

Document no: N3465=12-0155
Date: 2012-10-29
Revises: N2882=09-0072
Project: Programming Language C++
Reply to: Joaquín M^a López Muñoz
joaquin@tid.es

Adding heterogeneous comparison lookup to associative containers for TR2 (Rev 2)

C++11 has extended the applicability of binary search algorithms so that they accept keys and comparison operators of types other than those used for sorting the ranges being searched, provided that some compatibility conditions are met¹. For instance, this extension allows us to write the following:

```
struct name_entry
{
    std::string family_name;
    std::string given_name;
};

bool operator<(const name_entry& x, const name_entry& y)
{
    // lexicographical order on (family_name, given_name)

    if(x.family_name<y.family_name) return true;
    if(y.family_name<x.family_name) return false;
    return x.given_name<y.given_name;
}

struct comp_family_name
{
    bool operator()(const name_entry& x, const std::string& y) const
    {
        return x.family_name<y;
    }
    bool operator()(const std::string& x, const name_entry& y) const
    {
        return x<y.family_name;
    }
};

int main()
{
    std::vector<name_entry> names;
    ... // populate names;
    std::sort(names.begin(), names.end());
    // look for all Smiths
    std::equal_range(
        names.begin(), names.end(),
        std::string("Smith"), comp_family_name());
}
```

Conceptually, the extension consists in replacing the original formulation based on strict weak orderings with one relying on the notion of *sequence partitioning*, as first proposed by David Abrahams². Unfortunately, this extension process has not been carried out for the lookup operations of associative containers, which are still

¹ “Binary search requirements overly strict”, LWG issue 270, C++ Standard Library Defect Report List, <http://www.open-std.org/Jtc1/sc22/wg21/docs/lwg-defects.html#270>

² David Abrahams, “Binary Search with Heterogeneous Comparison”, J16-01/0027 = WG21 N1313, 2001.

formulated in terms of strict weak orderings and thus do not allow for heterogeneous comparison. So, if in the example above we had used a set rather than a vector we could not do this:

```
int main()
{
    std::set<name_entry> names;
    ... // populate names;
    // this does not compile
    names.equal_range(std::string("Smith"), comp_family_name());
}
```

and would have to resort to

```
int main()
{
    std::set<name_entry> names;
    ... // populate names;
    // look for all Smiths
    std::equal_range(
        names.begin(), names.end(),
        std::string("Smith"), comp_family_name());
}
```

which, since set iterators are bidirectional, has linear complexity, when set lookup operations are logarithmic. We propose to replace the current `equal_range` operations of associative containers

```
pair<iterator, iterator>          equal_range(const key_type& x);
pair<const_iterator, const_iterator> equal_range(const key_type& x) const;
```

with the following ones based on sequence partitioning concepts:

```
template<typename T>
    pair<iterator, iterator> equal_range(const T& x);
template<typename T, typename Compare>
    pair<iterator, iterator> equal_range(const T& x, Compare comp);
template<typename T>
    pair<const_iterator, const_iterator> equal_range(const T& x) const;
template<typename T, typename Compare>
    pair<const_iterator, const_iterator>
        equal_range(const T& x, Compare comp) const;
```

and similarly for the other lookup operations (`find`, `count`, `lower_bound` and `upper_bound`).

Implementation

At least for realizations of associative containers based on red-black trees, implementing the proposed extension is entirely trivial. For instance, starting from a canonical implementation of `lower_bound`

```
iterator lower_bound(const key_type& x)
{
    Node* top = root();
    Node* y = header();

    while(top) {
        if(!comp(top->value, x)) { // comp is the internal comparison object
            y = top;
            top = top->left;
        }
    }
}
```

```

        else top = top->right;
    }

    return iterator(y);
}

```

we can easily derive the partitioning-based extension:

```

template<typename T, typename Compare>
iterator lower_bound(const T& x, Compare comp)
{
    Node* top = root();
    Node* y = header();

    while(top){
        if(!comp(top->value, x)){ // comp is provided by the user
            y = top;
            top = top->left;
        }
        else top = top->right;
    }

    return iterator(y);
}

```

Note that the code remains exactly the same, except that we substitute the user-provided `comp` for the internal comparison object used before. This nice property, for which we provide a formal justification in an annex to this paper, holds for the rest of lookup operations as well.

Existing practice

Some of the components of the Boost MultiIndex library³ provide lookup facilities with heterogeneous comparison in a manner similar to that described here. The author has received some reports pointing to this functionality as reason alone to use Boost.MultiIndex in place of standard associative containers, leaving aside the more prominent multi-indexing capabilities offered by the library.

Proposed resolution

1. Change 23.2.4 [associative.reqmts] paragraph 8 from:

[...] `k` denotes a value of `X::key_type` and `c` denotes a value of type `X::key_compare`. [...]

to:

[...] `k` denotes a value of `X::key_type` and `c` denotes a value of type `X::key_compare`; `kl` is a value such that `a` is partitioned (25.4) with respect to `c(r, kl)`, with `r` the key value of `e` and `e` in `a`; `kcl` is a value and `cl` a copy constructible value such that that `a` is partitioned with respect to `cl(r, kcl)`; `ku` is a value such that `a` is partitioned with respect to `!c(ku, r)`; `kcu` is a value and `cu` a copy constructible value such that that `a` is partitioned with respect to `!cu(kcu, r)`; `ke` is a value

³ Joaquín M^a López Muñoz, Boost Multi-index Containers Library, http://www.boost.org/libs/multi_index

such that a is partitioned with respect to $c(r, ke)$ and $!c(ke, r)$, with $c(r, ke)$ implying $!c(ke, r)$; kce is a value and ce a copy constructible value such that a is partitioned with respect to $ce(r, kce)$ and $!ce(kce, r)$, with $ce(r, kce)$ implying $!ce(kce, r)$. [...]

2. Replace the following entries from Table 102 of section 23.2.4 [associative.reqmts]:

Expression	Return type	Assertion/note pre- / post-condition	Complexity
<code>a.find(k)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to an element with the key equivalent to <code>k</code> , or <code>a.end()</code> if such an element is not found	logarithmic
<code>a.count(k)</code>	size_type	returns the number of elements with key equivalent to <code>k</code>	$\log(\text{size}()) + \text{count}(k)$
<code>a.lower_bound(k)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to the first element with key not less than <code>k</code> , or <code>a.end()</code> if such an element is not found.	logarithmic
<code>a.upper_bound(k)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to the first element with key greater than <code>k</code> , or <code>a.end()</code> if such an element is not found.	logarithmic
<code>a.equal_range(k)</code>	pair<iterator, iterator>; pair<const- iterator, const- iterator> for constant <code>a</code> .	equivalent to <code>make_pair(a.lower_bound(k), a.upper_bound(k))</code> .	logarithmic

with:

Expression	Return type	Assertion/note pre- / post-condition	Complexity
<code>a.find(ke)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to an element with key <code>r</code> such that $!c(r, ke) \ \&\& \ !c(ke, r)$, or <code>a.end()</code> if such an element is not found	logarithmic
<code>a.find(kce, ce)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to an element with key <code>r</code> such that $!ce(r, kce) \ \&\& \ !ce(kce, r)$, or <code>a.end()</code> if such an element is not found	logarithmic
<code>a.count(ke)</code>	size_type	returns the number of elements with key <code>r</code> such that $!c(r, ke) \ \&\& \ !c(ke, r)$	$\log(\text{size}()) + \text{count}(ke)$
<code>a.count(kce, ce)</code>	size_type	returns the number of elements with key <code>r</code> such that $!ce(r, kce) \ \&\& \ !ce(kce, r)$	$\log(\text{size}()) + \text{count}(kce, ce)$
<code>a.lower_bound(kl)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to the first element with key <code>r</code> such that $!c(r, kl)$, or <code>a.end()</code> if such an element is not found.	logarithmic
<code>a.lower_bound(kcl, cl)</code>	iterator; const- iterator for constant <code>a</code> .	returns an iterator pointing to the first element with key <code>r</code> such that $!c(r, kcl) \ \&\& \ !c(kcl, cl)$, or <code>a.end()</code> if such an element is not found.	logarithmic

	iterator for constant a.	that !cl(r, kcl), or a.end() if such an element is not found.	
a.upper_bound(ku)	iterator; const_iterator for constant a.	returns an iterator pointing to the first element with key r such that c(ku, r), or a.end() if such an element is not found.	logarithmic
a.upper_bound(kcu, cu)	iterator; const_iterator for constant a.	returns an iterator pointing to the first element with key r such that cu(kcu, r), or a.end() if such an element is not found.	logarithmic
a.equal_range(ke)	pair<iterator, iterator>; pair<const_iterator, const_iterator> for constant a.	equivalent to make_pair(a.lower_bound(ke), a.upper_bound(ke)).	logarithmic
a.equal_range(kce, ce)	pair<iterator, iterator>; pair<const_iterator, const_iterator> for constant a.	equivalent to make_pair(a.lower_bound(kce, ce), a.upper_bound(kce, ce)).	logarithmic

3. In 23.4.4.1 [map.overview], 23.4.5.1 [multimap.overview], 23.4.6.1 [set.overview] and 23.4.7.1 [multiset.overview], replace:

```

iterator      find(const key_type& x);
const_iterator find(const key_type& x) const;
size_type     count(const key_type& x) const;

iterator      lower_bound(const key_type& x);
const_iterator lower_bound(const key_type& x) const;
iterator      upper_bound(const key_type& x);
const_iterator upper_bound(const key_type& x) const;

pair<iterator, iterator>
  equal_range(const key_type& x);
pair<const_iterator, const_iterator>
  equal_range(const key_type& x) const;

```

with:

```

template<typename T>
  iterator find(const T& x);
template<typename T, typename Compare>
  iterator find(const T& x, Compare comp);
template<typename T>
  const_iterator find(const T& x) const;
template<typename T, typename Compare>
  const_iterator find(const T& x, Compare comp) const;
template<typename T>
  size_type count(const T& x) const;
template<typename T, typename Compare>
  size_type count(const T& x, Compare comp) const;

template<typename T>
  iterator lower_bound(const T& x);
template<typename T, typename Compare>
  iterator lower_bound(const T& x, Compare comp);
template<typename T>
  const_iterator lower_bound(const T& x) const;
template<typename T, typename Compare>
  const_iterator lower_bound(const T& x, Compare comp) const;

```

```

template<typename T>
    iterator upper_bound(const T& x);
template<typename T, typename Compare>
    iterator upper_bound(const T& x, Compare comp);
template<typename T>
    const_iterator upper_bound(const T& x) const;
template<typename T, typename Compare>
    const_iterator upper_bound(const T& x, Compare comp) const;

template<typename T>
    pair<iterator, iterator> equal_range(const T& x);
template<typename T, typename Compare>
    pair<iterator, iterator> equal_range(const T& x, Compare comp);
template<typename T>
    pair<const_iterator, const_iterator> equal_range(const T& x) const;
template<typename T, typename Compare>
    pair<const_iterator, const_iterator>
    equal_range(const T& x, Compare comp) const;

```

Impact on existing code

There are pathological situations where this extension can break valid code or result in modified behavior; for instance, if `c` is an associative container, `key_type` is its key type and `x` a value of a type other than `key_type` that is implicitly convertible to `const key_type&`, the expression

```
c.find(x);
```

is currently equivalent to

```
c.find(static_cast<const key_type&>(x));
```

whereas under this proposal the conversion to `const key_type&` would not take place.

Additional considerations

Extending `erase`. It seems natural to apply this extension to another member function where comparison is used:

```
size_type erase(const key_type& x);
```

There are some difficulties here, though; extending this member function would clash with the homonym

```
iterator erase(const_iterator position);
```

provoking potential backwards compatibility problems (e.g. if `erase(x)` is invoked where `x` is a value of a type implicitly convertible to `const_iterator`, the extended key-based `erase` member function template would take precedence over the iterator-based `erase`). This issue is akin to that described in the section “**Impact on existing code**”, though probably a little less pathological. It can be argued that the problematic situations are unlikely to happen in real code and they could in any case be alleviated by carefully crafting the `requires` section of the extended `erase`.

Monomorphism of `std::less`. In the extension of `<algorithm>` functions used as a reference for this paper, those functions relying on operator `<`, such as

```
template<typename Iter, typename T>
Iter lower_bound(Iter first, Iter last, const T& value);
```

are typically more powerful than their equivalent member functions under the current proposal:

```
template<class T>
iterator lower_bound(const T& x);
```

due to the fact that `<` is inherently polymorphic, while in the case of associative containers an internal comparison object is used whose type `key_compare` is usually the monomorphic `std::less<key_type>`. Although this is probably beyond the scope of the proposal, it would be interesting to investigate the possibility that associative containers used a polymorphic type rather than `std::less` for their default comparison type, for instance:

```
struct polymorphic_less{
    template <typename T, typename Q>
        bool operator()(const T& x, const Q& y) const{ return x<y; }
};
```

Unordered associative containers. Although not considered in this paper, an analog extension of lookup facilities can be applied to unordered associative containers as well. Whereas for associative containers external keys are compatible with a range if they properly partition it, in the case of unordered associative containers the compatibility criterion is: elements of the range deemed equivalent (and only those) are equal to the key (in the context of the equality predicate used) and have the same associated hash value.

Acknowledgements

José Daniel García has kindly reviewed this version of the paper. Beman Dawes and Kevin Sopp reviewed former versions.

Annex

As we have seen before, usual lookup algorithms on a sorted range are formally equivalent to the extended algorithms needed to accommodate partitioning-based semantics: it only takes to utilize the user-provided heterogeneous comparison object in place of the internal comparison predicate used to sort the range. To prove this fact we need the following

Proposition. Let T be an arbitrary set with an associated strict weak order $<_T$ and L, U subsets of T such that

$$\begin{aligned} L \cap U &= \emptyset, \\ \forall a, b \in T \quad a \in L, b <_T a &\rightarrow b \in L, \\ \forall a, b \in T \quad a \in U, a <_T b &\rightarrow b \in U. \end{aligned}$$

We create the set Q by augmenting T with an additional element x , and define the binary relationship $<_Q$ on Q as follows:

$$\begin{aligned}
a <_Q b &:= a <_T b, \\
a <_Q x &:= a \in L, \\
x <_Q a &:= a \in U, \\
x <_Q x &:= \text{false},
\end{aligned}$$

for all $a, b \in T$. Under these conditions, $<_Q$ is a strict weak order on Q . (Proof trivial.)

Returning to our original scenario, let $[first, last)$ be a range of values of a type T sorted by some strict weak ordering, x a value of a type other than T and $comp$ a heterogeneous comparison object such that $[first, last)$ is partitioned (in the C++0x sense) both with respect to $comp(\cdot, x)$ and $!comp(x, \cdot)$, with $comp(e, x)$ implying $!comp(x, e)$. Now, we can regard $[first, last)$ as a range of elements of the set $Q = \{e \text{ in } [first, last)\} \cup \{x\}$ which is sorted with respect to an extended strict weak order defined on Q in the manner shown in the proposition above (with $L = \{e \text{ in } [first, last) : comp(e, x)\}$, $U = \{e \text{ in } [first, last) : !comp(x, e)\}$). So, any lookup algorithm operating on $[first, last)$ with input values of type T is formally a valid algorithm when input values are taken from Q and the proper strict weak order, with which $comp$ is compatible, is observed.