Modular Rollback through Control Logging

A Pair of Twin Functional Pearls

Olin Shivers
Northeastern University
shivers@ccs.neu.edu

Aaron Turon
Northeastern University
turon@ccs.neu.edu

Abstract

We present a technique, based on the use of first-class control operators, enabling programs to maintain and invoke rollback logs for sequences of reversible effects. Our technique is modular, in that it provides complete separation between some library of effectful operations, and a client, "driver" program which invokes and rolls back sequences of these operations. In particular, the checkpoint mechanism, which is entirely encapsulated within the effect library, logs not only the library's effects, but also the client's control state. Thus, logging and rollback can be almost completely transparent to the client code.

This separation of concerns manifests itself nicely when we must implement software with sophisticated error handling. We illustrate with two examples that exploit the architecture to disentangle some core parsing task from its error management. The parser code is completely separate from the error-correction code, although the two components are deeply intertwined at run time.

Categories and Subject Descriptors D.1.1 [*Applicative (Functional) Programming*]

General Terms Algorithms, Design, Languages

Keywords checkpoint, delimited control, error repair

Prologue

We all make mistakes. What counts is how we handle them.

There's little less pleasant in programming than watching beautiful code turn ugly. And nothing causes more contortion than catching and correcting input errors.

The problem is that error handling is not naturally modular: it is context-dependent, often relying on the details and current state of input processing. Extricating error checking into a separate prepass often entails duplicating some of the processing intended for correct input. For complicated artifacts like typecheckers, which are hard enough to get right in the first place, code to produce friendly error messages can obscure or even dwarf the code to check types [9]. When input is provided and processed interactively, checking and processing *must* be interleaved.

We tell a story in two parts, the moral of which is: error handling can be separable, even beautiful. The hero of our tale is `call/cc`, for to err is human, but to forgive requires first-class control.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP'11, September 19–21, 2011, Tokyo, Japan.
Copyright © 2011 ACM 978-1-4503-0865-6/11/09. . . $10.00

Part one takes place in an "eager" Scheme REPL, which parses complete s-expressions as the user types a closing parenthesis, long before a carriage return. Yet, despite advancing the state of the parser, the user can interactively alter erroneous text by "backing up" with the delete key and re-entering corrected text. So here the error-handling is done, in part, by the user, but done in a way that's invisible to the input-processing code.

Part two retells the tale in ML, using a similar technique to automatically discover and repair the "true" source of a syntax error, perhaps far removed from the location where the parser got stuck. It's the functional programming equivalent to Burke-Fisher error repair [2]. Again, the key benefit is the complete separation between the parsing algorithm and the repair algorithm: you write a parser in any style you like, and send it through our `BurkeFisher` functor to get a new parser that intelligently repairs errors.

We use first-class control in two crucial ways:

- First, to provide rollback logging. We layer an interface on top of `call/cc`, constructed so that the very performance of a reversible side-effect has an additional effect: it logs the appropriate reversal. Control effects and other effects smoothly interleave, whether going forward or backward.
- Second, to provide modularity: we sneak checkpointing through an unsuspecting parser. Although parsing and error correction *functionality* is interleaved, the *code* is not. Parsing is done via standard recursive descent, with no concern for error correction; likewise, error correction needs only understand the structure of the input to the parser.

We hope that, while savoring the stories below, you will see that these two techniques transcend them.

1. Prompt reading and effect logging in Scheme

In our first example, we show how to recreate, with modern tools, a bit of "lost art" from classic LISP input systems: how to intertwine input parsing with terminal-entry text editing without having to intertwine the code that carries out these separate tasks. As we do this, we'll discover a general technique for constructing "command" systems that allows us to transparently *log effects* in a way that permits them to be rolled back on demand, without having to compromise the simplicity of client code that issues sequences of these commands.

That's the big picture; let's begin by diving into the specifics of our example. S-expressions, the standard syntactic form for programs and data in LISP and Scheme, have an interesting property: they are frequently self-delimiting. For example, we do not need to look past the terminating right parenthesis in the string "`(+ x y)`" to know that we have reached the end of the form. By way of contrast, this is not true of numerals: if an s-expression reader sees the string "`387`", it must look ahead one character to know if this string

comprises an entire term, or if it is simply a prefix of a larger term. If the following character is a left parenthesis, then the string represents the integer three hundred eighty-seven (which happens to be followed by some list to be read at a later date); if the following character is another digit, then our string is simply the first three digits of some longer numeral.

Here is the breakdown for standard Scheme s-expressions, showing which kinds are self-terminating, and which aren't:

Self-terminating	Lists:	`()`, `(+ x 3)`, ...
	Vectors:	`#()`, `#(-2 7)`, ...
	Strings:	`"fred"`, `"dog"`, ...
	Booleans:	`#f`, `#t`
	Characters:	`#\z`, `#\q`, ...
Not self-terminating	Symbols:	`fred`, `dog`, ...
	Numerals:	`387`, `0`, ...

The details vary across dialects of LISP and Scheme, but the general property remains.

Interactive systems that read commands or other structured input from a user frequently take their inputs in the form of an s-expression—the classic example is the LISP read-eval-print loop. In such a system, the human user enters some command, expression or definition at the terminal; the operating system passes this text to the LISP process, which parses it into an s-expression, and then computes the command, expression, definition, *etc.* that it describes.

It turns out that old, 1970's- and 1980's-era LISP systems (*e.g.*, Maclisp and Lisp Machine Lisp) took advantage of the self-delimiting property enjoyed by their s-expression-based interaction languages. When reading s-expressions, these interactive systems would read characters from the terminal eagerly—that is, they turned off the operating system's buffering machinery, so that each character was presented to the reader as soon as the user hit the key on the keyboard. Thus, the parse would complete the instant the terminating close-parenthesis was entered, without requiring the user to enter a redundant newline. It's a small but pleasant convenience that, in the post-Lisp-Machine age, has become lost from the text-interaction modality.

What makes providing this capability tricky is the task of error handling. Human typists make mistakes, which they'd like to correct as they enter their text. This is usually handled by the so-called "line discipline" code in the operating system's terminal device driver. This code buffers terminal input and provides a simple text editor for operating upon this buffered text: most characters are taken as input from the user, and echoed back to the terminal device, to make them visible on entry. A few keys, however, are reserved for this editing task. On a Unix system, for example:

- The delete key "backs up" one character: the OS removes the last character from the input buffer and outputs a backspace-space-backspace sequence to the terminal to erase it from the screen.
- The control-U character typically erases the entire line of text, erasing the entire buffer's worth of input.
- The characters control-C and control-Z direct the OS to send an asynchronous control signal to the reading process.
- Control-V turns off any special treatment of the following character.
- Control-D represents end-of-file.

The operating system does not release the buffered input to a process attempting to read from the terminal device until the user enters a newline character: this commits the input sequence.

Now, in order to provide eager reading, we have to turn this OS-provided machinery off: the user process must now be responsible for echoing input and implementing the interactive editing that humans use to back up through mistaken input and make corrections. This is not so terrible; we can encapsulate the editing code in a small library and be done with it.

What's tricky is that the parsing computation is now *interleaved* with the text-editing computation. Formerly, these things were split into distinct phases. First, we enter a string, with no parsing going on at all; during this phase, we can perform editing on each line as we enter it. Once we strike the newline key, the final version of that line is shipped off to the parser phase (*i.e.*, the interpreter), which consumes the text, produces a parse tree (that is, an s-expression), and proceeds with the requested task.

In our new (well, 1970's) world, the application is carrying out the parse as text is entered by the user. If we change our minds and wish to erase some of the text we've entered, as we back up through the input text, we also have to *rewind the parse computation*. For example, an s-expression reader is typically a little recursive-descent parser. When it is reading the elements of a nested list and it encounters a close-parenthesis character, it returns the accumulated list to its recursive caller. If we were to subsequently decide to *delete* that close parenthesis, we'd have to back the parse computation back down into that formerly completed call.

What we need is the ability to take "snapshots" of the computation at various points as we parse. Roughly speaking, every time the parser calls out to the "read character" routine, that routine should first checkpoint the parse computation and save the checkpoint away on a list. If the user tries to delete some previously entered characters, the input routine can pull the appropriate checkpoint from its saved history and reset the parse computation to that previous character-reading state.

By this point, it should be clear that the right tool for this job is the Scheme `call-with-current-continuation` (or, `call/cc`) procedure: creating computational checkpoints is its *raison d'être*.

The rubout port and the reset protocol

In Scheme, we do I/O on *ports*; our task is to construct a new kind of "rubout port" for reading from a terminal. We'll construct our rubout port using a low-level mechanism in Scheme48 [7]/scsh [12] that permits programmers to define their own input ports, which can be passed to the system input procedures like any other input port. To define one of our custom rubout ports, we provide the extended-port constructor two things. The first is a fixed suite of "port methods," described by a record whose fields are procedures the input system will use to read a character, "peek" at a character, close the port, and a few other less-important operations. We also provide a data value which encapsulates the port state; this value will be passed to the method procedures when operations are performed on the port. Here is the definition of the rubout port's data record:

```
(define-record rubout-port-data
  ttyin              ; input tty port
  ttyout             ; output tty port, for echoing
  (prev-tty-info #f) ; tty's original echo state
  (peek-char     #f) ; "lookahead" peek cache
  (checkpoint    #f)); most recent checkpoint
```

This piece of Scheme code defines a procedure
`(make-rubout-port-data` *ttyin ttyout*`)`
which constructs a record from a pair of ports connected to the terminal device; the other three fields of the record are initialised to false, `#f`. We also get field-accessor procedures, such as
`(rubout-port-data:ttyin` *rpd*`)`
`(rubout-port-data:ttyout` *rpd*`)`
and so forth, and some analogous field-assignment procedures with names like `(set-rubout-port-data:peek-char!` *rpd char*`)`.

Note the `checkpoint` field. Our intention is that every time application code does a (`read-char` rp) operation that causes the rubout-port to actually get a character from the terminal, before returning the character to the caller, the code will also grab a continuation with `call/cc` and stow it away in the port's `checkpoint` field. If, during a later attempt to read, the rubout-port code reads, say, a "delete" character and decides to undo the previous read, it will be able to do so by fetching this saved value from the `checkpoint` field of the rubout-port's data record and invoking it—this will reset the entire computation back to the previous read point.

The key piece of design we must pin down is the protocol used to invoke the checkpointed continuation. Specifically, a checkpoint is a Scheme continuation that must be applied to a single boolean argument: (`checkpoint just-1?`). If the `just-1?` argument is true, we wish to rewind the computation back just one step—that is, to the immediately prior read operation. If the argument is false, then we wish to rewind all the way back to the beginning of the entire read session.

The core effects: reading and echoing

We can now define the pair of low-level procedures that perform our system's two primitive side-effects: reading a character from the keyboard, and echoing an input character to the display. Since each of these procedures carries out an effect, it must log the effect as it performs it—that is, we must update the port's checkpoint to add a rollback handler that will undo the effect if we rewind back through this point in the computation. Here is the code for echoing:

```
(define (echo c pd)
  (let* ((oport (rubout-port-data:ttyout pd))
         (reset (rubout-port-data:checkpoint pd)))
    (write-char c oport)     ; 1: Echo the character.

    ;; 2: Set a checkpoint to undo the echo.
    (set-rubout-port-data:checkpoint pd
      (call/cc (λ (ret)
                 (let ((just-1? (call/cc ret)))
                   (print-rubout-sequence oport)
                   (reset just-1?)))))))

;;; Output backspace/space/backspace to the port.
(define (print-rubout-sequence oport)
  (write-string "\b \b" oport))
```

The rubout-port machinery calls `echo` whenever it needs to echo a character just read to the terminal. The procedure writes the character out on line 4, and then logs the effect in the last half of the procedure by updating the port's checkpoint. This is the first piece of serious continuation manipulation we've performed, so we will trace through its execution carefully. The first, outer `call/cc` creates a return point `ret`; applying `ret` to some value cp will cause cp to be installed as the port's new checkpoint. The second, inner `call/cc` creates the actual checkpoint continuation; `call/cc` passes this continuation to `ret`, so, just as we described, this continuation gets installed as the current checkpoint.

Now, consider what happens when this checkpoint is invoked at some later time, by fetching it from the rubout port's data record and applying it to some boolean argument. We'll reset the computation back to *now*: the inner `call/cc` will return the boolean value, so it will be bound by the `let` form to the variable `just-1?`. Then we'll proceed into the body of the `let`, which contains the rollback action: we write out a backspace/space/backspace character sequence to the terminal, which will erase the character we previously echoed back on line 4 of the code, then we'll pass the `just-1?` boolean on to *our* checkpoint. We do this because we are in the process of rewinding back to some prior input—the echo action is an output effect, not an input effect, so we need to continue rewinding the computation.

Our next procedure (which appears in Figure 1) is the primitive character-input procedure. When the rubout-port machinery needs to actually read a character, it applies `%read-char` to the port's data record. This code is, essentially, our "line discipline" driver code, written in Scheme. The procedure sits in a loop (the named-let function `lp`), which does an input operation on the actual terminal, binding the variable `c` to the character entered by the user. Then we perform a variety of actions, depending on the character. Inducing rollback is easy: if the character is the delete character, we apply the port's current checkpoint continuation `rubout` to true, which will abort what we're doing, and reset to the previous read (and also undo any echoing we might have logged in-between). That previous read will then be able to get a new character from the user and return it to the parser in place of the original character it had input back when it first ran. (We'll see the code that does this—the code at the target end of our reset-continuation's non-local jump—in just a moment.)

If the character is the line-kill character, control-U, then we apply the checkpoint to false; by the rules of the checkpoint invocation protocol, this will induce a rewind all the way back to the beginning of the entire parse session, clearing previously echoed characters from the display as we rewind.

If the user entered control-C or control-Z, the code sends the current process the appropriate OS signal, and then loops by tail-calling (`lp`), to continue trying to read a character. (The process signals itself in a context that resets the terminal's echoing and buffering state to its original settings, but this is a fine point we can skip.)

These first few cases describe the text-editing functionality of our "device driver," where characters input by the user are not intended as data to be passed on to the process, but are rather interpreted directly by the input-port system as requesting various actions. The final case (the `else` clause of the `cond`) is the actual-input case: we've read an ordinary character `c` (or we've encountered the end of file, or we've read the "knockdown" character, control-V, followed by any character at all). Since we can now be considered to have successfully accomplished an input side-effect, before we return `c` to our caller, we must first log the operation by updating the port's checkpoint. Just as with `echo`, this is done with a nested, double `call/cc` pattern. The first, outer `call/cc` simply captures the context we'll use to return `c` and proceed with the parse by returning from `%read-char`. The second, inner `call/cc` creates the checkpoint continuation, naming it `ckpt`. We then install `ckpt` into the port's data record, and apply `ret` to the character `c`, which will lead us to produce `c` as the return value of `%read-char`. What happens when, at some later time, the checkpoint continuation is invoked? What we want to accomplish, if this happens, is to rewind the computation back to *now*, then get a new character c' from the terminal and return that *instead* of the character `c` we just now produced.

That's exactly what happens. Suppose some future call to `%read-char` gets a delete character and so fetches the checkpoint continuation we just now created from the port's data record and applies it to true (*i.e.*, executes line 7 of the procedure). When the checkpoint continuation is applied to true, we'll rewind to "now," and the second, inner `call/cc` will return the true value—so it will be let-bound to the variable `just-1?`, and we'll proceed into the body of the `let` form, whose `if` conditional will jump back to the top of our original loop, `lp`, continuing our current read. Any character c' we get will be returned to our caller, so the net effect is that, after rewinding the computation to here, we'll proceed with c' instead of our originally-read character. In short, we undid the input effect and replaced it with a new one.

```
(define (%read-char port-data)
  (let ((rubout (rubout-port-data:checkpoint port-data))
        (iport  (rubout-port-data:ttyin port-data)))
    (let lp ()
      ;; Assert: (rubout-port-data:peek-data port-data) = #F.
      (let ((c (read-char iport))) ; real input!
        (cond ((rubout-char?   c)  (rubout #t)) ; DEL => rubout a char
              ((linekill-char? c)  (rubout #f)) ; ^U  => kill the whole sexp.

              ((interrupt-char? c) ; ^C => interrupt
               ;; Undo the tty's raw/no-echo modes before raising the signal.
               (with-rubout-port-restored port-data
                 (signal-process 0 signal/int))
               (lp)) ; The process survived the signal => keep reading.

              ((suspend-char? c) ; ^Z => suspend
               ;; Undo the tty's raw/no-echo modes before raising the signal.
               (with-rubout-port-restored port-data
                 (signal-process 0 signal/tstp))
               (lp))

              (else (let ((c (cond ((knockdown-char? c) ; ^V => turn off special handling.
                                    (read-char iport))
                                   ((eof-char? c) *eof-object*) ; ^D => EOF
                                   (else c))))                  ; Normal char

                      ;; Outstanding! We have read an actual character, C.
                      ;; Before returning it to the parser, set a new
                      ;; checkpoint, so that if we later wish to rub C out,
                      ;; we can come back here, get a new character, and resume.
                      (call/cc                             ; Therein lies the rubout.
                       (λ (ret)
                         (let ((just-1? (call/cc (λ (ckpt) ; Set the checkpoint & return C:
                                                   (set-rubout-port-data:checkpoint port-data ckpt)
                                                   (ret c)))))
                           ;; On rewind, we'll bind JUST-1? & come here.
                           (if just-1? (lp)           ; Either re-do the current read,
                               (rubout #f)))))))))))) ; or keep rewinding.
```

Figure 1. The checkpointing character reader

On the other hand, suppose at some point in the future the user enters the "line-kill" character, Control-U. Then the checkpoint continuation will be applied to false (by line 8 of that future invocation of `%read-char`). The computation will be reset to "now," and the inner `call/cc` will produce the false value, which will be let-bound to `just-1?`. So, instead of looping in our current read, we'll keep rewinding by applying *our* checkpoint to false, in the very last line of the procedure: `(rubout #f)`.

In other words, we don't have to keep an explicit stack of checkpoints. The "stack" is implicit in the *environment structure* of the various checkpoints: each checkpoint continuation has access to the next checkpoint back in the timeline: it is bound to the variable `rubout` in the current checkpoint's lexical scope. (Similarly, it is bound to the variable `reset`, in the checkpoint created by the echo procedure.)

Peeking and reading

All the heavy lifting in our system is accomplished by the echo and `%read-char` primitives. The rubout-port system calls them from its peek-character and read-character "method" procedures. Note a policy decision encoded in these procedures: a character is *not* echoed as soon as we get it from the keyboard: peeks don't count. This code doesn't echo a character to the display until the parser actually consumes it with a `read-char` operation.

```
(define (rubout-peek-char port-data)
  ;; A "subsequent" peek: don't checkpoint
  (or (rubout-port-data:peek-char port-data)
      ;; A "first" peek -- checkpoint it
      (let ((c (%read-char port-data)))
        ;; then set the peek cache:
        (set-rubout-port-data:peek-char port-data c)
        c))) ; then return the character.

(define (rubout-read-char port-data)
  ;; If the peek does input, it will be logged:
  (let ((c (rubout-peek-char port-data)))

    ;; If we echo, it will be logged:
    (if (not (eof-object? c)) (echo c port-data))

    ;; Clear the peek cache.
    (set-rubout-port-data:peek-char port-data #f)
    c))
```

Note also that these two higher-level procedures do not concern themselves with setting checkpoints or logging effects. One of the nice properties of our design is that the effect-logging code is tightly bound to the code that commits the effects. Thus, doing a side effect and logging its rollback handler are welded together.

Client code, like the two procedures above, or the application's parser, just commit side-effects at will; it's impossible for them to break the pairing of effects and rollbacks. (This nice property is only extended to the side-effects that we included in our design, of course. Parser clients must be aware that they cannot perform other side-effects and expect them to be undone during rollback.)

Failure is not an option

All that remains is to define `with-rubout-session*`, the procedure used to delimit a single parse session; it appears in Figure 2. We can perform a rubout-enabled parse with something like:

```
(with-rubout-session* rubout-port
  (λ () (read rubout-port)))
```

The central body of this procedure executes in a dynamic context, established by the R5RS Scheme `dynamic-wind` procedure, that saves and restores the rubout port's checkpoint if we should throw out and then back into the session's dynamic extent by invoking saved continuations. It also serves another purpose: the `dynamic-wind`'s exit thunk clears the checkpoint from the rubout port, which permits the checkpoint to be garbage collected even if the rubout port itself continues to remain alive.

Note that the rubout-session's thunk executes in an exception-handler context that treats parser-syntax errors in a clever way: we refuse to accept input from the terminal that triggers a syntax-error exception. If the user should enter a character that causes the parser to raise this error, the exception handler "rings the bell" (in a visual manner), then discards the bad character by invoking the rubout port's current checkpoint, passing it true for its `just-1?` parameter. This rewinds the parse computation back to the read operation where the user entered the offending character (erasing the character as we rewind, if it had been echoed); we then resume the parse by reading an alternate character from the user.

For example, if the user tries to type in an ill-formed dotted-pair that has two dots, *e.g.*, `(a . b . c)`, the parser will stubbornly refuse to accept the second dot, ringing the bell every time the user attempts to enter it to signal that we have departed from the syntactic straight and narrow. Note that this facility is completely independent of the grammar we are parsing, or the details of the parser: the parser is just a piece of code that raises an error exception as soon as it encounters an illegal character.

The final task the rubout-session procedure must perform before invoking the client's parser computation `thunk` is to set the rubout port's initial checkpoint. Again, we see the nested, double `call/cc` pattern. Tracing through the `restart` loop shows that each time we rewind back to the initial checkpoint, we recreate and reinstall it into the port, then begin the parse (that is, call `thunk`) all over again. Thus, invoking the initial checkpoint has the effect of restarting the whole parse.

Discussion

The big idea The first thing to note about this rubout-handler system is that the technique is not at all limited to interactive text entry. The general idea here is that we have a library of effectful (but reversible) operations. Some client performs a series of effects by issuing a sequence of these operations; as the client computes, the application may decide to rewind the driving computation to some previous operation and do something different—perhaps (as in our rubout-handler scenario) in response to error conditions.

The design pattern we propose here is to instrument the effectful operation library to construct a rollback log for the operations it performs. More specifically, *by constructing this rollback log from continuations, we capture not only the library's effects, but also the client's control state at each operation-invocation point.* This is what permits completely transparent rollback of the client computation: it's all due to the power of `call/cc` to package up a general control state.

Delimiting our checkpoints Let's motivate our next point by considering an extension to our rubout-handler system. Many text-input systems have some kind of "history" mechanism: they save previous lines of input, which the user can recall by entering some special control key that cycles through previous entries. Suppose we wished to provide this kind of functionality—and, of course, keep our rubout-handling capability. Unfortunately, the saved continuations that make up our checkpoints capture *too much control state:* they capture not only the parse computation, but also the computational state of the client that called the parser. If we were to reset to one of the checkpoints from a previous read in, *e.g.*, a Scheme interpreter's read-eval-print loop, we'd reset the entire interpreter back to that earlier state!

This is a problem that is exactly handled by delimited-control operators such as Felleisen's `prompt` [4, 13] or Danvy and Filinski's `shift` [3]. We need only delimit the parse computation carried out by the thunk passed to `with-rubout-session*` and we're set. We leave this modification to our code as a (fun) exercise for the interested reader.

The dark side of Church encodings Scheme provides continuations encoded as procedures: we perform a non-local control transfer to a captured continuation by applying it to some argument. This has been a long-standing source of subtle problems using Scheme's continuations. The central issue is that, if we want to call procedure `p`, with argument `a` and continuation `k`, we cannot do the following: `(k (p a))`. This goes wrong because `k`'s underlying continuation does not actually become the *current* continuation until `p` returns; the whole time `p` is executing, the continuation is an extension of the one extant when we began evaluating the entire `(k (p a))` expression. So if `p` raises an exception, we will resolve it using the exception handlers of that continuation, not `k`'s handlers. If we were hoping to reclaim the original continuation's stack by installing `k` as the current continuation, this won't happen while `p` is executing—if it happens to be some long-running thread (such as a web server), then the original continuation's stack will *never* become free, leading to a subtle space leak. These kinds of space leaks aren't just theoretical oddities; they actually occur when programmers build thread schedulers based on continuations [1].

The real problem here is that when we Church-encode a kind of data, we can only do one thing with the data: apply it to an argument—that is the only operation permitted on procedures. In the case of continuations, we need to do two things: perform a function call with the given continuation for the call, and compose an $\alpha \to \beta$ function onto the "end" of a β-accepting continuation, producing an α-accepting continuation. These two procedures provide the necessary interface:

```
(with-continuation cont thunk)
(compose-continuation β-cont α->β)
```

Scheme doesn't have this kind of continuation mechanism, so we have to code in awkward ways to work around the limitations of `call/cc`.

This problem rears its head in our rubout-handler system. Consider the subtle, double-`call/cc` code that creates the new checkpoint when we read a character in `%read-char`. Couldn't we eliminate the inner `call/cc` and just use a simple procedure, instead of an exotic continuation, with the following?

```
(define (with-rubout-session* rubout-port thunk)
  (let ((pd (extensible-input-port-local-data rubout-port))
        (suspended-checkpoint #f))

    (with-raw-mode-rubout-port pd       ; Turn off tty buffering & echoing.
      (dynamic-wind
        ;; Dynamic-wind pre:
        ;;   If we are throwing back into the parse, restore the checkpoint
        ;;   we saved away when we threw out.
        (λ () (set-rubout-port-data:checkpoint pd suspended-checkpoint))

        ;; Dynamic-wind body:
        (λ () (with-syntax-error-handler
                ;; Here's the error handler. If the parser raises a syntax error,
                ;; ring the bell, clear the peek-char cache, then trigger a 1-step
                ;; rewind to rubout the last char, which triggered the error.
                (λ (c) (visible-bell (rubout-port-data:ttyout pd))
                       (set-rubout-port-data:peek-char pd #f)
                       ((rubout-port-data:checkpoint pd) #t))

                ;; Set the initial checkpoint, which comes back here and
                ;; restarts the whole parse.
                (call/cc (λ (go-start-parse)
                           (let restart ()
                             (call/cc (λ (icp)
                                        (set-rubout-port-data:checkpoint pd icp)
                                        (go-start-parse)))
                             ;; Come here when initial checkpoint ICP is invoked:
                             (restart))))

                (thunk))) ; Do the parse.

        ;; Dynamic-wind post:
        ;;   When we're done, clear out the port's checkpoint, so the port
        ;;   won't keep the checkpoint from being gc'd. But... we might not be
        ;;   done. We might be throwing out and later throwing back in. So
        ;;   save away the current checkpoint in case we later throw back in.
        (λ () (set! suspended-checkpoint (rubout-port-data:checkpoint pd))
              (set-rubout-port-data:checkpoint pd #f))))))

(define (visible-bell oport)    ; A simple "visible bell:"
  (write-char #\! oport)        ;   print out a !,
  (sleep 1)                     ;   pause one second, then
  (print-rubout-sequence oport));   erase it.
```

Figure 2. Delimiting a rubout-handler session.

```
(call/cc (λ (ret)
           (set-rubout-port-data:checkpoint port-data
             (λ (just-1?) (if just-1?
                              (ret (lp))
                              (rubout #f))))
           c))
```

Unfortunately, no. If some future delete command applies this checkpoint to true, we will perform the `(lp)` retry *in that future context*; we won't throw back to the previous control state until the `lp` call returns to the Scheme-style Church-encoded continuation `ret`. All the time that `(lp)` is running, it is running with the wrong exception-handler context, the wrong dynamic-wind undo/redo handlers, and so forth. Furthermore, the run-time system can't reclaim the triggering read's stack and other control context until `lp` returns; we want to reclaim it when `lp` *starts*.

The double-`call/cc` pattern establishes the proper context, then captures it with the inner `call/cc`, so things work out properly. It would have been easier and simpler to write this code if we'd had the alternate continuation functionality described above. Our checkpoint-setting code would then look like this:

```
(call/cc
 (λ (ret)
   (let ((ckpt (compose-continuation ret
                 (λ (just-1?)
                   (if just-1? (lp) (rubout #f))))))
     (set-rubout-port-data:checkpoint port-data ckpt))
   c))
```

An observant reader might similarly have wondered why we didn't use a simple procedure as the checkpoint for the `echo` procedure. This is why. As a matter of style, checkpoints in our rubout-handler system—that is, things stored in the `checkpoint` field of a rubout port's data record—are always and only continuations—that is, things created by `call/cc`.

This issue matters less in the case of `echo`'s checkpoints. If we'd written the checkpoint code for `echo` with the simpler

```
(set-rubout-port-data:checkpoint pd
  (λ (just-1?)
    (print-rubout-sequence oport)
    (reset just-1?)))
```

then the only code that would run in the wrong context would be the `(print-rubout-sequence oport)` call, which is short and presumably terminates with no other control effects such as raising an exception. Still, we preferred to work with the discipline of *only* using continuations for checkpoints.

Backwards and forwards An observant reader might also be wondering: old LISP systems didn't have `call/cc` or general continuations. So how did these systems provide prompt reading with rubout handling?

The answer is rather ingenious [10]. These systems couldn't back up to prior checkpoints by invoking saved continuations. Instead, the port machinery would log all the characters read during a session. If the user requested a single-character delete, the rubout handler would raise an exception, which would be caught by an exception handler established at the beginning of the rubout-handling session. The exception handler would delete the last character from the log, and then *completely restart* the entire parse. During the new parse attempt, the port would be in a special "replay" mode: whenever the parser requested a character from the port, the port would get the next item from the log, instead of going to the terminal to read a new character. When the log was exhausted, the port would revert to doing actual input from the terminal. So, to delete a character, the rubout handler would simply reread and reparse everything *but* the deleted character. Instead of going one step back, the system would do a complete restart and then go $n - 1$ steps forward.

Monads and dependent types The technique of using continuations to construct rollback logs for command sequences, where we wish to capture client control state as well, seems nicely suited for expressing in a Haskell monad that uses both continuations and state.

We should mention that Conor McBride was able to code up a functional equivalent to our rubout handler in a dependently-typed extension to Haskell. Although McBride was, of course, able to capture many system invariants by means of a rich static type system, we should point out that his implementation required almost sixty lines of code to achieve this feat.

Why `call/cc`? Our final point is the following. Many programmers think of `call/cc` and delimited control operators as being the province of effete language theorists. To which we really must reply[1]: "*Au contraire!*" Continuations are rather tools for hearty, robust systems programmers: we've used them to write a line-discipline driver to replace the one provided by the Unix kernel, and we only needed about 200 lines of code… and our version provides more functionality.

2. The Burke-Fisher functor in SML

Enough s-expressions. Let's look at some *real* syntax:

$$
\begin{array}{rcl}
\langle exp \rangle & ::= & \langle exp \rangle \texttt{ + } \langle exp \rangle \\
& | & \langle num \rangle \\
& | & \langle id \rangle \\
& | & \ldots \\
\langle decl \rangle & ::= & \texttt{val } \langle id \rangle \texttt{ = } \langle exp \rangle \texttt{ ;} \\
& | & \texttt{fun } \langle id \rangle \texttt{ (} \langle id \rangle \texttt{) = } \langle exp \rangle \texttt{ ;}
\end{array}
$$

This is just enough grammar to get us into some interesting trouble.

[1] In a hearty, robust manner, that is.

Imagine for a moment being a novice programmer sitting at the REPL. You type

```
> val f(x) = x + 1;
```

expecting to define your first function. But, of course, you are instead greeted by

```
SYNTAX ERROR: at (1:6), got '(', expected '='
```

Being a novice, you probably haven't seen the grammar above, which anyway would be much larger in practice. From the error, you only know that the first five characters drove the parser into a state from which `'='` is the only way out. If you also believe that `val` declarations are the sole way of creating bindings—or have forgotten whether it's `fun`, `fn`, `proc` or `function`—you're in for a long and frustrating REPL session.

A more helpful parser might instead respond with

```
> val f(x) = x + 1;
SYNTAX ERROR: at (1:1), did you mean 'fun'?
```

Startlingly, the *location* of the syntax error has changed in this interaction. The parser is recommending that `val` be replaced by `fun`, even though `val` is permitted by the grammar. This recommendation seems more helpful than replacing `(` with `=`, but why?

There is a general principle at work, one elucidated in a classic paper by Burke and Fisher [2]. The principle, loosely stated, is that

> *Syntax errors are usefully explained by finding minimal, nearby edits that allow the parser to substantially progress.*

Looking at our novice REPL session, we see

- If we replace `(` with `=`, the parser will encounter another syntax error almost immediately: the now-unbalanced closing parenthesis. To recover from *that* error, we would probably delete several tokens—the initial replacement leads us straight into a syntactic quagmire.
- If instead we replace `val` by `fun`, which requires *backing up* from the location where the error was detected, that single-token change results in a grammatically-correct input.

In short, the `fun` replacement is better because it explains more of the programmer's original input.

Burke and Fisher outline two requirements for "practicality":

1. Error handling should not substantially increase the space or time needed to parse *correct* input.
2. Error analysis and recovery should take constant time.

The technique outlined in their paper relies on the details of explicit-stack LL and LR parsing, and works by maintaining *two* parser states. The first parser state represents the "real" parser, and is used to detect syntax errors. The second parser state lags some fixed number k of tokens behind. When the real parser encounters an error, it can be repeatedly reverted to the state of the lagging parser as different modifications to the input are tried. The winning modification is, roughly, the smallest one yielding the greatest amount of progress beyond the original error.[2]

While ingenious, Burke-Fisher error repair is also a vivid example of too-tight coupling between error handling and core logic: the error-handling code relies on complete knowledge of the representation and algorithm of the underlying parser.

Using similar techniques to the eager REPL, we can liberate Burke-Fisher error repair from its assumptions about parsing, achieving clean separation between processing valid input and handling erroneous input.

[2] To work in constant time, a test of a repair should only venture some fixed number of tokens beyond the original error.

Plot summary

The basic trick is to slide in error correction between a parser and its source of input, a lexer. As with the eager REPL, we do this by feeding the parser an instrumented, checkpointing lexer. To keep things interesting, we switch to SML, where we can use the module system to be very clear about the separation of concerns. In the end, we write a single functor, `BurkeFisher`, that can add error repair to *any* module of signature `PARSER`.

We use two extensions to the language: delimited control and higher-order functors, both of which can be had in SML/NJ [5, 8]. Why? Read on.

Setting the scene: the signature of a parser

We start with the lexer:

```
signature LEXER =
sig
  type tok
  type tok_stream
  val lex : tok_stream -> (tok * tok_stream) option
end
```

The `LEXER` signature characterizes the source of input to a parser: a stream of tokens (type `tok`). We know nothing about the internal structure of token streams, but can observe the next token and remaining stream, if any, using `lex`.

A `PARSER` is parameterized by its `LEXER`:

```
signature PARSER =
sig
  type tok
  type result
  functor ForLexer(L : LEXER where type tok = tok):
  sig
    exception ParseError of L.tok_stream list
    val parse : L.tok_stream -> result
  end
end
```

Here is our first use of SML/NJ's higher-order module system: a module implementing `PARSER` must include a functor, `ForLexer`, which can be applied to token-compatible lexers. `ForLexer` in turn provides a `parse` function specialized to the token stream of the given lexer. Of course, all this could be replaced by an appropriate use of polymorphism, but then, so can SML's module system [11]. Explicit structuring via signatures, modules and functors allows us to *name* the separate concepts with which we wish to work.

Because `result` is defined outside the `ForLexer` functor, the result type of a parser cannot depend on its lexer.

On the other hand, the `ParseError` it raises on detecting a syntax error is parameterized by the token stream—in fact, it takes a *list* of token streams. This is the one concession a parser must make to error repair: it must signal the detection of an error by throwing an exception holding the stream just after the error was detected. If the parser performs backtracking choice, but every choice fails, it should throw the exception with a list containing *each* failure point.

A simple parser

The module `DeclRecognizer` in Figure 3 implements a recognizer for the ⟨*decl*⟩ grammar, providing an example realization of the `PARSER` signature. Being a recognizer, there are only two possible outcomes of `parse`: a unit value if successful, and an exception if not. The implementation of `DeclParser` uses a few parser combinators [6], which provide interaction with both the lexer and, ultimately, error repair.

The combinator `want` simply checks that the given token is the next one on the stream, returning an advanced stream if so and raising an exception if not:

```
fun want t = fn s =>
  case L.lex s
   of NONE => raise ParseError [s]
    | SOME (t', s') =>
        if t=t' then s'
        else raise ParseError [s']
               (* stream /after/ the error *)
```

Usually we compose recognizers sequentially, letting errors propagate:

```
infix >>
fun p >> q = fn s => q (p s)
```

To handle nonterminals like ⟨*decl*⟩, which have multiple productions, we use backtracking choice:

```
infix <|>
fun p <|> q = fn s =>
  p s handle ParseError ss =>
  q s handle ParseError ss' =>
  raise ParseError (ss @ ss')
```

With choice we see concretely the concession the parser must make to error repair: it must do a bit of work to bundle together the possible error locations.

Checkpointing a lexer

Parsers are conveniently parameterized over lexers, leaving the perfect loophole through which to checkpoint parser state. Unlike the eager REPL, we will employ *delimited* continuations for checkpointing. A delimited continuation captures the evaluation context only up to the most recent delimiter. Delimitation will allow us to tentatively run the rest of a parse with various repairs, searching for the best one, without going on to execute the rest of the program.

We use an implementation of delimited control [3], due to David Herman in an ICFP pearl [5], that is parameterized by a single *answer* type:

```
signature CONTROL =
sig
  type ans
  val shift : (('a -> ans) -> ans) -> 'a
  val reset : (unit -> ans) -> ans
end
```

Delimiters are inserted using `reset`, which expects an answer-producing thunk. On the other hand, `shift` aborts to the nearest delimiter, but reifies the evaluation context up to the delimiter as a reusable function. Thus,

```
> (reset (fn () => shift (fn k => 1) + 1)) * 2;
2
> (reset (fn () => shift (fn k => k (k 1)) + 1)) * 2;
6
```

In the first case, the captured context `[ ] + 1` is simply discarded; the `reset` is replaced by the answer `1`, which is then doubled. In the second case, the captured context `[ ] + 1` is applied, as a function, twice; the `reset` is replaced by the answer `3`, which is then doubled.

For our purposes, a single delimiter surrounding a parse will suffice. The checkpoints captured by the lexer will correspond to the remaining execution of the parser, starting from its request for a token at some point in the stream, and continuing to the point

```
structure DeclRecognizer =
struct
  datatype tok = VAL | FUN | LPAREN | RPAREN | ID | EQ | NUM | PLUS | SEMI
  type result = unit

  functor ForLexer (L : LEXER where type tok = tok) =
  struct
    exception ParseError of L.tok_stream list

    (* ... definition of combinators want, >> and <|> as given in text ... *)

    val wantExp  = (* ... *)
    val wantDecl =
         (want FUN >> want ID >> want LPAREN >> want ID >> want RPAREN >> want EQ >> wantExp >> want SEMI)
     <|> (want VAL >> want ID >> want EQ >> wantExp >> want SEMI)
    fun parse s = let val _ = wantDecl s in () end
  end
end
```

Figure 3. A combinator-style recognizer for ⟨*decl*⟩

```
fun lex (s, PASSTHRU) =
    (case L.lex s
      of NONE => NONE
       | SOME (t, s') => SOME (t, (s', PASSTHRU)))

  | lex (s, CHECKPOINT w) =
    (case L.lex s
      of NONE => NONE
       | SOME (t, s') => SOME (C.shift (fn k =>
           k (t, (s', CHECKPOINT (Window.push w                  (* first time through, yield t *)
                       (fn t' => k (t', (s', PASSTHRU))))))))) (* on checkpoint invocation, yield t' *)
```

Figure 4. The checkpointing lexer

where it produces a final answer. The type `ans` will be `P.result` for a parser `P`.

The checkpointing instrumentation is performed by a `LEXER`-transforming functor. Using SML/NJ's `funsig` form, we can give it the following signature:

```
funsig LEXER_WRAPPER (C : CONTROL) (L : LEXER) =
                     LEXER where type tok = L.tok
```

As the signature indicates, the instrumented lexer produces the same type of tokens as the original one, so it will be compatible with the same parsers. What changes is the internal representation of token streams, which includes a `wrapper` component:

```
functor WrappedLexer (C : CONTROL) (L : LEXER) =
struct
  type tok = L.tok
  type checkpoint = L.tok -> C.ans
  datatype wrapper
    = CHECKPOINT of checkpoint Window.window
    | PASSTHRU
  type tok_stream = L.tok_stream * wrapper

  fun lex (* ... see figure ... *)
end
```

An instrumented lexer can operate in one of two modes: checkpointing or passthrough. Checkpointing is used during initial parsing, until a syntax error is found. The last k checkpoints are maintained using a `window`, so that instrumentation imposes only a constant-bounded space overhead:

```
signature WINDOW =
sig
  type 'a window
  val empty : 'a window
  (* keeps only last k pushes *)
  val push  : 'a window -> 'a -> 'a window
  val list  : 'a window -> 'a list
end
```

Figure 4 shows the implementation of instrumented lexing. In passthrough mode, instrumentation has no effect. In checkpointing mode, the lexer uses `shift` to abort while capturing the continuation up to the end of parsing. The delimited continuation `k` is immediately invoked—effectively undoing the abort—yielding the same token `t` that the underlying lexer would. However, a checkpoint is pushed that, when invoked, re-runs the parser from the point of the `shift`, providing an alternative token `t'` and switching to passthrough mode.

An instrumented lexer provides `wrap` to inject an underlying lexer, `unwrap` to project the underlying lexer, and `checkpoints` to extract the last k checkpoints:

```
fun wrap s = (s, CHECKPOINT Window.empty)
fun unwrap (s, _) = s
fun checkpoints (_, CHECKPOINT w) = Window.list w
  | checkpoints (_, PASSTHRU)     = []
```

Putting it together: the Burke-Fisher functor

Nearly all the ingredients are now in place. However, to search for repairs, we'll need to know something about the available tokens:

```
signature TOK =
sig
  type tok
  val toks : tok list
  val toString : tok -> string
end
```

The list `toks` provides an instance of each token that should be used for replacement or insertion during error repair.

With that, we can specify the shape of functional Burke-Fisher error repair:

```
funsig BURKE_FISHER
         (T : TOK)
         (P : PARSER where type tok = T.tok) =
  PARSER where type tok = T.tok
```

Figure 5 gives an implementation—the `BurkeFisher` functor. To keep the presentation short, the functor is not quite true to the original algorithm:

- It only performs a single repair.
- It requires the repair to make the entire input valid, rather than choosing a repair that makes maximal (but bounded) progress. This means, in particular, that repair is not constant-time.
- It only attempts repairs of Hamming distance 1, that is, single-token replacements.

These limitations can be easily addressed by modifying the functor, without any change to the underlying parser.

When applied to a parser `P`, `BurkeFisher` produces a parser with result type `P.result * string option`. The string component is a description of the single repair (if any) performed.

Internally, the functor uses Herman's `GreatEscape` functor [5], which provides delimited control by using SML/NJ's call/cc facility. Crucially, this implementation of delimited control interacts properly with exceptions: delimited continuations captured by `shift` also capture exception handlers up to the delimiter—no more, no less. Thus, as Herman writes,

```
reset
  (fn _ =>
      (shift (fn k => (k 0)
                      handle Fail _ => 1))
    + (raise Fail "uncaught"))
handle Fail _ => 2
```

returns 1 rather than 2. Since we use exceptions to communicate parse failures, and thus exception handlers to institute repair, we rely on this behavior.

Once `BurkeFisher` is applied to a particular parser and lexer, it instantiates the parser with an instrumented version of the lexer, yielding a module `UP` (for "underlying parser"). To parse an underlying stream `us`, it wraps the stream and feeds it to the underlying parser—within a `reset`, which sets the delimiter for the checkpointing lexer. The result of the parse is paired with `NONE`, meaning no repair was necessary, which is returned if all goes well. But *outside* this pair (in particular, outside the `reset`), there is an exception handler that kicks off the repair process. Placing the handler outside the `reset` guarantees it won't be captured in the checkpoints.

AT LAST—when the underlying parser throws a syntax error, `BurkeFisher` catches it and its list of wrapped streams `wss`. Each wrapped stream represents a possible syntax error within a choice. Concatenating the checkpoints from each stream, the repairing `parse` uses `tryCPs` to attempt a repair at each checkpoint in turn. If successful, `tryCPs` returns a pair of the `result` from the underlying parser and the replacement token it used to get that successful parse. Otherwise `tryCPs` returns `NONE`, and `parse` re-raises the parse error using the underlying lexer streams.

Supposing we've written modules `Tok` and `SimpleLexer` (which represents streams as lists of tokens), we can sit down at the REPL and see

```
> structure RP  = BurkeFisher(Tok)(DeclRecognizer);
> structure RPP = RP.ForLexer(SimpleLexer);
> RPP.parse [VAL,ID,LPAREN,ID,RPAREN,EQ,
             NUM,PLUS,NUM,SEMI];
((), SOME "Did you mean 'fun'?")
```

A caveat: choice points are commit points

There's a subtlety regarding choice: if we factor out " = ⟨*exp*⟩ ; ", which is common to both forms of ⟨*decl*⟩, the repair no longer works.

The problem is this: once a choice has been successfully parsed, only the checkpoints of the taken branch are retained. The effective result is that repair checkpoints do not always have access to all possible branches.

We used exceptions carrying a list of streams to join together the checkpoints of an unsuccessful choice. To deal with successful choices, we need to maintain all checkpoints even when backtracking—we need the backtracking control structure to play nicely with the checkpointing control structure. An easy and fairly pleasing way to do this is to parameterize the underlying parser by an implementation of backtracking choice—which `BurkeFisher` can then provide—which also allows `ParseError` exceptions to carry just a single stream:

```
fun choose p q = fn (us, w) =>
  p s handle ParseError (us', w') =>
  q (us, w')
```

Discussion

The big idea As with the eager REPL, the key benefit of control operators is the ability to hook into—even manipulate—the computation of a parser, without any understanding of how that computation is structured.

There are also two significant differences from the eager REPL. First, we assume that parsing and lexing are free from observable side effects, so we do not use effect logging (and, in fact, avoid state—see below). Thus, we are *only* logging control state. Second, by having a modicum of knowledge about the structure of input to the parser (via the `Tok` signature), we are able to intelligently explore the space of input repairs, while keeping the computational structure of the parser abstract.

Look Ma, no refs! Contra the eager Scheme REPL, we haven't used mutable state in implementing Burke-Fisher repair. We were able to avoid it because the `PARSER` signature requires parsers to use their lexers *functionally*—in particular, to thread the token stream through the parsing processes. The downside to this approach is that checkpointing is sensitive to the way the parser threads the token stream, which as we noted can cause problems when the parser itself backtracks.

But, of course, an imperative interface is also workable, and it has the benefit that *every* use of the lexer is checkpointed, regardless of how the token stream is threaded. What's more, the token streams can be dropped from the `ParseError` exception, further shrinking the interface between the parser and the repair functor.

```
functor BurkeFisher (T : TOK)
                    (P : PARSER where
                           type tok = T.tok) =
struct
  type tok = T.tok
  type repair = string
  type result = P.result * string option

  (* Herman's delimited control from call/cc *)
  structure C = GreatEscape(type ans = P.result)

  functor ForLexer (L : LEXER
                    where type tok = tok) =
  struct
    exception ParseError of L.tok_stream list
    structure WL = WrappedLexer (C) (L)
    structure UP = P.ForLexer (WL)

    fun tryReps k []      = NONE
      | tryReps k (t::ts) = SOME (k t, t)
          handle UP.ParseError _ => tryReps k ts

    fun tryCPs []      = NONE
      | tryCPs (k::ks) = case tryReps k T.toks
                          of NONE    => tryCPs ks
                           | SOME rt => SOME rt

    fun parse us =
      (C.reset (fn () =>
         UP.parse (WL.wrap us)), NONE)
      handle UP.ParseError wss =>
        case tryCPs (List.concat
                      (map WL.checkpoints wss))
         of NONE =>
              raise ParseError (map WL.unwrap wss)
          | SOME (r, t) =>
              (r, SOME (concat ["Did you mean ",
                                T.toString t,
                                "?\n"]))
  end
end
```

Figure 5. The Burke-Fisher functor

To type-checking, and beyond Having generalized Burke-Fisher repair to arbitrary parsers, it's natural to wonder if we can go even farther. By encapsulating the notion of input streams and local repairs in a separate module taken as an extra parameter, the `BurkeFisher` functor could be pared down to handling just the checkpointing and error-catching process. We suspect, for example, that Lerner *et al.*'s SEMINAL tool [9], which attempts to explain type errors by searching for repairs, could be restructured to use Burke-Fisher-style local search—becoming just one more instantiation of the functor.

The real deal The `BurkeFisher` functor came out of work done building new lexing and parsing infrastructure for SML/NJ; it is in use in production code, and available with recent versions of the compiler.[3] The implementation includes the full suite of Burke-Fisher repairs, allows for multiple repairs, and judges goodness of repairs using a sophisticated metric.

[3] See `http://smlnj.org/`

Acknowledgements

This article reports on work supported by the Defense Advanced Research Projects Agency under Air Force Research Laboratory (AFRL/Rome) Contract No. FA8650-10-C-7090 and Cooperative Agreement No. FA8750-10-2-0233. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

Turon is currently supported by a Microsoft fellowship, and was supported by NSF award CNS-0454136 while writing the ML-ANTLR parser at the University of Chicago.

John Reppy provided much guidance. Matthias Felleisen provided encouragement and insightful comments to Shivers while he was working out the rubout-handler system; Mark Feeley provided Shivers an early opportunity to air the design at a Scheme workshop. The quotation on line 31 of Figure 1 is due to a prince in Denmark. Kent Pitman and Dan Weinreb were fonts of wisdom concerning sophisticated details of classic LISP technology. Thanks go to Vincent St-Amour, Sam Tobin-Hochstadt, and Jesse Tov for feedback on a draft of the paper.

Shivers would like to dedicate this set of twin pearls to another such set: Olin and Avery.

References

[1] E. Biagioni, K. Cline, P. Lee, C. Okasaki, and C. Stone. Safe-for-space threads in Standard ML. *Higher Order and Symbolic Computation*, 11: 209–225, Sept. 1998.

[2] M. G. Burke and G. A. Fisher. A practical method for LR and LL syntactic error diagnosis and recovery. *ACM Transactions on Programming Languages and Systems*, 9(2):164–197, Mar. 1987.

[3] O. Danvy and A. Filinski. Abstracting control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming (LFP'90)*, pages 151–160, 1990.

[4] M. Felleisen. The theory and practice of first-class prompts. In *Conference Record of POPL 1988: The Fifteenth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 180–190, 1988.

[5] D. Herman. Functional pearl: The great escape, or how to jump the border without getting caught. In *Proceedings of the Twelfth ACM SIGPLAN International Conference on Functional Programming (ICFP'07)*, 2007.

[6] G. Hutton and E. Meijer. Monadic parsing in Haskell. *Journal of Functional Programming*, 8(4):437–444, July 1998.

[7] R. Kelsey and J. Rees. The Scheme48 system. `http://s48.org`.

[8] G. Kuan. *A true higher-order module system*. PhD thesis, University of Chicago, 2011.

[9] B. S. Lerner, M. Flower, D. Grossman, and C. Chambers. Searching for type-error messages. In *Proceedings of the 2007 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'07)*, 2007.

[10] K. M. Pitman. Ambitious evaluation: a new reading of an old issue. *Lisp Pointers*, VIII(2), May 1995.

[11] A. Rossberg, C. Russo, and D. Dreyer. F-ing modules. In *Proceedings of the ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI'10)*, 2010.

[12] O. Shivers, B. Carlstrom, M. Gasbichler, and M. Sperber. The scsh manual, release 0.6.6. `http://scsh.net`, Mar. 2004.

[13] D. Sitaram and M. Felleisen. Control delimiters and their hierarchies. *Lisp and Symbolic Computation*, 3:67–99, May 1990.