

Document number: D0769R0
Date: 2017-08-25
Author: Dan Raviv <dan.raviv@gmail.com>
Audience: LEWG => LWG

Add shift to <algorithm>

I. Introduction

This paper proposes adding shift algorithms to the C++ STL which shift elements forward or backward in a range of elements.

II. Motivation and Scope

Shifting elements forward or backward in a range is a basic operation which the STL should allow performing easily. An important use case is [time series analysis](#) algorithms used in scientific and financial applications.

The scope of the proposal is adding the following function templates to <algorithm>:

```
template<class ForwardIt>
ForwardIt shift_left(
    ForwardIt first, ForwardIt last,
    typename std::iterator_traits<ForwardIt>::difference_type n
);

template<class ExecutionPolicy, class ForwardIt>
ForwardIt shift_left(
    ExecutionPolicy&& policy, ForwardIt first, ForwardIt last,
    typename std::iterator_traits<ForwardIt>::difference_type n
);

template<class ForwardIt>
ForwardIt shift_right(
    ForwardIt first, ForwardIt last,
    typename std::iterator_traits<ForwardIt>::difference_type n
);

template<class ExecutionPolicy, class ForwardIt>
ForwardIt shift_right(
    ExecutionPolicy&& policy, ForwardIt first, ForwardIt last,
    typename std::iterator_traits<ForwardIt>::difference_type n
);
```

A sample implementation which uses `std::move` to implement `shift_left` for forward iterators and `std::move_backward` to implement `shift_right` for bidirectional iterators can be found in https://github.com/danra/shift_proposal, though it's possible more efficient implementations could be made, since elements are guaranteed to be moved within the same range, not between two different ranges.

The sample implementation also implements a non-trivial algorithm for `shift_right` of forward, non-bidirectional iterators.

III. Possible Objections and Responses

- 1) Objection: Shifting can be done by using `std::move` (in <algorithm>).
 Response: Which of `std::move` or `std::move_backward` must be used depends on the shift direction, which is error-prone. It also makes for less readable code; consider
`std::shift_right(v.begin(), v.end(), 3);`
 vs.
`std::move_backward(v.begin(), v.end() - 3, v.end());`

In addition, `std::shift_right` and `std::shift_left` may be implemented more efficiently than `std::move` and `std::move_backward`, since elements are guaranteed to be moved within the same range, not between two different ranges.

Also, ranges of forward, non-bidirectional iterators cannot be shifted right using either `std::move` or `std::move_backward`. Such ranges are possible to shift right, though, in $O(N)$ time and constant space, as shown in the sample implementation.

- 2) Objection: Instead of shifting a range, you can use a circular buffer.
 Response: A circular buffer is a valid alternative. However, it should not be forced on the programmer, and it does have its own limitations:
- In case there are multiple indices into the buffer, all must be updated in some way.
 - Similarly, in the common case where there is some mask applied to the buffer which should not cycle with the data, the mask indices need to be updated whenever the buffer is cycled.
 - A programmer might need to shift elements in non-circular buffers provided by a 3rd-party library.
- 3) Objection: There's already `std::rotate` which is similar in functionality.
 Response: Shifting just the desired elements would allow for both a more efficient implementation and clearer semantics in case rotation is not needed.

IV. Impact On the Standard

The only impact on the standard is adding the proposed function templates to <algorithm>.

V. Design Decisions

1) `shift_left` and `shift_right` are provided as separate function templates instead of just a single `shift` function template to maximize performance and minimize compiled code size. Since shifting left and shifting right may have significantly different implementations (as is the case in the sample implementation), implementing both shift directions in a single `shift` function template would both require extra conditional logic and inline less easily than the specific direction shifts.

Given that both `shift_left` and `shift_right` are provided, it would still be possible to provide `shift` as well, for convenience, but it seems redundant.

2) `std::shift_left` should return an iterator to the new end of the shifted range. The beginning of the shifted range would always be equal to the beginning of the range before the shift, so there is no need to also return an iterator to the beginning of the shifted range. (This is similar to how `std::move` only returns an iterator to the end of the moved range).

Similarly, `std::shift_right` should return an iterator to the new beginning of the shifted range. The end of the shifted range would always be equal to the end of the range before the shift, so there is no need to also return an iterator to the end of the shifted range. (This is similar to how `std::move_backward` only returns an iterator to the end of the moved range).

3) After shifting a range by n elements, either to the right or to the left, exactly n elements would be left “empty”, with their previous values having been shifted to other elements but with no new values shifted into them. It’s been suggested to provide function template overloads which with an extra `filler` value parameter which would set all such “empty” elements to its value. However, this has been decided to be redundant, since the iterator returned from `shift_left` / `shift_right` (see (2) above) can be passed to `std::fill` (along with unchanged `begin`/`end` of the range) to fill the empty values.

4) `std::shift_left` without an execution policy or with the standard `sequenced_policy` execution policy moves the shifted elements (those which would still present in the range after the shift) in order, similar to how `std::move` moves elements in order.

Similarly, `std::shift_right` of bidirectional iterators without an execution policy or with the standard `sequenced_policy` execution policy moves the shifted elements (those which would still present in the range after the shift) in reverse order, similar to how `std::move_backward` moves elements in reverse order.

There is no guarantee what order `std::shift_right` of forward iterators shifts elements in.

5) Shifting a range by more than its length (`std::distance(first, last)`), either to the left or to the right, has the effect of shifting out all of the elements, the same as shifting a range by exactly the length of the range. This could be defined as undefined behavior instead, but there is would probably be no extra cost to handle larger shifts, seeing as an implementation would have to handle a shift by exactly the length of the range anyway.

6) Shifting a range by a negative n is undefined behavior. (in case shift by zero is decided to be defined behavior, see open issues below, shifting by negative n could also be defined behavior and do nothing, since it would have no extra cost. However, it’s not really important for a shift by negative n to be defined behavior, so it shouldn’t be a factor for deciding on shift by zero being defined or not).

VI. Open Issues

1) Should shift by zero be undefined behavior?

Pro for undefined behavior: Could simplify implementation and optimize performance. For example, in the [sample implementation](#), since both `std::move` and `std::move_backward` have undefined behavior when moving a range exactly onto itself, an extra `n==0` condition check must be done before performing either of them for a shift by zero to have defined behavior.

Pro for defined behavior: It is reasonable to expect a shift by zero to do nothing, so it is programmer-error-prone to make it undefined behavior.

2) It would be preferable for `std::shift_left` and `std::shift_right` to have more generic names; the fact that the first element in a range is the left-most is a matter of convention which is not specified in the standard, and some programmers may think of the first element as the right most, or maybe the top-most, etc.

However, `std::shift_backward`, `std::shift_back` and `std::shift_forward` are probably all out of the question, since other algorithms exist, e.g., `std::move_backward` and `std::copy_backward`, in which backward means performing the operation starting from the back of the range, instead of from its front.

`std::shift_to_front` and `std::shift_to_back` come to mind. Perhaps there are better names; ideas would be welcome.

VII. Proposed Wording

TODO

VIII. Acknowledgements

- Special thanks to Walter E. Brown for his comments and guidance writing the paper.
- Special thanks to Casey Carter for his comments and generous code contribution, including a `shift_right` implementation for forward, non-bidirectional iterators.
- Thanks to Alexander Zaitsev, Arthur O'Dwyer, Bryce Glover, Howard Hinnant, Nicol Bolas and Ray Hamel for their helpful comments on the paper.