

Doc No: P0339r0
Date: 2016-05-31
Audience: LEWG
Authors: Pablo Halpern, Intel Corp. <phalpern@halpernwrightsoftware.com>
Dietmar Kühl <<mailto:dkuhl@bloomberg.net>>

polymorphic_allocator<void> as a vocabulary type

Contents

1	Abstract	1
2	Motivation.....	2
3	Proposal Overview.....	3
4	Alternatives Considered	5
5	Future directions	6
6	Formal Wording	6
6.1	Document Conventions	6
6.2	Undo changes to uses-allocator construction	6
6.3	Remove erased_type from the TS	6
6.4	Changes to std::experimental::function.....	6
6.5	Changes to type-erase allocator.....	8
6.6	Definition of polymorphic_allocator<void>	9
6.7	Changes to class template promise	11
6.8	Changes to class template packaged_task	12
7	References	12

1 Abstract

The `pmr::memory_resource` type, recently added to the C++17 working draft, provides a way to control the memory allocation for an object without affecting its compile-time type – all that is needed is for the object’s constructor to accept a pointer to `pmr::memory_resource`. The `pmr::polymorphic_allocator<T>` adaptor class allows memory resources to be used in all places where allocators are used in the standard: uses-allocator construction, scoped allocators, type-erased allocators, etc.. For many classes, however, the `T` parameter does not make sense.

In this paper, we propose an explicit specialization of `pmr::polymorphic_allocator` for use as a vocabulary type. This type meets the requirements of an allocator in the standard, but is easier to use in contexts where it is not necessary or desirable to fix the allocator type at compile time. The use of `pmr::polymorphic_allocator<void>` also simplifies the definition of *uses-allocator construction* in the TS and situations where allocator type-erasure would otherwise be used, including in `std::function`.

This proposal is targeted for the next release of the Library Fundamentals technical specification.

2 Motivation

Consider the following class that works like `vector<int>`, but with a fixed maximum size determined at construction:

```
class IntVec {
    std::size_t m_size;
    std::size_t m_capacity;
    int *      m_data;
public:
    IntVec(std::size_t capacity);
    : m_size(0), m_capacity(capacity), m_data(new int[capacity]) { }
    ...
};
```

Suppose we want to add the ability to choose an allocator. One way would be to make the allocator type be a compile-time parameter:

```
template <class Alloc = std::allocator<int>> class IntVec ...
```

But that has changed our simple class into a class template, and introduced all of the complexities of writing classes with allocators, including the use of `allocator_traits`. The constructor for this class template looks like this:

```
IntVec(std::size_t capacity, Alloc alloc = {} )
    : m_size(0), m_capacity(capacity), m_alloc(alloc)
    , m_data(std::allocator_traits<Alloc>::allocate(m_alloc, capacity)) { }
```

Our next attempt removes the templating by using `pmr::memory_resource` to choose the allocation mechanism at run time instead of at compile time, thus avoiding the complexities of templates and ensuring that all `IntVec` objects are of the same type:

```
IntVec(std::size_t capacity,
        std::pmr::memory_resource *memrsrc = std::pmr::get_default_resource())
    : m_size(0), m_capacity(capacity), m_memrsrc(memrsrc)
    , m_data(memrsrc->allocate(capacity*sizeof(int), alignof(int))) { }
```

This solution works very well in isolation, but suffers from a number of drawbacks:

1. Does not conform to the Allocator concept

The pointer type, `std::pmr::memory_resource*`, does not meet the requirements of an allocator, and so does not fit into the facilities within the standard designed for allocators, such as *uses-allocator construction* (section 20.9.7.2 in the standard working draft, N4582).

The original proposal for `memory_resource`, [N3916](#), included modifications to the definition of *uses-allocator construction* in order to address this deficiency. Those changes were not added to the C++17 working draft with the rest of the Fundamentals TS version 1

2. Lack of reasonable value-initialization

The result of default-initialization of a pointer is indeterminate, and the result of value initialization is a null pointer, neither of which is a useful value for storing in the

class. The programmer must explicitly call `std::pmr::get_default_resource()`, as shown above. It is easily forgotten and is verbose.

3. Danger of null pointers

Any time you pass a pointer to a function, you must contend with the possibility of a null pointer. Either you forbid it (ideally with a precondition check or `assert`), or you handle it some special way (i.e., by substituting some default). Either way, there is a chance of error.

4. Inadvertent reseating of the memory resource

Idiomatically, neither move assignment nor copy assignment of an object using an allocator or memory resource should move or copy the allocator or memory resource. With rare exceptions, the memory resource used to construct an object should be the one used for its entire lifetime. Changing the resource can result in a mismatch between lifetime of the resource and the lifetime of the object that uses it. Also, assigning to an element of a container would result in breaking the homogenous use of a single allocator for all elements of that container, which is crucial to safely and efficiently applying algorithms like `sort` that swap elements within the container. Raw pointers encourage blind moving or copying of member variables during assignment, which can be dangerous.

Issues 2, 3, and 4 were addressed by another paper, [P0148](#), which proposed a new type, `memory_resource_ptr`, that provided a default constructor, and which was not assignable. This proposal, however, was withdrawn in Jacksonville when we (the authors of that paper) discovered that there was a simpler and more complete solution possible without introducing a completely new type: by using `polymorphic_allocator`. That discovery was the genesis of this paper.

3 Proposal Overview

We observed that a `polymorphic_allocator` object, which is basically a wrapper around a `memory_resource` pointer, can be used just about anywhere that a raw `memory_resource` pointer can be used, but does not suffer from the drawbacks listed above. Consider a minor rewrite of the `IntVec` class (above):

```

class IntVec {
public:
    typedef std::pmr::polymorphic_allocator<int> allocator_type;

private:
    std::size_t    m_size;
    std::size_t    m_capacity;
    allocator_type m_alloc;
    int *          m_data;
public:
    IntVec(std::size_t capacity, allocator_type alloc = {} );
        : m_size(0), m_capacity(capacity), m_alloc(alloc)
        , m_data(alloc.allocate(capacity)) { }
    ...
};

```

Let's consider the deficiencies of using a raw `memory_resource` pointer, one by one, to see how this new approach compares to the previous one:

1. The definition of the `allocator_type` nested type and the constructor taking a trailing allocator argument allows `IntVec` to play in the world of *uses-allocator construction*, including being passed an allocator when inserted into a container that uses a `scoped_allocator_adaptor`.
2. Value-initializing the allocator causes the default memory resource to be used, simplifying the default allocator argument and reducing the chance of error. If `IntVec` had a default constructor, the allocator would, again, use the default memory resource, with no effort on the part of the programmer.
3. A `polymorphic_allocator` is not a pointer and cannot be null. Attempting to construct a `polymorphic_allocator` with a null pointer violates the preconditions of the `polymorphic_allocator` constructor. This contract can be enforced by a single contract assertion in the `polymorphic_allocator` constructor, rather than in every client.
4. [P0335](#) proposes that the assignment operators for `polymorphic_allocator` should be deleted. If this recommendation is accepted, then the problem of accidentally reseating the allocator would no longer exist for `polymorphic_allocator`. The deleted assignment operators would prevent the incorrect assignment operations from being generated automatically, forcing the programmer to define them, hopefully with the correct semantics. See [P0335](#) for more details.

The above list shows that `polymorphic_allocator` can be used idiomatically to good effect. The novel feature of this paper is not this idiomatic use, therefore, but a new specialization for `polymorphic_allocator<void>`. Unlike `std::allocator<void>`, which does not actually meet the requirements of an allocator, `polymorphic_allocator<void>` is designed to be a complete allocator type. It is similar to `polymorphic_allocator<char>`, but has certain features to conveniently expose the capabilities of the underlying `memory_resource` pointer.

In addition to normal allocator functions, `polymorphic_allocator<void>` provides the following features:

- Being completely specialized, `polymorphic_allocator<void>` does not behave like a template, but like a class. This fact can prevent inadvertent template bloat in client types.
- It can allocate objects of any type without needing to use `rebind`. Allocating types other than `value_type` is common for node-based and other non-vector-like containers.
- It can allocate objects on any desired alignment boundary. For example, `VecInt` might choose to align its data array on a SIMD data boundary.
- It provides member functions to allocate and construct objects in one step.
- It provides a good alternative to type erasure for types that don't have an allocator template argument. Note that an important part of this proposal is to simplify `std::function` to avoid the problematic two-dimensional type erasure that has caused problems since C++11.

In addition to the definition of `polymorphic_allocator<void>` itself, we propose the following significant simplifications to the memory section of the Library Fundamentals TS:

- Remove changes to the definition of *uses-allocator construction* and the `uses_allocator` trait. (Section 2 of the TS is completely removed.)
- Rewrite the **Type-erased allocator** section in terms of `polymorphic_allocator<void>` instead of `memory_resource*`. Eliminate the `erased_type` struct.
- Eliminate type-erased allocators from the `function` class template, replacing it with `polymorphic_allocator<void>`. (Note that the type-erased allocator was not implemented by any major standard-library supplier.)
- Update `promise` and `packaged_task` to use the new type-erased allocator idiom.

4 Alternatives Considered

[P0148](#) proposed a new type, `memory_resource_ptr`, which provided many of the benefits described for `polymorphic_allocator<void>`. The `memory_resource_ptr` type did not, however, conform to *allocator requirements* and did less to smooth the integration of `memory_resource` into the allocator ecosystem than does `polymorphic_allocator<void>`. P0148 was withdrawn in favor of this proposal.

5 Future directions

We should consider using `polymorphic_allocator` in the interface to `std::experimental::any`.

6 Formal Wording

6.1 Document Conventions

All section names and numbers are relative to the **March 2016 draft of the Library Fundamentals TS, N4584**. Note that major sections of the TS have been moved into the C++17 WD. Section numbers are, therefore, subject to significant change.

Existing working paper text is indented and shown in dark blue. Edits to the working paper are shown with ~~red strikeouts for deleted text~~ and green underlining for inserted text within the indented blue original text.

Comments and rationale mixed in with the proposed wording appears as shaded text.

Requests for LWG opinions and guidance appear with light (yellow) shading. It is expected that changes resulting from such guidance will be minor and will not delay acceptance of this proposal in the same meeting at which it is presented.

6.2 Undo changes to `uses_allocator` construction

Remove section 2 from the TS, which would have made changes to sections 20.7.7.1, [allocator.uses.trait] and 20.7.7.2 [allocator.uses.construction] of the standard.

6.3 Remove `erased_type` from the TS

Remove section 3.1 [utility] from the TS, which defines `struct erased_type`. The changes to type-erased allocators, below, make this `struct` no longer necessary.

6.4 Changes to `std::experimental::function`

In section 4.1 [header.functional.synop] of the TS, remove the specialization of `uses_allocator` from the end of the <functional> synopsis:

```
template<class R, class... ArgTypes, class Alloc>  
struct uses_allocator<experimental::function<R(ArgTypes...)>, Alloc>{
```

In section 4.3 [func.wrap.func] of the TS, modify `allocator_type` and all of the constructors that take an allocator in `std::experimental::function`:

```
template<class R, class... ArgTypes>  
class function<R(ArgTypes...)> {  
public:  
    typedef R result_type;  
    typedef T1 argument_type;  
    typedef T1 first_argument_type;  
    typedef T2 second_argument_type;
```

```

typedef erased_type pmr::polymorphic_allocator<void> allocator_type;

function() noexcept;
function(nullptr_t) noexcept;
function(const function&);
function(function&&);
template<class F> function(F);
template<class A> function(allocator_arg_t,
                           const A_allocator_type&) noexcept;
template<class A> function(allocator_arg_t,
                           const A_allocator_type&, nullptr_t) noexcept;
template<class A> function(allocator_arg_t,
                           const A_allocator_type&, const function&);
template<class A> function(allocator_arg_t,
                           const A_allocator_type&, function&&);
template<class F, class A> function(allocator_arg_t,
                                     const A_allocator_type&, F);

```

And replace `get_memory_resource()` with `get_allocator()`:

```

pmr::memory_resource* get_memory_resource();
allocator_type get_allocator() const noexcept;
};

```

In sections 4.2.1 [func.wrap.func.con] and 4.2.2 [func.wrap.func.mod], eliminate all references to type erasure and memory resources:

4.2.1 function construct/copy/destroy [func.wrap.func.con]

When a function constructor that takes a first argument of type `allocator_arg_t` is invoked, the second argument ~~is treated as a type-erased allocator (8.3)~~ shall be a polymorphic allocator (C++14 §20.11.3 [memory.polymorphic_allocator.class] or LFTS §8.6 [memory.polymorphic_allocator.class]). A copy of the allocator argument is used to allocate memory, if necessary, for the internal data structures of the constructed function object, otherwise `pmr::polymorphic_allocator<void>{} is used`. If the constructor moves or makes a copy of a function object (C++14 §20.9), including an instance of the `experimental::function` class template, then that move or copy is performed by *using-allocator construction* with allocator ~~`get_memory_resource()`~~ `get_allocator()`.

~~In the following descriptions, let `ALLOCATOR_OF(f)` be the allocator specified in the construction of function `f`, or `allocator<char>()` if no allocator was specified.~~

```
function& operator=(const function& f);
```

Effects: `function(allocator_arg, ALLOCATOR_OF(*this) get_allocator(), f).swap(*this);`

Returns: `*this`.

```
function& operator=(function&& f);
```

Effects: `function(allocator_arg, ALLOCATOR_OF(*this) get_allocator(), std::move(f)).swap(*this);`

Returns: `*this`.

```
function& operator=(nullptr_t) noexcept;
```

Effects: If `*this != nullptr`, destroys the target of this.

Postconditions: `!(*this)`. The `memory_resource` allocator returned by `get_memory_resource()` after the assignment is equivalent to the `memory_resource` allocator before the assignment. [*Note:* the address returned by `get_memory_resource()` might change — *end note*]

Returns: `*this`.

```
template<class F> function& operator=(F&& f);
```

Effects: `function(allocator_arg, ALLOCATOR_OF(*this)get_allocator(), std::forward<F>(f)).swap(*this);`

Returns: `*this`.

Remarks: This assignment operator shall not participate in overload resolution unless `declval<decay_t<F>&&>()` is Callable (C++14 §20.9.11.2) for argument types `ArgTypes...` and return type `R`.

```
template<class F> function& operator=(reference_wrapper<F> f);
```

Effects: `function(allocator_arg, ALLOCATOR_OF(*this)get_allocator(), f).swap(*this);`

Returns: `*this`.

4.2.2 function modifiers [func.wrap.func.mod]

```
void swap(function& other);
```

Requires: `*this->get_memory_resource() == *other.get_memory_resource()
this->get_allocator() == other.get_allocator().`

Effects: Interchanges the targets of `*this` and `other`.

Remarks: The allocators of `*this` and `other` are not interchanged.

Add a new section describing the `get_allocator()` function:

```
allocator_type get_allocator() const noexcept;
```

Returns: A copy of the allocator specified at construction, if any; otherwise a copy of `allocator_type{} evaluated at the time of construction of this object.`

6.5 Changes to type-erase allocator

Make the following changes to section 8.3 Type-erased allocator [memory.type.erased.allocator]:

8.3 Type-erased allocator [memory.type.erased.allocator]

A type-erased allocator is an allocator or memory resource, `alloc`, used to allocate internal data structures for an object `X` of type `C`, but where `C` is not dependent on the type of `alloc`. Once `alloc` has been supplied to `X`

(typically as a constructor argument), a copy of alloc can be retrieved from X only as a pointer ~~rptr~~ of static type `std::experimental::pmr::memory_resource*` (8.5) via an object named (for exposition) `pmr_alloc` of type `pmr::polymorphic_allocator<void>` (C++14 §20.11.3 [memory.polymorphic_allocator.class] or LFTS §8.6 [memory.polymorphic_allocator.class]). The process by which ~~rptr~~ `pmr_alloc` is ~~computed~~ initialized from alloc depends on the type of alloc as described in Table 13:

Table 13 — Initialization of type-erased allocator

If the type of alloc is	then the value of rptr <code>pmr_alloc</code> at X construction time is
non-existent — no alloc specified	The value of <code>experimental::pmr::get_default_resource()</code> at the time of construction <u>value initialized.</u>
<code>nullptr_t</code>	The value of <code>experimental::pmr::get_default_resource()</code> at the time of construction <u>value initialized.</u>
a pointer type convertible to <code>pmr::memory_resource*</code>	<code>static_cast<experimental::pmr::memory_resource*>(alloc)</code> <u>initialized with alloc</u>
<code>pmr::polymorphic_allocator<U></code>	<u>initialized with</u> <code>alloc.resource()</code>
any other type meeting the Allocator requirements (C++14 §17.6.3.5)	<u>initialized with</u> a pointer to a value of type <code>experimental::pmr::resource_adaptor<A></code> where A is the type of alloc. rptr <code>pmr_alloc</code> remains valid only for the lifetime of X.
None of the above	The program is ill-formed.

Additionally, class C shall meet the following requirements:

- `C::allocator_type` shall be identical to `std::experimental::erased_type pmr::polymorphic_allocator<void>`.
- `X.get_memory_resource()` ~~X.get_allocator()~~ returns ~~rptr~~ `pmr_alloc`.

6.6 Definition of `polymorphic_allocator<void>`

In section 8.4 [memory.resource.synop] of the TS, add the `void` specialization of `polymorphic_allocator` to the synopsis for `<experimental/memory_resource>`:

```
template <class Tp> class polymorphic_allocator;
template <> class polymorphic_allocator<void>;
```

Editorial note: Since section 8.4 has been copied into the C++17 WD, it is not clear if it will remain in the TS. Some editorial re-arrangement of these changes may be necessary.

In section 8.6.1 [memory.polymorphic_allocator.overview], add the following specialization immediately after the general definition of `polymorphic_allocator<Tp>` (underline highlighting omitted for ease of reading):

The specialization of `polymorphic_allocator<void>` provides additional member functions for managing memory in bytes, providing convenient access to the facilities in the underlying `memory_resource`. Except where specified, the definition of member functions and constructors is identical to that of the primary template.

```

template <>
class polymorphic_allocator<void> {
    memory_resource* m_resource; // For exposition only

public:
    typedef void value_type;

    polymorphic_allocator() noexcept;
    polymorphic_allocator(memory_resource* r);

    polymorphic_allocator(const polymorphic_allocator& other) = default;

    template <class U>
        polymorphic_allocator(const polymorphic_allocator<U>& other) noexcept;

    polymorphic_allocator&
        operator=(const polymorphic_allocator& rhs) = delete;

    void* allocate(size_t bytes);
    void* allocate(size_t bytes, size_t alignment);
    void deallocate(void* p, size_t bytes);
    void deallocate(void* p, size_t bytes, size_t alignment);

    template <class T, class CtorArgs...>
        T* new_object(CtorArgs&&... ctor_args);
    template <class T>
        void delete_object(T* p);

    template <class T, class... Args>
        void construct(T* p, Args&&... args);

    // Specializations for pair using piecewise construction
    template <class T1, class T2, class... Args1, class... Args2>
        void construct(pair<T1,T2>* p, piecewise_construct_t,
            tuple<Args1...> x, tuple<Args2...> y);
    template <class T1, class T2>
        void construct(pair<T1,T2>* p);
    template <class T1, class T2, class U, class V>
        void construct(pair<T1,T2>* p, U&& x, V&& y);
    template <class T1, class T2, class U, class V>
        void construct(pair<T1,T2>* p, const std::pair<U, V>& pr);
    template <class T1, class T2, class U, class V>
        void construct(pair<T1,T2>* p, pair<U, V>&& pr);

    template <class T>
        void destroy(T* p);

    // Return a default-constructed allocator (no allocator propagation)
    polymorphic_allocator select_on_container_copy_construction() const;

    memory_resource* resource() const;
};

```

Add descriptions for specialized member functions after section 8.6.3 [memory.polymorphic_allocator.mem] (underline highlighting omitted for ease of reading):

8.6.4 polymorphic_allocator<void> specialized functions [memory.polymorphic_allocator.voidalloc]

Relative to the primary template, the specialization of polymorphic_allocator with a void template parameter has slightly different semantics for the allocate and deallocate member functions, as well as providing additional overloads of those functions. The void specialization also provides two additional member functions, new_object and delete_object.

```
void* allocate(size_t bytes);
```

Returns: equivalent to m_resource->allocate(bytes, m), where m is the smallest alignment suitable for any non-over-aligned object with a size of bytes.

```
void* allocate(size_t bytes, size_t alignment);
```

Returns: equivalent to m_resource->allocate(bytes, alignment).

```
void deallocate(void* p, size_t bytes);
```

Effects: Equivalent to m_resource->deallocate(p, bytes, m), where m is the smallest alignment suitable for any non-over-aligned object with a size of bytes.

Throws: Nothing.

```
void deallocate(void* p, size_t bytes, size_t alignment);
```

Effects: Equivalent to m_resource->deallocate(p, bytes, alignment).

Throws: Nothing.

```
template <class T, class CtorArgs...>  
T* new_object(CtorArgs&&... ctor_args);
```

Effects: Allocates and constructs an object of type T as if by

```
void* p = allocate(sizeof(T), alignof(T));  
try {  
    new (p) T(std::forward<CtorArgs>(ctor_args)...);  
} catch (...) {  
    m_resource->deallocate(p, sizeof(T), alignof(T));  
    throw;  
}
```

Returns: The address of the newly constructed object (i.e., p).

```
template <class T>  
void delete_object(T* p);
```

Effects: Equivalent to p->~T(); deallocate(p, sizeof(T), alignof(T)).

6.7 Changes to class template promise

Make the following changes to the class definition of promise in section 11.2 [futures.promise] of the TS, consistent with the change in type-erased allocators:

```

template <class R>
class promise {
public:
    typedef erased_typepolymorphic_allocator<void> allocator_type;
...
    pmr::memory_resource* get_memory_resource();
    pmr::polymorphic_allocator<void> get_allocator() const noexcept;
};

```

6.8 Changes to class template packaged_task

Make the following changes to the class definition of `packaged_task` in section 11.3 [futures.task], consistent with the change in type-erased allocators:

```

template <class R, class... ArgTypes>
class packaged_task<R(ArgTypes...)> {
public:
    typedef erased_typepolymorphic_allocator<void> allocator_type;
...
    pmr::memory_resource* get_memory_resource();
    pmr::polymorphic_allocator<void> get_allocator() const noexcept;
};

```

7 References

[N4584](#) *Working Draft, C++ Extensions for Library Fundamentals, Version 2*, Geoffrey Romer, editor, 2016-03-08.

[N3916](#) *Polymorphic Memory Resources - r2*, Pablo Halpern, 2014-02-14.

[P0148](#) *memory_resource_ptr: A Limited Smart Pointer for memory_resource Correctness*, Pablo Halpern and Dietmar Kühl, 2015-10-14.

[P0335](#) *Delete operator= for polymorphic_allocator*, Pablo Halpern, 2016-05.