then } t \text{ else } t & \rightsquigarrow \text{raise } v \\
 \text{raise (raise } v) & \rightsquigarrow \text{raise } v \\
 \text{try raise } v \text{ handle } x.t & \rightsquigarrow [v/x]t \\
 \text{try } v \text{ handle } x.t & \rightsquigarrow v
 \end{aligned}$$

Do these rules look familiar? Each of them corresponds to a `raise` rules of the abstract machine. The first 4 are quite boring: throw away stack frames that are not try-handle blocks. The second-to-last one handles

the exception that was raised. The last one propagates a value through a handler. Having extended the notions of reduction as above, we could elaborate this into a slightly different abstract machine semantics. All of the above rules would become `reduce` rules, rather than `raise` rules, and `raise` configurations would transition to `reduce` configurations

There are yet more succinct ways to express these exceptions as reduction semantics. We can represent the process of raising an exception and handling it as one big reduction step. To do this, we first define the subset of evaluation contexts that do not handle exceptions:

$$F \in \text{RCtxt} \subseteq \text{ECtxt}$$

$$F ::= \square \mid F[\square + t] \mid F[v + \square] \mid F[\text{if } \square \text{ then } t \text{ else } t] \mid F[\text{raise } \square]$$

Then given this definition, raising and handling an exception can be captured in a single notion of reduction:

$$\text{try } F[\text{raise } v] \text{ handle } x.t \longrightarrow [v/x]t$$

Here,  $F[\text{raise } v]$  represents some expression that can be broken down into a restricted evaluation context (containing no try-handle frames) and an expression that is raising a value. Now in one fail swoop, one step of the reduction semantics, the exception is raised and handled.

## Exceptions with Resumptions

In every mainstream language with exceptions, the process of raising an exception throws away part of the control context (i.e. stack, i.e. evaluation context) until arriving at a handler for the exception. Some language design researchers have wondered whether this behavior was strictly necessary, or whether it would make sense to handle an exception and go right back to whatever was going on when the exception was raised. This model of exceptions is called *exceptions with resumptions*, and we can easily model this variation<sup>1</sup>.

The syntax of this kind of language is exactly the same as before. What changes is the *behavior*. We can capture this change of behavior simply by altering what the abstract machine does when an exception is raised. In order to resume the computation that was happening when the exception was raised, the machine must remember all the information that the other machine simply discarded. To this end, we change the `raise` configuration to keep track of the part of the context that used to be discarded

$$C ::= \dots \mid \langle E, \text{raise } v, E \rangle_{\text{raise}}$$

and the machine transitions that involve `raise` configurations are updated to save and restore context information.

$$\begin{aligned} \langle E[\text{raise } \square], v \rangle_{\text{return}} &\longrightarrow \langle E, \text{raise } v, \square \rangle_{\text{raise}} \\ \langle E_1[\square + t], \text{raise } v, E_2 \rangle_{\text{raise}} &\longrightarrow \langle E_1, \text{raise } v, \square[\square + t] ++ E_2 \rangle_{\text{raise}} \\ \langle E_1[v_1 + \square], \text{raise } v, E_2 \rangle_{\text{raise}} &\longrightarrow \langle E_1, \text{raise } v, \square[v_1 + \square] ++ E_2 \rangle_{\text{raise}} \\ \langle E_1[\text{if } \square \text{ then } t \text{ else } t], \text{raise } v, E_2 \rangle_{\text{raise}} &\longrightarrow \langle E_1, \text{raise } v, \square[\text{if } \square \text{ then } t \text{ else } t] ++ E_2 \rangle_{\text{raise}} \\ \langle E_1[\text{raise } \square], \text{raise } v, E_2 \rangle_{\text{raise}} &\longrightarrow \langle E_1, \text{raise } v, \square[\text{raise } \square] ++ E_2 \rangle_{\text{raise}} \\ \langle E_1[\text{try } \square \text{ handle } x.t], \text{raise } v, E_2 \rangle_{\text{raise}} &\longrightarrow \langle E_1[\text{try } \square \text{ handle } x.t] ++ E_2, [v/x]t \rangle_{\text{refocus}} \end{aligned}$$

To define these transitions, we rely on the function  $E_1 ++ E_2$  (also called *append*), which sticks the  $E_2$  context into  $E_1$ 's hole (I leave it's definition as an exercise).

The `return` configuration that raises an exception must now make space to remember the part of the context between the exception and its handler. Then, each step of `raise` that does not handle the exception pops the current frame, but stores it in the second context. Finally when an exception handler is reached, the exception value is passed to the handler, using substitution, which is executed back in the original context from which the exception was raised.

As with the traditional exception semantics, we can recast exceptions with resumption as a reduction semantics where exception raising and handling happens in one reduction step:

$$\text{try } F[\text{raise } v] \text{ handle } x.t \longrightarrow \text{try } F[v/x]t \text{ handle } x.t.$$

<sup>1</sup>As it turns out, exceptions with resumptions have been deemed a **bad idea**!

Notice how this semantics differs from traditional exceptions. The intermediate context, including the exception handler, is preserved. What happens in this step is that the raised exception value is substituted into a copy of the exception handler. This behavior is very similar to the process of looking up a function (i.e.  $\lambda x.t$ ) and applying it to  $v$ .