

User-extensible sequences in Common Lisp

Christophe Rhodes
Goldsmiths, University of London
New Cross Road, London, SE14 6NW
c.rhodes@gold.ac.uk

ABSTRACT

Common Lisp is often touted as the programmable programming language, yet it sometimes places large barriers in the way, with the best of intentions. One of those barriers is a limit to the extensibility by the user of certain core language constructs, such as the ability to define subclasses of built in classes usable with standard functions: even where this could be achievable with minimal penalties. We introduce the notion of user-extensible sequences, describing a protocol which implementations of such classes should follow. We show examples of their use, and discuss the issues observed in providing support for this protocol in a Common Lisp, including ensuring that there is no performance impact from its inclusion.

1. INTRODUCTION

Common Lisp is recognized as being an extremely flexible language: one in which linguistic experimentation can take place, one where the method of solving a problem in a domain is first to write an interpreter or a compiler for a language to express concepts in the domain of interest, and then to use that domain-specific language to solve the problem. To some extent, then, it might be surprising to find some of the ways in which extension of Common Lisp (by the implementor or user) is limited. This is at least partially explained by the aims of the standardizers, which included “stricter standardization for portability” (Pitman and Chapman, 1994, section 1.1.2): codification of existing practice was a large part of the X3J13 committee’s work, and it is difficult to ensure portability with a highly-extensible language core.

However, the standardization process was not intended to close the door to language development: merely to provide a stable platform which could be agreed on. For example, the CLOS system for object orientation was standardized without very much scope for extensibility; however, it was intentionally (Steele, 1990, chapter 28) upwardly compatible with something close to the Metaobject Protocol described

in Kiczales et al. (1991), often (though not always) supported by contemporary Lisp implementations.

Meanwhile, other languages and language environments have not stood still; many of Lisp’s once-unique features are now to be found in other languages (Norvig, 2002), though some are still not¹; additionally, sometimes these languages have features not found in any available Lisp implementation. In some cases, such as Aspect-Oriented Programming (Kiczales et al., 1997) these features can be straightforwardly implemented by any interested user – this itself is one of Lisp’s unique features not often found elsewhere – but some features need implementation support for them to be used to maximum effect: addition of features by the Lisp programmer can suffice for certain needs, but they are not pervasive in the way that one would want; such extensions, even upwardly-compatible extensions, need to be explicitly used by third-party code.

One such extension is the ability for the user to define new sequence types. For the purposes of this paper, a sequence is a finite ordered collection; more concretely, a sequence has a number of elements, addressable by an integer between 0 (inclusive) and the sequence’s length (exclusive). The motivation for allowing the definition of new sequence types is primarily to eliminate the need for new names for essentially the same concept: examples of such from real-world code include `climacs-buffer:size` and `flexichain:nbelements` (Strandh et al., 2004) as analogues for `length` on user-defined data structures, and `tabcode-syntax:buffer-position-if` (Rhodes et al., 2005) for `position-if`; these multiple names are a barrier to clear expression and understanding of Common Lisp code which defines objects which are conceptually sequences. Additionally, the implementation of sequence functionality is tedious and error-prone², and it is more convenient to require of the datastructure implementor only a few simple method definitions to achieve all the standard sequence functionality.

We describe in this paper a proposal for implementations to allow users to define new sequence classes which interoperate seamlessly with the standard Common Lisp sequence functions. The basic mode of operation of our design is as

¹and it has been argued that their inclusion would convert their host language into a Lisp of some form.

²The author, in developing this extension, wrote the implementation first and test cases second: the test cases revealed five implementation errors.

```

(defun foo (x)
  (handler-case
    (etypecase (ignore-errors (make-sequence x 42))
      (null ...) ; make-sequence threw an error
      (list ...)
      (vector ...))
    (type-error (c)
      ;; we get here if make-sequence returned a non-list non-vector
      (error "BUG: system threw ~S on ~S" c 'make-sequence))))

```

Figure 1: How to detect formally nonconforming implementations of `make-sequence`. If calling `foo` with any argument causes the `type-error` handler to be called, the Common Lisp implementation is, strictly speaking, nonconforming.

follows: implementors of sequence classes write methods on generic functions named by symbols in the `sequence` package; users of sequence classes call the standard functions named by symbols in the `common-lisp` package: it is intended that sufficiently generic existing code may be run unmodified using extended sequences. The implementor of this proposal is responsible for arranging that the calls to the standard functions invoke applicable methods on the appropriate functions in the `sequence` package. Further, users wishing to construct new functions involving iteration over sequences should call the iteration operators included in this proposal, so that their new functions will work on user-defined sequence classes.

The rest of this paper is organized as follows: we introduce the design considerations of this proposal in section 1.1; then, after discussing compatibility issues and related work in sections 1.2 and 1.3, we give an introduction for the prospective user of extensible sequences in section 2, and some examples in section 3, and some details of our implementation are given in section 4. As a snapshot of our work in progress, a more formal specification of the protocols is presented in appendix A, intended in the first instance to stimulate discussion rather than to be the definitive specification.

1.1 Design considerations

Our design for this extension as a whole was influenced by a number of desiderata:

- We wish to minimize incompatibility with the existing Common Lisp standard, introducing such incompatibility only where strictly necessary. There are only a few corner cases where the existing standard document needs to be clarified or reinterpreted; see section 1.2 for discussion of this point.
- We wish to make defining new sequence classes simple and convenient for users, so that simply defining a subclass of `sequence` and a few methods specialized on that class suffices for the entire sequence functionality to be available, and also to provide enough hooks for sequence class implementors to be able to customize the behaviour for their own needs.
- We wish not to overly constrain implementations of this protocol, allowing implementors to make trade-offs to reflect the needs of their users. In particular, we wish to allow both a simple implementation (where

the Common Lisp implementation supports this protocol natively for all sequence types) and an implementation where existing code suffers no run-time performance penalty; additionally, we do not wish to render a portable implementation impossible, though we believe that such an implementation would necessarily have visible seams.

While it is possible to recommend ways of accessing and iterating over both Common Lisp sequences and other objects in `userspace`³, there is no way of having those other objects (conceptually sequences though not of type `sequence`) seamlessly interoperate with the standard Common Lisp sequence functionality, or with third-party code which does not follow such a recommendation.

The desire not to constrain implementations of this proposal leads to the creation of a new package name to name the operators providing the extensible functionality, allowing dispatch from the standard function to the extensible one through trampolining or `defadvice`-like mechanisms. However, an implementation of Common Lisp wishing to support this proposal natively may choose to implement some of its CL symbols as direct imports from the new `sequence` package.

In what follows, there will be reference to new operators without explicit package prefixes: in such cases, the symbols' package should be assumed to be the `sequence` package. A number of new operators have the same name as standardized operators of Common Lisp; except in appendix A, the operator without an explicit prefix should be taken to mean the Common Lisp operator, while the new operators will have an explicit `sequence:` prefix.

1.2 Compatibility

The types `[sic]` `vector` and the type `list` are disjoint subtypes of type `sequence`, but are not necessarily an exhaustive partition of `sequence`.

Pitman and Chapman (1994, System Class `sequence`)

Although the quote above might suggest that there would be no problem from the point of view of formal standard con-

³By 'userspace' in this paper, we mean the realm of conforming or de-facto acceptable code written by users of an implementation, as opposed to code written by implementors themselves to provide an implementation with extensions.

<code>sequence:length</code>	<code>sequence:elt</code>	<code>(setf sequence:elt)</code>
<code>sequence:adjust-sequence</code>	<code>sequence:make-sequence-like</code>	

Table 1: The protocol functions which must be implemented for a sequence class. Three correspond directly to primitive sequence operations in Common Lisp; `adjust-sequence` and `make-sequence-like` provide support for components of higher-level destructive and consing operations respectively (see table 2).

formance for an implementation to offer non-standard types of `sequence`, the outline code in figure 1, as a result of Issue `CONCATENATE-SEQUENCE` (Pitman, 1991), is defined by ANSI CL never to get to the `type-error` branch of the `handler-case`, which was probably not intended by the standardizers. For a fuller discussion of this issue, and a suggestion for a resolution, see Rhodes (2006).

An additional issue with formal compliance is that ANSI CL specifies for the `length` function and the `elt` accessor that their sequence argument should be a *proper sequence*, which is defined in the normative glossary as *a sequence which is not an improper list; that is, a vector or a proper list*; again, this definition was probably not intended to prohibit non-standardized sequences, even though, interpreted strictly, the second half of it implies that no non-standard sequence is a proper sequence.

We must also consider the issue of cultural compatibility with the body of Common Lisp code available in the wild. Perhaps because of the historical lack of generic sequences offered by Common Lisp implementations, there seems to be little in the way of packages written to exploit the abstract sequence data type; instead, packages choose a concrete sequence type and implement their functionality atop that, manipulating a particular kind of sequence rather than sequences in general⁴. However, it is likely that if user-extensible sequences become available, code will be modified or written afresh to take advantage of it.

It would be possible to implement this proposal (or something like it) entirely portably (in ‘userspace’), by defining a new package shadowing many of the standard Common Lisp functions and macros, and implementing the generic functionality on the standard types by trampolining to the standard functions. However, such an implementation strategy has drawbacks:

- Implementation-specific compiler optimizations (such as compiler macros) for the shadowed functions would essentially be lost. As an example, if an implementation has a specialized implementation of `map-into` for arguments known at compile-time to be certain types of vector, that optimization will go unused in a putative userspace implementation in function calls to `sequence-cl:map-into`.
- The userspace implementation of extensible sequences would not interoperate with third-party Common Lisp code: programs already written with generic (not nec-

essarily user-extensible) sequences in mind, using the `common-lisp` package, will not be able to interoperate cleanly with sequences defined using this portable implementation.

- Innocuous-looking uses of explicit package prefixes (for example, `cl:sequence`) would have surprising and potentially difficult-to-debug effects. Other maintenance headaches include how to support both native and portable implementations in library code.

A userspace implementation might be better than nothing, for a transition period, but support from the Lisp implementation is required for seamless operation.

That said, there is an issue regarding portability of libraries using extensible sequences: until it is ubiquitously implemented, use of this facility in bodies of code render those bodies of code unportable in practice. Whether this unportability itself becomes a problem in practice is largely a matter for the user community to decide.

1.3 Related Work

Much inspiration for this proposal was drawn from the Dylan (Shalit, 1996) iterator protocol, which provides for iteration over collections with the open generic function `forward-iteration-protocol`. Where our design is similar to this protocol, it is largely for the same reasons: we do not wish to impose unnecessary run-time overhead (in space or speed) for those uses which need high performance. However, in this proposal, we aim to provide a little more of a layer of convenience for the user who does not need to minimize overhead.

Many other languages provide some form of iteration or collection protocol. Python, for instance, allows the implementor of a collection to define a method on `__iter__()` to return an iterator object, which itself must have methods on `__iter__()` and `next()`. Apparently for reasons of efficiency, the Python iterator protocol (Yee and van Rossum, 2001) provides no explicit means for checking for termination of an iteration, instead requiring the iterator to signal an exception when `next()` is called on an iterator representing a terminated iteration.

The Scheme language (Kelsey et al., 1998) definition has little built-in support for user-extensible or even generic sequences. Its community has made one attempt (Miller, 2004) at defining an interface and conventions for collections; however, this SRFI has apparently not seen many Scheme implementations decide to support it natively; as of two years after its finalization, the SRFI status page reports no implementation as supporting it.

⁴There are exceptions: for instance, the `split-sequence` library, designed by readers of `news:comp.lang.lisp` in 2001, works without modification on generic sequences of the form described in this paper.

Common Lisp Function	Extensible Generic Function
copy-seq, subseq	sequence:copy-seq, sequence:subseq
reduce, fill	sequence:reduce, sequence:fill
reverse, nreverse	sequence:reverse, sequence:nreverse
sort, stable-sort	sequence:sort, sequence:stable-sort
count, find, position	sequence:count, sequence:find, sequence:position
count-if, count-if-not	sequence:count-if, sequence:count-if-not
find-if, find-if-not	sequence:find-if, sequence:find-if-not
position-if, position-if-not	sequence:position-if, sequence:position-if-not
search, mismatch, replace	sequence:search, sequence:mismatch, sequence:replace
substitute, nsubstitute	sequence:substitute, sequence:nsubstitute
substitute-if, substitute-if-not	sequence:substitute-if, sequence:substitute-if-not
nsubstitute-if, nsubstitute-if-not	sequence:nsubstitute-if, sequence:nsubstitute-if-not
remove, delete	sequence:remove, sequence:delete
remove-if, remove-if-not	sequence:remove-if, sequence:remove-if-not
delete-if, delete-if-not	sequence:delete-if, sequence:delete-if-not
remove-duplicates, delete-duplicates	sequence:remove-duplicates, sequence:delete-duplicates

Table 2: Common Lisp functions, and the corresponding generic functions specified to be extensible in this protocol. Calling the Common Lisp function with an extended sequence as an argument will cause applicable methods on the extensible generic function to be invoked; whether this is achieved through defadvice, trampolining or function identity is left unspecified.

The Factor language’s sequence interface is conceptually very similar to that described in the following sections for Common Lisp, including the distinction between sequence protocol (defining the fundamental operations that must be implemented for sequence objects) and ‘utility words’, analogous to functions performing computations over sequences (Pestov, 2006, Sequences), which will work on any object implementing the sequence protocol.

In the Common Lisp world itself, an early (pre-CLOS) attempt to provide generic sequence functionality was presented in Haible (1988) for the GNU CLISP implementation; however, that proposal was never formally exported or documented (Haible, 2006); the author’s primary concern with this attempt was in supporting the necessary operations efficiently. Some of these issues of efficiency remain in this proposal, though we believe that they might be resolved in a Lisp implementation supporting sealing and inlining of methods.

2. USER-DEFINED SEQUENCES

In this section, we must draw the distinction between the implementor of a sequence class, and the implementation of Common Lisp which supports this user-extensible sequence facility. Most of the time, we will use ‘user’ to mean the implementor of a sequence class, and ‘implementor’ to mean implementor of a Common Lisp implementation, and we hope that it will be clear from context when this does not apply.

The names of the various operators have been chosen to maximize both backward- and forward-compatibility (with Common Lisp as standardized, and with potential related extensions to Common Lisp such as a collections protocol); we specify operators corresponding to standardized functions such as `find` to be named like `sequence:find`, so that an implementation of Common Lisp ‘natively’ implementing this proposal can simply import the relevant symbols from the `sequence` package⁵. Thus, we specify `sequence:adjust-`

⁵According to some interpretations of the standard, the im-

plementation may not simply make `sequence` a nickname for the `common-lisp` package, as the list of package nicknames is standardized as `cl` only.

`sequence` rather than (`setf sequence:length`), so that an implementation can import `sequence:length` into the standard `common-lisp` package without inadvertently causing there to be a `setf` function for `length`, in contravention of the constraints on implementations (Pitman and Chapman, 1994, Section 11.1.2.1.1).

We also aim to be not incompatible with similar extensions to Common Lisp, such as a protocol for accessing and iterating over general collections; in this proposal we are dealing with user-defined sequence classes because Common Lisp as standardized has a large library of functions acting on generic sequences, while it has no functions acting on generic collections – and so a ‘userspace’ implementation of a collections protocol would not pose the interoperability problems discussed in section 1.2.

2.1 Sequence Datatypes

Under this proposal, the user may define a direct subclass of `sequence` using `defclass`, specifying `sequence`⁶ as one of the superclasses (but not the only one: for code portable between implementations of this proposal, `standard-object` must be in the superclasses list too). The resulting class is subtype `sequence`, and instances of the class are `typep sequence`.

The fundamental operations defined in Common Lisp relating to sequences are `length`, `elt`, and (`setf elt`); in order to support these, the user of the extensible sequence facility (the implementor of the sequence class) should define methods on `sequence:length`, `sequence:elt` and (`setf sequence:elt`).

plementation may not simply make `sequence` a nickname for the `common-lisp` package, as the list of package nicknames is standardized as `cl` only.

⁶Reviewers of this paper have raised concerns about the use of `sequence` as a direct superclass, suggesting instead a protocol class, maybe `sequence-object`, as the user-specializeable class. The author wishes to acknowledge reviewer concerns in this matter, while continuing to present the protocol as currently implemented.

The set of functions making up the sequence protocol of this proposal (see table 1) consists of the analogues to the three Common Lisp functions above, along with two others: `sequence:adjust-sequence`, which is similar to the `adjust-array` function, but generalized to sequences; and `sequence:make-sequence-like`, which creates a sequence of the the same kind⁷ as its first argument, with length given by its second argument.

The rationale for `make-sequence-like` is to support the `subseq` and `copy-seq` operators, and the sequence functions which are defined to return a freshly-constructed sequence (such as `substitute`). Additionally, this operator is easier both to specify and to implement efficiently than one like the `make-sequence` function, which requires a full understanding of the Common Lisp type system; `make-sequence-like` can operate as a simple generic function specialized on its sequence argument.

To support the `delete` and `delete-duplicates` functions, we provide `adjust-sequence`, which adjusts various properties, including the length, of a sequence. Users of this protocol are encouraged to document the circumstances under which their methods on this operator will preserve the identity of its sequence argument, and must arrange that if a new sequence is allocated, the original sequence argument must not be modified; however, in general, callers (either explicit or implicit) may not assume that the sequence returned from `adjust-sequence` is the same sequence as its argument.

Once methods for these generic functions have been implemented for a sequence class, all the regular Common Lisp sequence functionality will work as expected. However, some sequences may not admit implementations of all these operations, or indeed might offer means to implement certain operations more efficiently, so in the following sections we describe the protocol by which the Lisp implementation provides the sequence functionality.

2.2 Iteration

It is common to iterate over sequences, both in the conception of many of the Common Lisp sequence functions, and in user-defined operations. In this section, we describe the protocol and interface for iterating over sequence contents, both built-in and user-defined. The protocol discussed here has a set of default methods specialized on the `sequence` class, so that any sequence class for which methods on the protocol functions of section 2.1 have been implemented will obey the protocol. However, these default methods cannot take into account the characteristics of any particular sequence class, and so for efficiency users may wish to override them for their own classes.

The essential concept is the use of an iterator object to represent the current state of an iteration. It is not necessary for this iterator object to have a distinguished class, as the sequence over which it is iterating is present in all function calls in this protocol (and therefore its class can be used for specialization of methods); indeed, the conceptual object is

⁷Neither ‘type’ nor ‘class’ is quite right here, as `cons` and `null` are distinct types and classes, while both being subtypes of the sequence type `list`.

represented in this protocol by three objects and six functions. This iteration protocol is close to the one available in Dylan (Shalit, 1996).

The `make-sequence-iterator` operator constructs one of these iterator objects: after the required sequence argument, it accepts keyword `:start`, `:end` and `:from-end` arguments, and returns nine values. The first three of those values are an iterator state, a limit and the from-end argument; the remaining six are functions which, respectively, return a state one step ahead, if possible; test the state against the limit for termination; retrieve the element at the current iteration state from the sequence; set the element at the current iteration state to a new value; return the index corresponding to the current iteration state; and return a distinct iteration state representing a copy of the current one.

The default method on `make-sequence-iterator` is intended for convenience: for most uses, it is unnecessary to construct the nine return values; instead, the default method (specialized to `sequence`) on `make-sequence-iterator` generates the first three of the return values by calling `make-simple-sequence-iterator`, and returns in addition six protocol generic functions: `#'iterator-step` (which advances the iteration state); `#'iterator-endp`, testing an iteration for termination; the `#'iterator-element` reader and the `#'(setf iterator-element)` writer; `#'iterator-index`, returning the sequence index corresponding to the iterator state; and finally `#'iterator-copy`, returning a copy of the iteration state. These functions have methods specialized to `list` and `vector` to provide iterators for the built-in sequence classes.

While implementors of sequence classes may choose to use this CLOS-based iterator protocol (at the potential loss of efficiency through generic function dispatch at each step), users of the iteration protocol (who define functions which perform iterations over sequences) may not assume that the sequence class implementor has done so, and so must call `make-sequence-iterator` or the operators discussed below.

A small dose of syntactic sugar around `make-sequence-iterator` is provided by `with-sequence-iterator`, which binds as if by `multiple-value-bind` the variables in its first argument to the result of applying `make-sequence-iterator` to its second argument, and then executes the body. A slightly simpler macro to use is `with-sequence-iterator-functions`, which binds the six names in its first argument to six local functions (which have dynamic extent and close over the return values from `make-sequence-iterator`) which perform the various iterator manipulations.

For programmer convenience, we also provide a `dosequence` macro, behaving as `dolist` (but for arbitrary sequences), and an extension for `loop` using the loop keywords `element` and `elements`, in a similar fashion to the `for-as-package` loop path (Pitman and Chapman, 1994, Section 6.1.2.1.7).

It is intended that the iteration protocol described here will be compatible with an iteration protocol for general collections (including hash tables, trees and other similar data structures); at the minimum, any such protocol should be

some	every	notany
notevery	map	map-into
concatenate	merge	
coerce	make-sequence	

Table 3: Common Lisp functions applicable or otherwise relevant to sequences which are not specified as extensible in this protocol; however, the implementation’s versions are specified to be applicable to arbitrary sequence types, and to produce the expected result.

able to specify its behaviour for sequences in terms of the operators described here.

2.3 Sequence Functions

For most of the functions in the sequences chapter of the standard, there is an analogous generic function with the same argument list which can be specialized for user-defined sequence classes. The complete list of generic functions specified to be extensible, and their corresponding Common Lisp function, is given in table 2. Implementations of this proposal may choose to make the Common Lisp function `eq1` to the extensible generic function, or even make the symbols themselves `eq1` to each other; however, users may not assume this.

2.4 Other Affected Functions

The Common Lisp functions in table 3 are not extensible in this proposal. We discuss the interaction of this proposal with the behaviour of those functions below; in summary, where ANSI CL specifies that they accept sequence arguments, the implementation must accept any user-defined sequence, and where ANSI CL specifies that they accept a sequence type-specifier, they must accept a user-defined sequence name or class object (see also Rhodes (2006) for a discussion of the implications of this with respect to the standard behaviour).

Six functions in the SEQUENCES chapter of the ANSI CL standard do not have the structure for user-extensibility in the same straightforward way as the functions discussed in section 2.3: they are `concatenate`, `map`, `merge`, `make-sequence`, `coerce` and `map-into`.

Taking `map-into` first, we simply specify that an implementation of user-extensible sequences must implement the `map-into` function such that target sequences and source sequences of arbitrary subtype of `sequence` are supported given a complete implementation of the sequence and iteration protocol of sections 2.1 and 2.2. An implementation is not prohibited from providing a means of customizing the behaviour of `map-into`, but neither is it required to.

A means of extending the behaviour of `concatenate`, `map`, `merge`, `make-sequence` and `coerce` is likewise left unspecified, though again an implementation of this document must provide for the functionality of these functions to be available for objects of arbitrary sequence type and type specifiers naming a concrete subtype of `sequence`; for details of how this interacts with the ANSI CL specified behaviour of

these functions, see Rhodes (2006). It is likely that any documented fashion of extending these five functions will use the same mechanism for all of them.

As for functions which are not in the SEQUENCES chapter of the ANSI CL specification, the intent is that where a particular behaviour is specified for a sequence argument, that behaviour should be implemented for the user-extensible sequences. For instance, the short-circuiting quantifiers `some`, `every`, `notevery` and `notany` should accept sequence arguments of arbitrary class; `make-array`’s `:initial-contents` argument (and the `#a` array reader) should accept a sequence of sequences, as specified, including user-extended sequences. Note that we do not specify a different behaviour of the `equalp` function from the standard-mandated one (of using `eq` on objects of type `standard-object`): although `equalp` is specified to descend arrays and conses, it does not do so in a sequence-like way, and in particular a list and a vector with the same content are not considered to be `equalp`. Equality over sequence contents can be computed with `mismatch`.

Particular attention must be paid to the `read-sequence` and `write-sequence` functions. Although Gray (1989) does not suggest that these functions be extensible in the manner of the other stream functions, it is likely that this is because `read-sequence` and `write-sequence` were added to the language after the extensible streams proposal was made, and indeed Common Lisp implementations have provided `stream-read-sequence` and `stream-write-sequence` generic functions for user customizeability. Lisps providing both an implementation of this proposal and extensible streams based on Gray (1989) should document the effect of calling `read-sequence` and `write-sequence` with arguments being instances of non-standardized classes. The simplestreams (Franz Inc., 2006, Chapter 83) extensible analogue is documented to signal an error for all sequence types other than strings and octet vectors; if this restriction is relaxed, there are similar interactions as with Gray streams over which extensibility (sequence or stream) takes precedence.

3. EXAMPLES

In this section, we present two examples of uses for the protocols discussed in this paper. Firstly, in section 3.1, we present the implementation of a distinct sequence type along with the method definitions to allow it to interoperate; then, in section 3.2, we demonstrate the definition of a mixin class and methods specialized on it to provide additional functionality to generic sequences implemented as in this paper.

3.1 Queue

As an example of a non-standard sequence, consider implementing a queue data structure, which supports the operations `enqueue` and `dequeue`. Figure 2 shows one way in which such a queue could be implemented, using a list as the storage for the data, and keeping a reference to the cons cell at the back of the queue. Purely for interest, we also implement a variant of the queue which is also funcallable, and arrange so that calling the funcallable queue with no arguments performs dequeuing, while with one argument it enqueues that argument.

An implementation of the sequence and iteration protocol

```

(defclass queue (sequence standard-object)
  ((%data :accessor %queue-data) (%pointer :accessor %queue-pointer)))
(defmethod initialize-instance :after ((o queue) &key)
  (let ((head (list nil)))
    (setf (%queue-data o) head (%queue-pointer o) head)))

(defgeneric enqueue (data queue)
  (:argument-precedence-order queue data)
  (:method (data (o queue))
    (setf (cdr (%queue-pointer o)) (list data) (%queue-pointer o) (cdr (%queue-pointer o))
    o))
(defgeneric dequeue (queue)
  (:method ((o queue))
    (progn1 (cadr (%queue-data o))
      (setf (cdr (%queue-data o)) (cddr (%queue-data o))))))

(defclass funcallable-queue (queue funcallable-standard-object)
  () (:metaclass funcallable-standard-class))
(defmethod initialize-instance :after ((o funcallable-queue) &key)
  (flet ((fun (&optional (new nil new-p)) (if new-p (enqueue new o) (dequeue o))))
    (set-funcallable-instance-function o #'fun)))

```

Figure 2: Basic definitions of a queue data structure, including a funcallable variant. The layout of the code in this and subsequent figures is not idiomatic for reasons of space.

```

(defmethod sequence:length ((o queue)) (length (cdr (%queue-data o))))
(defmethod sequence:elt ((o queue) index) (elt (cdr (%queue-data o)) index))
(defmethod (setf sequence:elt) (new-value (o queue) index) (setf (elt (cdr (%queue-data o)) index) new-value))
(defmethod sequence:make-sequence-like ((o queue) length &key (initial-element nil iep) (initial-contents nil icp))
  (let ((result (make-instance (class-of o))))
    (cond
      ((and iep icp)
        (error "supplied both ~S and ~S to ~S" :initial-element :initial-contents 'make-sequence-like))
      (icp (unless (= (length initial-contents) length)
        (error "length mismatch in ~S" 'make-sequence-like))
        (setf (cdr (%queue-data result)) (coerce initial-contents 'list)
          (%queue-pointer result) (last (%queue-data result))))
      (t (setf (cdr (%queue-data result)) (make-list length :initial-element initial-element)
        (%queue-pointer result) (last (%queue-data result))))
    result))
(defmethod sequence:adjust-sequence ((o queue) length &key initial-element (initial-contents nil icp))
  (cond
    ((= length 0)
      (setf (cdr (%queue-data o)) nil (%queue-pointer o) (%queue-data o)))
    (t (sequence:adjust-sequence (%queue-data o) (1+ length) :initial-element initial-element)
      (setf (%queue-pointer o) (last (%queue-data o)))
      (when icp (replace (%queue-data o) initial-contents :start1 1) o)))

(defmethod sequence:make-simple-sequence-iterator ((q queue) &rest args &key from-end start end)
  (declare (ignore from-end start end))
  (apply #'sequence:make-simple-sequence-iterator (cdr (%queue-data q)) args))
(defmethod sequence:iterator-step ((q queue) iterator from-end)
  (sequence:iterator-step (cdr (%queue-data q)) iterator from-end))
(defmethod sequence:iterator-endp ((q queue) iterator limit from-end)
  (sequence:iterator-endp (cdr (%queue-data q)) iterator limit from-end))
(defmethod sequence:iterator-element ((q queue) iterator)
  (sequence:iterator-element (cdr (%queue-data q)) iterator))
(defmethod (setf sequence:iterator-element) (new-value (q queue) iterator)
  (setf (sequence:iterator-element (cdr (%queue-data q)) iterator) new-value))
(defmethod sequence:iterator-index ((q queue) iterator)
  (sequence:iterator-index (cdr (%queue-data q)) iterator))
(defmethod sequence:iterator-copy ((q queue) iterator)
  (sequence:iterator-copy (cdr (%queue-data q)) iterator))

```

Figure 3: Implementation of the sequence protocol for the queues of figure 2. The first five method definitions are sufficient to allow all sequence functionality to run correctly; the methods on the iterator functions reduce the complexity for our simple queue implementation from $O(N^2)$ to $O(N)$ for most operations.

```
(length #[a b c d]) ; => 4
(count 1 #[1 2 3]) ; => 1
(remove-if-not #'oddp #[1 2 3]) ; => #[1 3]
(remove-duplicates #[1 2 3 4 5] :end 4
                  :key #'oddp :from-end t)
; => #[1 2 5]
#2a#[#[1] #(2)] ; => #2A((1) (2))
```

Figure 4: Examples of using the queues as sequences. For the `#[` and `#{` reader macros and queue print functions, see figure 7 in the appendix.

for a queue implemented in this fashion is shown in figure 3. The first five method definitions are required under this proposal, for implementing the primitive sequence operations. The default method on `make-sequence-iterator` would suffice to make all Common Lisp sequence functionality work; however, an index-based iteration scheme as provided by that default method would suffer from poor performance for our implementation of queues based on lists, so we define methods for the iterator functions of section 2.2, calling the system-provided methods for lists (assumed to be efficient at least for forward iteration).

Given the code in figure 3, queues can be used wherever Common Lisp specifies that a sequence is acceptable. For instance, one can ask for the position of an element in the queue using `position`, with all the keywords (`:test`, `:key`, `:start` etc.) that the Common Lisp function accepts; some examples are given in figure 4. Furthermore, this queue implementation will potentially interoperate with bodies of code written for generic sequences, as long as there are no uses of the formal interpretation of the requirements of `make-sequence` and friends.

3.2 Undoable mixin

To illustrate possible uses of the sequence protocol discussed in this document, we present in figure 5 a mixin for implementing some undo functionality for a sequence. The `undoable-mixin` class contains a `record` slot, which records enough information to reconstruct the state of the sequence before an operation; if the user calls the `undo` function on a sequence with this class mixed in, the previous state of the sequence will be reconstructed using `undo-using-record` methods. A more sophisticated version of this might be useful as a component of an editor buffer implementation, for example.

For `(setf elt)`, we record the index and the value at that index before performing the mutation; clearly, this permits reconstructing the previous state of the sequence. For operations such as `fill` and `nreverse`, we could simply make do with this (and an analogous change for the setter function of the iterator protocol; see figure 8 in the appendix for that detail), but this would have the consequence that a single logical operation such as `fill` would require multiple calls to `undo` to undo the state.

Instead, therefore, for `fill`, we record the contents between the bounding index designators, while for `nreverse` we need record nothing, as it is a reversible operation⁸. An alter-

⁸We assume for expository purposes that the destructive

native strategy for grouping multiple primitive operations, used in the method for recording calls to `nsubstitute` and `delete`, is shown in figure 8 in the appendix.

4. IMPLEMENTATION DETAILS

We have implemented the above proposal in Steel Bank Common Lisp (Newman et al., 2000); the overall strategy was to modify the implementations of the standard Common Lisp sequence functions to call into their `sequence` package analogues instead of signalling an error, if their sequence arguments were neither lists nor vectors. This way of implementing this proposal was motivated by changing the run-time performance characteristics of the system as little as possible for existing code.

The first aspect of SBCL itself which needed to be modified was the type system: the system needed to be informed that `sequence` was no longer simply an alias in that implementation of `(or list vector)`, but was an unsealed (that is, subclassable) class with its own identity in the type system; since the SBCL compiler makes heavy use of type inference, both when compiling user code and when cross-compiling itself, it was also important for the compiler’s version of `subtypep` to know that the types `(and sequence vector)` and `(and sequence list)` are equivalent to `vector` and `list` respectively, irrespective of the extensibility of `sequence`.

Additionally, the declared return type of various sequence functions such as `copy-seq` needed to be altered: SBCL has the type `consed-sequence`, which was previously aliased to `(or list (simple-array * (*)))`, expressing that (in SBCL) freshly consed vectors do not have fill-pointers, are not displaced and are not adjustable. In an implementation supporting extensible sequences, this type alias needs to be changed to `(or (simple-array * (*)) (and sequence (not vector)))` to correctly represent the implementation’s behaviour⁹.

After these modifications to the type system, the implementation of the remainder of this proposal involved two distinct parts: modifying SBCL itself to insert calls to the sequence generic functions (say, `sequence:length`) from the standard functions (`length`), and implementing those generic functions in a userspace library. Incremental development of the SBCL internals was eased by the use of a `seq-dispatch` macro, taking a sequence argument and expanding into either two or three cases: one for when the argument is a list, one when it is a vector, and (optionally) one for neither: the third case is only executed when a generic sequence is encountered.

The various compiler transformations and optimizations of sequence functions (such as `find`, `position`, `map`, `coerce`) were largely unaffected, as either the various checks on their applicability were already restrictive enough, or else the optimizations performed were generic to all sequences. The implementation of `find` and `position` (and `-if` and `-if-not` relatives) needed a small amount of alteration to allow

sequence functions such as `nreverse` act in-place on the user-extended sequences in which this class will be mixed.

⁹SBCL has a sophisticated understanding of the Common Lisp type system, so such complex types do not cause difficulty in type inference.

```

(defclass undoable-mixin ()
  ((recording :initform nil :accessor recording) (record :initform nil :accessor record)))
(defclass setelt-record ()
  ((index :initarg :index :reader index) (value :initarg :value :reader value)))
(defclass fill-record ()
  ((start :initarg :start :reader start) (end :initarg :end :reader end)
   (contents :initarg :contents :reader contents)))
(defclass nreverse-record () ())
(defmacro without-undoing ((object) &body body)
  `(let ((old (recording ,object)))
     (unwind-protect (progn (setf (recording ,object) t) ,@body)
      (setf (recording ,object) old))))
(defmacro define-undo-method (name arglist &body body)
  `(defmethod ,name :around ,(substitute '(o undoable-mixin) 'o arglist)
    (if (recording o) (call-next-method) (without-undoing (o) ,@body (call-next-method))))
(define-undo-method (setf sequence:elt) (new-value o index)
  (push (make-instance 'setelt-record :index index :value (elt o index)) (record o)))
(define-undo-method sequence:fill (o item &key (start 0) end)
  (push (make-instance 'fill-record :start start :end end :contents (subseq (coerce o 'vector) start end))
    (record o)))
(define-undo-method sequence:nreverse (o) (push (make-instance 'nreverse-record) (record o)))
(defun undo (object)
  (undo-using-record object (car (record object))) object)
(defmethod undo-using-record ((o undoable-mixin) (r null)) (error "Nothing to undo"))
(defmethod undo-using-record :after ((o undoable-mixin) r) (pop (record o)))
(defmethod undo-using-record ((o undoable-mixin) (r setelt-record))
  (without-undoing (o) (setf (elt o (index r)) (value r))))
(defmethod undo-using-record ((o undoable-mixin) (r fill-record))
  (without-undoing (o)
   (with-accessors ((start start) (end end) (contents contents)) r (setf (subseq o start end) contents))))
(defmethod undo-using-record ((o undoable-mixin) (r nreverse-record))
  (without-undoing (o) (sequence:nreverse o)))

```

Figure 5: Illustration of a simple implementation of undo for sequences: an undoable-mixin class which can be mixed in to a sequence class, providing undo functionality (illustrated here for (setf elt), fill and nreverse).

the protocol for `sequence:find` and `sequence:position` described in section A.3 to be implemented, while the implementation of `map` was improved in the process of development, being rewritten to use `dynamic-extent` support, resolving a long-standing issue in the system.

Support for defining subclasses of `sequence` from the customized PCL (Bobrow and Kiczales, 1988) which SBCL incorporates to implement CLOS was as simple as modifying the system method for the MetaObject Protocol function `mop:validate-superclass` to allow `sequence` as a direct superclass of classes with metaclasses `standard-class` and `mop:funcallable-standard-class`. In addition, to support the operations involving sequence type specifiers (such as `map` and `merge`), we used `make-sequence-like` acting on the `mop:class-prototype` of the class named by the type specifier, if that class was a subclass of `sequence`.

SBCL’s `loop` facility is based on the MIT implementation, so it was straightforwardly extended using the `add-loop-path` operator as shown in figure 6.

In the implementation of this proposal for SBCL, there is no run-time performance impact for bodies of code which does not use the facilities provided by user-extensible sequences (see <http://sbcl.boinkor.net/bench/> for benchmark data: the relevant changeset is labelled 1.0.0.22). This is understandable, as all that has been done to the run-time code in this implementation is to convert cases which used to signal errors (usually the `type-error` for the datum

not being a `sequence`) into calls to the analogous function of the `sequence` package. Changes in compile-time performance are potentially present but unmeasurable: the only likely contributor is the additional work in computing type relationships between `sequence` and other types. There is an increase of approximately 1% of the heap size of a fresh SBCL instance, because of the extra code used in implementing this proposal; however, this extra space is in a tenured generation and behind a write barrier, and therefore does not significantly increase garbage-collection times.

5. CONCLUSIONS

We have described an extension to Common Lisp to allow users to define their own sequence classes in an interoperable manner, and have described the changes necessary to implement this extension in a contemporary implementation. Additionally, we have attempted to justify the need for this extension in terms of expressivity and parsimony, and given simple examples of using it.

We leave for future work the related task of defining a protocol for user-defined collections; Common Lisp sequences are a specialized form of collection, finite and ordered. However, preliminary work suggests that such a collection protocol can be defined to interoperate with the proposals in this document for sequences.

It is intended that a revised and expanded version of appendix A should be submitted to some suitable forum for standardization of some form; it is likely that feedback gath-

```

(defun loop-elements-iteration-path (variable data-type prep-phrases)
  (let (of-phrase)
    (loop for (prep . rest) in prep-phrases do
      (ecase prep
        ((:of :in) (if of-phrase
                      (sb-loop::loop-error "Too many prepositions")
                      (setq of-phrase rest))))))
    (destructuring-bind (it lim f-e step endp elt seq)
      (loop repeat 7 collect (gensym))
      (push '(let ((,seq ,(car of-phrase)))) sb-loop::*loop-wrappers*)
      (push '(sequence:with-sequence-iterator (,it ,lim ,f-e ,step ,endp ,elt) (,seq))
            sb-loop::*loop-wrappers*)
      '(((,variable nil ,data-type) () () nil (funcall ,endp ,seq ,it ,lim ,f-e)
        (,variable (funcall ,elt ,seq ,it) ,it (funcall ,step ,seq ,it ,f-e))))))
  (sb-loop::add-loop-path
   'element elements) 'loop-elements-iteration-path sb-loop::*loop-ansi-universe*
   'preposition-groups '(:of :in) :inclusive-permitted nil)

```

Figure 6: Extension of loop for a loop path to iterate over the elements of generic sequences.

ered from discussion of this paper will mean that there will be some changes in detail in the revision process, so users should not rely on the description in this paper remaining authoritative.

Acknowledgments

We thank Nikodemus Siivola, Marcus Pearce, Gábor Melis, Eric Marsden, Peter Housel, Cyrus Harmon, Peter Graves and Pascal Costanza for fruitful and helpful discussions, and acknowledge financial support from EPSRC grant number GR/S84750/01. We also wish to thank the reviewers for their valuable feedback.

References

- Bobrow, D. G. and Kiczales, G. (1988). The Common Lisp Object System Metaobject Kernel: A Status Report. In *Lisp and Functional Programming*, pages 309–315.
- Franz Inc. (2006). Allegro CL version 8.0 documentation. <http://tinyurl.com/3cfu66>.
- Gray, D. N. (1989). Issue STREAM-DEFINITION-BY-USER. Failed Issue, X3J13, ANSI. <http://tinyurl.com/2d4csd>.
- Haible, B. (1988). The Abstract Datatype *Sequence*. Technical report, University of Karlsruhe. <http://tinyurl.com/yy3eys>.
- Haible, B. (2006). personal communication.
- Kelsey, R., Clinger, W., and Rees, J. (1998). Revised⁵ Report on the Algorithmic Language Scheme. *Higher-Order and Symbolic Computation*, 11(1).
- Kiczales, G., des Rivières, J., and Bobrow, D. G. (1991). *The Art of the Metaobject Protocol*. MIT Press.
- Kiczales, G., Lamping, J., Menhdhekar, A., Maeda, C., Lopes, C., Loingtier, J.-M., and Irwin, J. (1997). Aspect-Oriented Programming. In *European Conference on Object-Oriented Programming*, pages 220–242.
- Miller, S. G. (2004). Collections. SRFI 44, [schemers.org](http://srfi.schemers.org). <http://srfi.schemers.org/srfi-44/srfi-44.html>.
- Newman, W. H. et al. (2000). SBCL User Manual. <http://www.sbcl.org/manual/>.

Norvig, P. (2002). Large-Scale Web Services, and the Programming Languages that Build Them. In *International Lisp Conference Proceedings*.

Pestov, S. (2006). Factor documentation. <http://factorcode.org/responder/help/>.

Pitman, K. and Chapman, K., editors (1994). *Information Technology – Programming Language – Common Lisp*. Number 226–1994 in INCITS. ANSI.

Pitman, K. M. (1991). Issue CONCATENATE-SEQUENCE. Issue 73, X3J13, ANSI. <http://tinyurl.com/2zt4d6>.

Rhodes, C. (2006). Revisiting CONCATENATE-SEQUENCE. Document 3, Common Lisp Document Repository. <http://cdr.eurolisp.org/document/3>.

Rhodes, C., Strandh, R., and Mastenbrook, B. (2005). Syntax Analysis in the Climacs Text Editor. In *International Lisp Conference Proceedings*.

Shalit, A. (1996). *The Dylan Reference Manual*. Addison-Wesley, Redwood City, CA, USA. <http://www.opendylan.org/books/drm/Title>.

Steele, Jr, G. L. (1990). *Common Lisp: The Language*. Digital Press, second edition.

Strandh, R., Villeneuve, M., and Moore, T. (2004). Flexichain: An editable sequence and its gap-buffer implementation. In *European Lisp and Scheme Workshop*. <http://tinyurl.com/yhqwp3>.

Yee, K.-P. and van Rossum, G. (2001). Iterators. PEP 234, Python Software Foundation. <http://www.python.org/dev/peps/pep-0234/>.

APPENDIX

A. SPECIFICATION

In the sections below, all of the generic functions and macros being specified are named by symbols external in the `sequence` package. In cases where the generic function being defined corresponds to a standardized Common Lisp function, it is not specified whether the corresponding Common Lisp function is distinct from the specified

```

;;; syntax: #[1 2 3 ...] for ordinary queues, #{1 2 3 ...} for funcallable ones
(macrolet ((def (name open close)
  (let ((reader-name (intern (format nil "~A~A" name 'reader))))
    '(progn
      (defun ,reader-name (stream char n)
        (declare (ignore char n))
        (let ((contents (read-delimited-list ,close stream t))
              (result (make-instance ',name)))
          (dolist (o contents result) (enqueue o result))))
      (set-dispatch-macro-character #\# ,open ',reader-name)
      (set-syntax-from-char ,close #\))
      (defmethod print-object ((o ,name) stream)
        (when (and *print-readably* (not (eq (class-of o) (find-class ',name))))
          (call-next-method))
        (format stream "#~C" ,open)
        (do ((data (cdr (%queue-data o)) (cdr data)) (i 0 (1+ i)))
            ((null data) (format stream "~C" ,close) o)
          (unless (or *print-readably* (not *print-length*))
            (when (= i *print-length*)
              (format stream "...~C" ,close)
              (return o)))
          (let ((*print-level* (and *print-level* (1- *print-level*))))
            (write (car data) :stream stream))
          (unless (null (cdr data)) (format stream " ")))))))
  (def queue #\[ #\]) (def funcallable-queue #{ #\}))

```

Figure 7: Implementation of reader macros and print functions for queues and funcallable queues. Note the support for correct behaviour in the presence of non-nil values of printer control variables.

generic functions, nor indeed whether the symbol in the `common-lisp` package is distinct from that in the `sequence` package; that is, `#'length` might or might not be `eql` to `#'sequence:length`, and `length` might or might not be `eql` to `sequence:length`. The `sequence` package may have additional, implementation-specific names; `sequence` need not be the `package-name` of the package.

It is not specified whether the methods specified on list and vector are in fact called when the Common Lisp function corresponding to the generic function is called on such data. It is not specified whether the methods specified on sequence are called when the Common Lisp function corresponding to the generic function is called on data of type vector or list.

In implementations which support the Metaobject Protocol as defined in Kiczales et al. (1991), suitable methods on `mop:validate-superclass` should be provided such that no error is signalled for user-defined sequence classes of metaclass `standard-class`; such implementations are additionally encouraged to allow `mop:funcallable-standard-class` as a compatible metaclass for user-defined sequence classes.

A.1 Basic sequence operations

We specify generic functions `length`, `elt` and `(setf elt)` to correspond to the Common Lisp functions with the same name. In each case, there are two primary methods with the sequence argument specialized on `list` and on `vector`, providing the standard-defined behaviour for the Common Lisp operator, and a third method with the sequence argument specialized on `sequence`, which signals an error of type `type-error`, for compatibility with the standard requirement of the sequence argument to be a *proper sequence*.

In addition, we specify two new generic functions and corresponding methods (with a similar pattern of behaviour)

described below.

Generic Function `make-sequence-like`

Syntax:

`make-sequence-like` *sequence length &key*
initial-element initial-contents

This generic function returns a sequence of the same class as its *sequence* argument, with the specified *length*. If neither *initial-element* nor *initial-contents* is supplied, the consequences are undefined if any element of the resulting sequence is read before being written. If *initial-element* is provided, it is used to initialize the resulting sequence; if *initial-contents* is provided, it must be a sequence of length *length*, which is used to initialize the resulting sequence. The consequences are undefined if both *initial-element* and *initial-contents* are supplied.

Primary Method `make-sequence-like` (*l* list) *length &key*
initial-element initial-contents

Primary Method `make-sequence-like` (*v* vector) *length*
&key initial-element initial-contents

No behaviour for these methods is specified beyond that for the generic function.

Primary Method `make-sequence-like` (*s* sequence) *length*
&key initial-element initial-contents

This method signals an error of type `type-error`.

Generic Function `adjust-sequence`

Syntax:

`adjust-sequence` *sequence length &key initial-element*
initial-contents

```

(defmethod sequence:make-sequence-iterator :around ((o undoable-mixin) &key from-end start end)
  (declare (ignore from-end start end))
  (multiple-value-bind (s l f step endp elt setelt index copy) (call-next-method)
    (values s l f step endp elt
            (lambda (nv s i)
              (if (recording s)
                  (funcall setelt nv s i)
                  (without-undoing (s)
                                   (push (make-instance 'setelt-record :index (funcall index s i) :value (funcall elt s i))
                                         (record s))
                                   (funcall setelt nv s i))))
            index copy)))

(defmacro with-compound-recording ((object) &body body)
  `(let ((old (recording ,object)) (or (record ,object)))
      (unwind-protect
        (progn (setf (recording ,object) nil) (setf (record ,object) nil) ,@body)
        (setf (recording ,object) old)
        (setf (record ,object) (cons (make-instance 'compound-record :records (record ,object)) or))))))

(defclass compound-record () ((records :initarg :records :reader records)))
(defmethod undo-using-record ((o undoable-mixin) (r compound-record))
  (dolist (r (records r)) (undo-using-record o r)))
(defmacro define-compound-undo-method (name arglist)
  `(define-undo-method ,name ,arglist
    (with-compound-recording (o) (call-next-method))))

(define-compound-undo-method sequence:nsubstitute (new old o &key from-end start end test test-not key count))
(define-compound-undo-method sequence:nsubstitute-if (new old o &key from-end start end key count))
(define-compound-undo-method sequence:nsubstitute-if-not (new old o &key from-end start end key count))

(defclass adjust-sequence-record ()
  ((previous :initarg :previous :reader previous) (discarded :initarg :discarded :reader discarded)))
(defmethod undo-using-record (o length &rest args)
  (push (make-instance 'adjust-sequence-record :previous (length o)
                    :discarded (when (< length (length o)) (coerce (subseq o length) 'vector)))
        (record o)))
(defmethod undo-using-record ((o undoable-mixin) (r adjust-sequence-record))
  (sequence:adjust-sequence o (previous r))
  (when (discarded r) (setf (subseq o (- (length o) (length (discarded r)))) (discarded r))))

(define-compound-undo-method sequence:delete (item o &key from-end (start 0) end test test-not key count))
(define-compound-undo-method sequence:delete-if (item o &key from-end (start 0) end key count))
(define-compound-undo-method sequence:delete-if-not (item o &key from-end (start 0) end key count))
(define-compound-undo-method sequence:delete-duplicates (o &key from-end (start 0) end test test-not key))
(define-undo-method sequence:replace (o sequence2 &key (start1 0) end1 (start2 0) end2)
  (push (make-instance 'fill-record :start start1 :end end1 :contents (coerce (subseq o start1 end1) 'vector))
        (record o)))

(define-undo-method sequence:sort (o predicate &key key)
  (push (make-instance 'fill-record :contents (coerce o 'vector) :start 0 :end nil) (record o)))
(define-undo-method sequence:stable-sort (o predicate &key key)
  (push (make-instance 'fill-record :contents (coerce o 'vector) :start 0 :end nil) (record o)))

```

Figure 8: Implementation details of the undoable mixin. The setter in the iterator protocol is wrapped by the `:around` method on `make-sequence-iterator`; a specialized undo record is defined for `adjust-sequence`; and undo records for `nsubstitute` and `delete` are represented as sequences of primitive undo records, treated as an atom for the purposes of running `undo`.

This generic function changes the length or elements of *sequence*. It returns a sequence of length *length*, with contents as specified by *initial-contents* if provided, otherwise with the original contents in their original position and *initial-element* as any additional elements. The return value may, but need not, share structure with or be identical to the *sequence* argument.

Primary Method `adjust-sequence` (*l* list) *length* &*key*
initial-element *initial-contents*

Primary Method `adjust-sequence` (*v* vector) *length* &*key*
initial-element *initial-contents*

This method functions in a similar manner to `adjust-array`, though the implementation may choose to preserve the identity of the argument if it has a fill pointer and the *length* argument is not greater than the size of the vector.

Primary Method `adjust-sequence` (*s* sequence) *length*
&*key* *initial-element* *initial-contents*

This method signals an error of type `type-error`.

A.2 Iteration

Generic Function `make-sequence-iterator`

Syntax:
`make-sequence-iterator` *sequence* &*key* *from-end* *start*
end

This generic function returns nine values: three values corresponding to an iterator state, a limit state and *from-end*, and six functions with signatures and functionality like the `iterator-` functions below.

Primary Method `make-sequence-iterator` (*s* sequence)
&*key* *from-end* (*start* 0) *end*

This method returns the three values returned by calling `make-simple-sequence-iterator`, along with the iterator functions `#'iterator-step`, `#'iterator-endp`, `#'iterator-element`, `#'(setf iterator-element)`, `#'iterator-index` and `#'iterator-copy` defined below.

For the seven following generic functions, we describe both the generic function and also in more detail the return values of the primary method, specialized on `sequence`. The primary methods described here are suitable for iteration across data structures with efficient random access, such as vectors, and less suitable for iteration across some other data structures, such as lists. We expect implementations of this protocol to provide additional methods to perform more efficient iteration on lists than is possible with the primary method described; however, we do not impose additional requirements on the behaviour of such methods.

Generic Function `make-simple-sequence-iterator`

Syntax:
`make-simple-sequence-iterator` *sequence* &*key*
from-end *start* *end*

This generic function returns three values: an iterator object, a limit state and *from-end*. These values, along with *sequence*, are to be used for calling to the iterator generic

functions, below; the consequences are unspecified if objects not returned by a call to `make-simple-sequence-iterator` are passed as arguments to the iterator functions. The primary method returns the three values *start*, *end* (or the length of *sequence* if *end* is `nil`) and `nil`, if *from-end* is *false*; if *from-end* is true, it returns the three values (1- (or *end* (length *sequence*))), (1- *start*), and *from-end*.

Generic Function `iterator-step`

Syntax:
`iterator-step` *sequence* *iterator* *from-end*

This generic function returns a new iterator state representing the advancement of the iteration across *sequence* in the direction indicated by *from-end*. The primary method returns (1- *iterator*) if *from-end* is true; otherwise, it returns (1+ *iterator*).

Generic Function `iterator-endp`

Syntax:
`iterator-endp` *sequence* *iterator* *limit* *from-end*

This generic function tests the iterator for having reached its end state for iteration across *sequence* indicated in the direction indicated by *from-end*. The primary method returns the value of (= *iterator* *limit*).

Generic Function `iterator-element`

Syntax:
`iterator-element` *sequence* *iterator*

This generic function returns the element of *sequence* at the point of iteration indicated by *iterator*. The primary method returns the value of (elt *s* *iterator*).

Generic Function (setf `iterator-element`)

Syntax:
(setf `iterator-element`) *new-value* *sequence* *iterator*

This generic function sets the element of *sequence* at the point of iteration indicated by *iterator* to *new-value*. The primary method returns the value of (setf (elt *s* *iterator*) *new-value*).

Generic Function `iterator-index`

Syntax:
`iterator-index` *sequence* *iterator*

This generic function returns the index of the point indicated by *iterator* in *sequence*. The primary method returns *iterator*.

Generic Function `iterator-copy`

Syntax:
`iterator-copy` *sequence* *iterator*

This generic function returns an iterator state identifying the same point in an iteration as *iterator*, such that changes to one do not affect the other. The primary method returns *iterator*.

Macro `with-sequence-iterator` (&*optional* *state* *limit*
from-end *step* *endp* *elt* *setelt* *index* *copy*) (*sequence* &*key*
from-end (*start* 0) *end*) &*body* *body*

This macro binds the names in its first argument list as if by `multiple-value-bind` to the values returned by `make-sequence-iterator` applied to its second argument list, and then executes *body*.

Macro `dosequence` (*var sequence-form* &optional *result-form*) &body *body*

This macro iterates over a sequence, in a similar fashion to `dolist` iterating over a list. The *body* is like a `tagbody`, consisting of a series of tags and statements.

`dosequence` evaluates *sequence-form*, which should produce a sequence. It then executes the body once for each element in the sequence, in the order in which the tags and statements occur, with *var* bound to the element. Then *result-form* is evaluated with *var* bound to `nil`.

An implicit block named `nil` surrounds the entire `dosequence` form. `return` may be used to terminate the loop immediately without performing any further iterations, returning zero or more values.

The scope of the binding of *var* does not include the *sequence-form*, but the *result-form* is included.

It is implementation-dependent whether `dosequence` establishes a new binding of *var* on each iteration or whether it establishes a binding for *var* once at the beginning and then assigns it on any subsequent iterations.

Loop Path `for-as-sequence`

This allows using an iteration variable, as with other loop `for-as-` clauses, as if the following clause were added to the loop grammar (Pitman and Chapman, 1994, Macro `loop`):

```
for-as-sequence ::=  
  var [type-spec] being {each | the}  
  {element | elements} {in | of} sequence
```

and as if the `for-as-subclause` clause permitted this `for-as-sequence` clause along with the standardized clauses. As with other iteration control clauses, the variable argument may be a destructuring list. The effect of this clause is to iterate through the contents of a sequence, starting from the zeroth element and terminating at the end of the sequence¹⁰.

A.3 Sequence Function Specifications

Except as specified here, it is implementation-dependent whether methods on these generic functions call other such generic functions or not. For each of these generic functions there is a method, called the “default method” below, where all sequence parameters are specialized on `sequence`, implementing the default behaviour given an implementation of the iteration and basic sequence protocol in the sections above. The user of this protocol is permitted to extend or override these generic functions, but is not permitted to specialize any non-sequence argument.

It is unspecified whether the generic functions specified below are called when the corresponding Common Lisp function has only `list` and `vector` arguments for sequences; however, the implementation is required to provide methods for these generic functions implementing similar behaviour, so that the user may call these generic functions on `list` and `vector` sequences.

For each Common Lisp function in table 2, we specify a generic function of the same name in the `sequence` package with the same argument list, each with a default method where all sequence arguments are specialized on `sequence`. Given that the sequence argument to these functions is named *sequence*, we enumerate the specified calling relationships between these functions as follows:

- the default method on `copy-seq` calls `subseq` with parameters *sequence* and 0;
- the default method on `substitute` calls `copy-seq` on *sequence*, and then calls `nsubstitute` on its arguments with the provided sequence replaced by the copy. This pattern is repeated for the `-if` and `-if-not` variants of `substitute`, and for `remove`, `remove-if`, `remove-if-not` and `remove-duplicates` calling the corresponding `delete-` function.
- the default method on `sort` behaves as if it constructs a vector with the same elements as *sequence*, calls `sort` on that vector, then replaces the elements of *sequence* with the elements of the sorted vector.

The rationale behind specifying these calling relationships is to make clear the work needed by a sequence class implementor to install specialized implementations of these functions for particular sequence classes.

¹⁰Note that this behaviour is different from that of the `for-as-across` loop clause on vectors with fill pointers.