

Doc No: N3858
Supersedes: N3722
Date: 2014-01-19
Reply to: Niklas Gustafsson <niklas.gustafsson@microsoft.com>
Deon Brewis <deonb@microsoft.com>
Herb Sutter <hsutter@microsoft.com>

Resumable Functions

While presenting a proposal that can be adopted or rejected in isolation, this document is related to N3857. The reader is advised to read both as a unit and to consider how the two build on each other for synergy. Reading them in their assigned numeric order is advised.

New in this version: minor changes to examples, explanatory text; make the section on generators clearer.

1. The Problem

The direct motivation for introducing resumable functions is the increasing importance of efficiently handling I/O in its various forms and the complexities programmers are faced with using existing language features and existing libraries.

The motivation for a standard representation of asynchronous operations is outlined in N3857 and won't be repeated here. The need for language support for resumable functions arises from the inherent limitations of the library-based solution described in that proposal.

Taking a purely library-based approach to composition of asynchronous operations means sacrificing usability and versatility: the development of an asynchronous algorithm usually starts with a synchronous, imperative expression of it, which is then manually translated into an asynchronous equivalent. This process is quite complex, akin to the reformulation of an imperative algorithm in a pure functional language, and the resulting code may be difficult to read.

A pure library-based approach leads to object lifetime management complexities and thus a different way of designing the objects that are to be used by and with asynchronous operation. An even bigger problem with solutions that avoid language support for asynchrony is the lack of ability to compose asynchronous operations using the rich variety of traditional control-flow primitives.

Consider this example, using the modified version of `std::future<T>` introduced in N3857:

```
future<int> f(shared_ptr<stream> str)
{
    shared_ptr<vector<char>> buf = ...;
    return str->read(512, buf)
        .then([] (future<int> op)           // lambda 1
            {
```

```

        return op.get() + 11;
    });
}

future<void> g()
{
    shared_ptr<stream> s = ...;
    return f(s).then([s](future<int> op)        // lambda 2
    {
        s->close();
    });
}

```

When `g()` is activated, it creates its stream object and passes it to `f`, which calls the `read()` function, attaches a continuation (lambda 1) to its result, and then returns the result of the `then()` member function call. After the call to `f()` returns, `g` attaches a continuation (lambda 2) to the result, after which it returns to its caller.

When the read operation finishes, lambda 1 will be invoked from some context and its logic executed, resulting in the operation returned from `f()` completing. This, in turn, results in lambda 2 being invoked and its logic executed. If you were to set a breakpoint in either of the lambdas, you would get very little context or information on how you got there, and a debugger would be hard-pressed to make up for the lack of information.

To make matters worse, the above code does not consider that the futures returned by `read()` and `f()` may already be completed, making the attachment of a continuation lambda unnecessary and expensive. To squeeze out all the performance possible, the code will wind up being quite complex.

In order to properly account for the fact that the lifetime of local objects passed to asynchronous operations (`s`, `buf`) is different from the scope in which they are declared the programmer has to allocate them on the heap and find some means of managing them (using a `shared_ptr<>` is often sufficient).

Contrast this with how the same algorithm, just as efficient and asynchronous, would look when relying on resumable functions:

```

future<int> f(stream str) resumable
{
    shared_ptr<vector<char>> buf = ...;
    int count = await str.read(512, buf);
    return count + 11;
}

future<void> g() resumable
{
    stream s = ...;
    int pls11 = await f(s);
    s.close();
}

```

Not only is this simpler, it is more or less identical to a synchronous formulation of the same algorithm. Note, in particular, that there is no need to manage the lifetime of locally declared objects by allocating them on the heap: the compiler takes care of the lifetimes, allowing the programmer to write code that looks almost identical to synchronous code.

The library-based approach gets even more complicated when our example includes control-flow such as conditional evaluation and/or loops. The language-based approach allows control-flow to remain identical to the synchronous formulation, including the use of try-catch blocks and non-reducible constructs such as goto, continue, and break.

The following example illustrates this.

While iterative composition is not covered in N3857, it is nevertheless expected that libraries need to provide higher-order compositional constructs to mimic the behavior of such things as loops and conditional expressions/statements. With the help of a `do_while()` construct (not described in N3857), we get this code:

```
auto write =
  [&buf](future<int> size) -> future<bool>
  {
    return streamW.write(size.get(), buf).then(
      [](future<int> op){ return op.get() > 0; });
  };

auto flse = [](future<int> op){ return future::make_ready_future(false);};
auto copy = do_while(
  [&buf]() -> future<bool>
  {
    return streamR.read(512, buf)
      .then(
        [](future<int> op){ return (op.get() > 0) ? write : flse;});
  });
```

With resumable functions, the same code snippet would be:

```
int cnt = 0;

do
{
  cnt = await streamR.read(512, buf);

  if ( cnt == 0 ) break;

  cnt = await streamW.write(cnt, buf);
} while (cnt > 0);
```

It is not necessarily a lot shorter, but undoubtedly easier to comprehend, more or less identical to a synchronous formulation of the same algorithm. Further, no special attention needs to be paid to object lifetimes.

Resumable functions are motivated by the need to adequately address asynchronous operations, but are not actually tied to the proposal for a standard representation of such operations. The definition in this proposal can be used with any types that fit the described patterns. For example, resumable functions may be used to implementing a system for fully synchronous co-routines.

That said, throughout this document, `future<T>` will be used as the primary example of usage and implementation.

2. The Proposal

2.1 Terminology

A resumable function is a function that is capable of split-phase execution, meaning that the function may be observed to return from an invocation without producing its final logical result or all of its side-effects. This act is defined as the function suspending its execution. The result returned from a function when it first suspends is a placeholder for the eventual result: i.e. a `future<T>` representing the return value of a function that eventually computes a value of type T.

After suspending, a resumable function may be resumed by the scheduling logic of the runtime and will eventually complete its logic, at which point it executes a return statement (explicit or implicit) and sets the function's result value in the placeholder.

It should thus be noted that there is an asymmetry between the function's observed behavior from the outside (caller) and the inside: the outside perspective is that function returns a value of type `future<T>` at the first suspension point, while the inside perspective is that the function returns a value of type T via a return statement, functions returning `future<void>/shared_future<void>` behaving somewhat different still.

Within the function, there are zero or more suspension points. A resumable function *may* pause when it reaches a suspension point. Given control-flow, it may or may not be the case that a resumable function actually reaches a suspension point before producing a value (of type T); conversely, a given suspension point may be reached many times during the execution of a function, again depending on its control-flow.

A resumable function *may* continue execution on another thread after resuming following a suspension of its execution.

2.2 Declaration and Definition

Resumable functions are identified by decorating the function definition and declaration with the `resumable` keyword following the formal argument list. In the grammar productions for function definitions and lambda expressions, the `resumable` keyword is placed right before the exception-specification.

Any function or lambda that can legally return a `future<T>/shared_future<T>` or `future<void>/shared_future<void>` may be a resumable function, regardless of scope.

2.3 Restrictions

Resumable functions cannot use a variable number of arguments. For situations where `varargs` are necessary, the argument unwrapping may be placed in a function that calls a resumable function after doing the unwrapping of arguments.

The return type of a resumable function must be `future<T>` or `shared_future<T>`. The restrictions on `T` are defined by `std::future`, not this proposal, but `T` must be a copyable or movable type, or 'void.' It must also be possible to construct a variable of `T` without an argument; that is, it has to have an accessible (implicit or explicit) default constructor if it is of a class type.

Await expressions may not appear within the body of an exception handler and should not be executed while a lock on any kind is being held by the executing thread.

2.4 Suspension Points

Within a resumable function, its resumption points are uniquely identified by the presence of the 'await', which is treated as a keyword or reserved identifier within resumable functions. It is used as a unary operator, which pauses the function and resumes it when its operand is available. The expression denoted by the 'await' keyword is called an "await expression."

The unary operator has the same precedence as unary operator '!' boolean. Therefore, these statements are equivalent:

```
int x = (await expr) + 10;  
int x = await expr + 10;
```

The operator may take an operand of any type that is `future<U>/shared_future<U>` or convertible to `future<U>/shared_future<U>`.

If `U` is not 'void,' it produces a value of type `U` by waiting for the future to be filled and returning the value returned by the future's 'get()' function. If `U` is 'void,' the await expression must be the term of an expression statement. That is, it cannot be the operand of another expression (since it yields 'void').

The `U` used in the operand of any given await expression in a function does not correspond to or have to be related to, the operands of other await expressions or the return type of the function itself: the types of what the function consumes (using await) and produces (using return) are independent of each other.

2.5 Return Statements

A resumable function produces its final value when executing a return statement or, only in the case of `future<void>/shared_future<void>` as the function return type, it reaches its end of the function without executing a return statement.

For example:

```
future<int> abs(future<int> i_op) resumable
{
    int i = await i_op;
    return (i < 0) ? -i : i;
}
```

Within a resumable function declared to return `S<T>`, where `S` is a future or `shared_future` and `T` is not 'void', *a return statement should accept an operand of type T*. In the case of a resumable function declared to return `future<void>` or `shared_future<void>`, a return statement should be of the form "return;" or be omitted. In such a function, reaching the end of the function represents termination of the function and filling of the result.

3. Interactions and Implementability

3.1 Interactions

Keywords

The proposal uses two special identifiers as keywords to declare and control resumable functions. These should cause no conflict with existing, working, code.

In the case of `resumable`, it appears in a place where it is not currently allowed and should therefore not cause any ambiguity. Introducing the use of `resumable` as an identifier with a special meaning only when it appears in that position is therefore not a breaking change.

In the case of `await`, it is globally reserved but meaningful only within the body of a function or within the argument of a `decltype()` expression.

A possible conflict, but still not a breaking change, is that the identifiers may be in use by existing libraries. In the case of `resumable` the context should remove the possibility of conflict, but 'await' is more difficult. When used with a parenthesized operand expression, it will be indistinguishable from a call to a function 'await' with one argument.

A second possible non-breaking conflict is if there is a macro of the name 'await,' in which preprocessing will create problems.

A quick search of the header files for the Microsoft implementation of the standard C++ and C libraries, the Windows 7 SDK, as well as a subset of the Boost library header files show that there are no such conflicts lurking within those common and important source bases.

Overload Resolution

From the perspective of a caller, there is nothing special about a resumable function when considering overload resolutions.

Expression Evaluation Order / Operator Precedence

This proposal introduces a new unary operator, only valid within resumable functions and `decltype()` expressions. The precedence of the operator is the same as that of the `!` operator, i.e. Boolean negation.

3.2 Implementation, #1: Resumable Side Stacks

The following implementation sketch is not intended to serve as a reference implementation. It is included for illustrative purposes only.

With this prototype implementation each resumable function has its own side stack. A side stack is a stack, separate from the thread's main stack, allocated by the compiler when a resumable function is invoked. This side stack lives beyond the suspension points until logical completion. The implementation is relatively simple and does not require a compiler to significantly alter the body of the function, but it requires the allocation of a separate contiguous stack for each outstanding resumable function.

Consider the following:

```
void foo()
{
    future<bool> t = bar();
}
future<bool> bar() resumable
{
    do_work();
    //possible suspension points and synchronous work here
    return true;
}
```

When `foo` (a non-resumable function) is invoked, the current thread's stack is allocated. Next, when the function `bar()` is executed, the `resumable` keyword is recognized and the compiler creates a side stack, allocating it appropriately. Next, the compiler switches from the thread's current stack, to the new side stack, this is the push side stack operation. Depending on the implementation, the side stack can run on the same thread or a new thread. `do_work()`, a synchronous operation executes on the resumable function's side stack.

Next consider the following resumable function which has a suspension point denoted by the `await` keyword:

```
void foo()
{
    future<bool> t = bar();
    do_work();
    bool b = t.get();
}

future<bool> bar() resumable
{
    do_work();
    await some_value;
}
```

```
        do_more_work();  
        return true;  
    }
```

From above, we know that resumable function `bar()` is currently running on its own side stack. After completing the synchronous work on the call stack, the function reaches a possible suspension point which is indicated with `'await'`. If the value at that point is not ready, `bar()` is suspended and a `future<T>`, (where `T` is the type of the function's return statement), is returned to the calling function. The side stack is popped and the compiler switches back to the thread's main stack. Function `foo()` continues to proceed until it gets blocked waiting for the future's state to be ready. After some time, when `some_value` is fulfilled, the function `bar()` resumes from where it left off (the suspension point) and `bar()`'s side stack is pushed. When the end of the resumable function is reached, the previously returned future's state is set to ready and its value is set to the function's return value. Once the resumable function `bar()` is completed, the side stack is popped off and deleted, and the compiler switches back to the main thread.

3.3 Implementation #2: Heap-allocated activation frames

A second example implementation requires considerably more “heavy lifting” from the compiler, but does not require allocation of a large, contiguous stack for the function. Activation frames for resumable functions are allocated in heap-based storage and are reference-counted.

3.2.1 Function Definition

The definition of a resumable function results in the definition of the locals frame structure and the added function, into which the body of the resumable method is moved before being transformed. The resumable method itself is more or less mechanically changed to allocate an instance of the frame structure, copy or move the parameters, and then call the added method.

It's worth pointing out that the frame structure used in this example is an artifact of our attempt to represent the transformations using valid C++ code. A “real” implementation would allocate a suitably large byte array and use that for storage of local variables and parameters. It would also run constructors and destructors at the correct point in the function, something that our source-code implementation cannot.

The definition:

```
future<int> f(future<double> g) resumable { return ceil(await g); }
```

results in:

```
struct _frame_f
{
    int _state;
    future<int> _resultF;
    promise<int> _resultP;
    _frame_f(future<double> g) : g(g), _state(0)
    {
        _resultF = future<int>(_resultP);
    }
    future<double> g;
};
future<int> f(double g)
{
    auto frame = std::make_shared<_frame_f>(g);
    _res_f(frame);
    return frame->_resultF;
}
void _res_f(const std::shared_ptr<_frame_f> &frame)
{
    return ceil(await frame->g);
}
```

Note that the body of the `_res_f()` function represents artistic license, as it represents a transitional state of the original body. It still needs to be transformed into its final form, as described in the next section.

3.2.2 Function Body

There are four main transformations that are necessary, not necessarily performed in the order listed: a) space for local variables needs to be added to the frame structure definition, b) the function prolog needs to branch to the last resumption point, c) await expressions need to be hoisted and then transformed into pause/resumption logic, and d) return statements need to be transformed to modify the `_result` field of the frame.

3.2.3 Allocating Storage

All variables (and temporaries) with lifetimes that statically span one or more resumption points need to be provided space in the heap-allocated structure. In the hand-translated version, their lifetimes are extended to span the entire function execution, but a real, low-level implementation must treat the local variable storage in the frame as just storage and not alter the object lifetimes in any way.

The heap-allocated frame is reference-counted so that it can be automatically deleted when there are no longer any references to it. In this source-code implementation, we're using `std::shared_ptr<>` for reference counting. Something more tailored may be used by a real implementation.

An implementation that cannot easily perform the necessary lifetime analysis before allocating space in the frame should treat all local variables and formal parameters of a resumable function as if their lifetimes span a resumption point. Doing so will increase the size of the heap-allocated frame and decrease the stack-allocated frame.

3.2.4 Function Prolog

The `_state` field of the frame structure contains an integer defining the current state of the function. A function that has not yet been paused always has `_state == 0`. With the exception of the initial state, there is a one-to-one correspondence between state identities and resumption points. Except for the value 0, the actual numerical value assigned to states has no significance, as long as each identity uniquely identifies a resumption point.

Each state is associated with one label (branch target), and at the prolog of the function is placed the equivalent of a switch-statement:

```
void _res_f(std::shared_ptr<_frame_f> frame)
{
    switch(frame->_state)
    {
        case 1: goto L1;
        case 2: goto L2;
        case 3: goto L3;
        case 4: goto L4;
    }
}
```

In the hand-coded version, special care has to be taken when a resumption point is located within a try-block; an extra branch is required for each try block nesting a resumption point: first, the code branches to just before the try-block, then we allow the code to enter the block normally, then we branch again:

```
void _res_f(std::shared_ptr<_frame_f> frame)
{
    switch(frame->_state)
    {
        case 1: goto L1_1;
        case 2: goto L2;
        case 3: goto L3;
        case 4: goto L4;
    }

L1_1:
    try
    {
        switch(frame->_state)
        {
            case 1: goto L1;
        }
L1:
        ...
    }
}
```

Depending on the implementation of try-blocks, such a chaining of branches may not be necessary in a low-level expression of the transformation.

3.2.5 Hoisting 'await' Expressions

Before transformation, each resumption point needs to be in one of these two forms:

```
x = await expr;
await expr;
```

In other words, embedded await operators need to be hoisted and assigned to temporaries, or simply hoisted in the case of void being the result type. The operand 'expr' also needs to be evaluated into a temporary, as it will be used multiple times in the implementation, before and after the resumption point.

Note that await expressions may appear in conditional (ternary) expressions as well as in short-circuit expressions, which may affect their hoisting in some compiler implementations.

3.2.6 Implementing await expressions

In our hand-coded implementation, the hoisted expression "t = await g;" is transformed thus:

```
if ( !frame->g.ready() )
{
    frame->_state = 1;
    frame->g.then(
        [=] (future<double> op)
```

```

        {
            __res_f(frame);
        });
    return;
}

```

L1:

```
t = frame->g.get();
```

In the case of 'await g' being used as the expression of an expression statement, i.e. the value is thrown away, the compiler must emit a call to 'wait()' after the resumption. Calling wait() when the result is not used gives the runtime a chance to raise any propagated exceptions that may otherwise go unobserved.

3.2.7 Transforming 'return' Statements

Return statements are simply transformed into calls to set the value contained in the _result objects, or overwrite it with a new object:

```

// return ceil(await g);

if ( frame->_state == 0 )
    frame->_resultF = make_ready_future<int>(ceil(t));
else
    frame->_resultP.set(ceil(t));

```

The test for _state == 0 is done to establish whether the function has ever been paused or not. If it has not, it means that the caller will not yet have been passed back the result instance and it is therefore not too late to replace it with a more efficient value holder. *It is an optimization and an implementation is not required to test for this condition.*

4. Technical Specifications / Proposed Wording for the C++ Standard

5.1 resumable

In 1.9 p15 add:

Every evaluation in the calling function (including other function calls **to non-resumable functions, as defined in 8.3.5/15**) that is not otherwise specifically sequenced before or after the execution of the body of the called function is indeterminately sequenced with respect to the execution of the called function. **The execution of a resumable function may appear to interleave with the calling function. When a resumable function is suspended at the await keyword, a placeholder of the eventual result is returned and the calling function continues its execution. After suspending, a resumable function may be resumed and will eventually complete its logic, at which point it executes a return statement filling in the value of the placeholder.**

To footnote 9 add:

In other words, function executions do not interleave with each other, with the exception of resumable functions, which may interleave with their caller.

In 2.11 p2 add to Table 3 – Identifiers with special meaning:

resumable

In 5.1.2 p1:

lambda-declarator:

(*parameter-declaration-clause*) mutable_{opt} resumable-specification_{opt}
exception-specification_{opt} attribute-specifier-seq_{opt} trailing-return-type_{opt}

In 5.1.2 p4 add:

If a *lambda-expression* does not include a *lambda-declarator*, it is as if the *lambda-declarator* were (). The lambda return type is auto, which is replaced by the *trailing-return-type* if provided and/or deduced from return statements as described in 7.1.6.4. If the lambda has the resumable specifier and no *trailing-return-type* is provided, the return type is *future<T>*, where T is the type deduced from return statements. [Example:

```
auto x1 = [](int i){ return i; }; // OK: return type is int
auto x2 = []{ return { 1, 2 }; }; // error: deducing return type from
braced-init-list
int j;
auto x3 = []()->auto&& { return j; }; // OK: return type is int&
auto x4 = []() resumable {
    int i = await read_stream();
    return i; }; // OK: return type is future<int>
—end example ]
```

In 8.0 p4:

parameters-and-qualifiers:

(*parameter-declaration-clause*) attribute-specifier-seq_{opt} cv-qualifier-seq_{opt}
ref-qualifier_{opt} resumable-specification_{opt} exception-specification_{opt}

In 8.3.5 p1:

D1 (*parameter-declaration-clause*) cv-qualifier-seq_{opt}
ref-qualifier_{opt} resumable-specification_{opt} exception-specification_{opt} attribute-specifier-seq_{opt}

In 8.3.5 p2:

D1 (*parameter-declaration-clause*) cv-qualifier-seq_{opt}
ref-qualifier_{opt} resumable-specification_{opt} exception-specification_{opt} attribute-specifier-seq_{opt} trailing-return-type

In 8.3.5 p2:

The *resumable-specification* is not a part of the function type.

In 8.4.1 p2:

D1 (*parameter-declaration-clause*) *cv-qualifier-seq_{opt}*

ref-qualifier_{opt} **resumable-specification_{opt}** *exception-specification_{opt}* *attribute-specifier-seq_{opt}* *trailing-return-type_{opt}*

In 8.3.5 add p15

15. The function specified with an resumable specifier is a *resumable function* (see definition below):

resumable-specification:

resumable

- A resumable function is a function that returns a placeholder for an eventually available value.
- It may be observed by its caller to return without filling the placeholder with a value.
- If the resumable function terminates, it will fill the placeholder with either a value or an exception.
- Some side-effects of the resumable function may be delayed until after its return.
- The caller of this function can resume its work without waiting for the *resumable function* to finish.

[Example:

```
int work1(int value);
void f(int value) {
    future<int> f = g(value);
    work2();
}

future<int> g(value) resumable {
    result = await std::async( [= ] {return work1(value)}); //
    return result;
}
```

- end example]

- If a declaration contains a *resumable-specification* then every subsequent redeclaration shall also contain a *resumable-specification*.
- A *resumable-specification* shall not appear in a typedef declaration or alias-declaration.
- A function declaration with the `resumable` specifier must return `future<T>` or `shared_future<T>`.
- The *parameter-declaration-clause* of a *resumable function* shall not terminate with an ellipsis.
- The result returned from a function when it first suspends is a placeholder for the eventual result: i.e. a `future<T>` representing the return value of a function that eventually computes a value of type T.
- The *parameter-declaration-clause* may terminate with a function parameter pack.
- The `resumable` keyword is only a reserved keyword in *the resumable-specification* position of a function's declaration. It has no special meaning if used elsewhere.

5.2 await

In 2.12 add to Table 4 – Keywords:

await

In 5.3 p1 add:

Expressions with unary operators group right-to-left.

unary-expression:

postfix-expression
++ cast-expression
-- cast-expression
unary-operator cast-expression
sizeof unary-expression
sizeof (type-id)
sizeof ... (identifier)
alignof (type-id)
noexcept-expression
new-expression
delete-expression

unary-operator: one of

* & + - ! ~ **await**

In 5.3 add:

5.3.8 await unary operator

A unary operator expression of the form:

await cast-expression

1. The `await` operator is only valid within *resumable functions* [8.3.5] and `decltype()` expressions.
2. When the `await` operator is applied to an operand in a *resumable function*, the execution of the *resumable function* is suspended until the operand completes.
3. The `cast-expression` shall be of class type `future<T>` or `shared_future<T>` or shall be implicitly convertible to `future<T>` or `shared_future<T>`.
4. `await` is globally reserved but meaningful only within the body of a function or within the argument of a `decltype()` expression.
5. The `await` operator shall not appear within the body of an exception handler¹.
6. The `await` operator shall not be executed while a lock [30.4.2] is being held.
7. The result of `await` is of type `T`, where `T` is the return type of the `get` function of the `future` or `shared_future` object. If `T` is `void`, then the `await` expression cannot be the operand of another expression.

¹ The motivation for this is to avoid interfering with existing exception propagation mechanisms, as they may be significantly (and negatively so) impacted should `await` be allowed to occur within exception handlers.

5.3 Return

In 6.6.3, add a paragraph 4:

- 4 Within a resumable function declared to return `future<T>` or `shared_future<T>`, where `T` may be `void`, any `return` statement shall be treated as if the function were declared to return a value of type `T`.

5. Generalizations

This section is meant to demonstrate two ways in which the proposal for resumable functions can eventually be generalized if there are use cases to motivate it.

5.1 Types other than `future<T>/shared_future<T>`

The proposal places strict requirements on the return type of resumable functions, as well as the operand of `await` expressions. These restrictions may be relaxed, independently, in the following ways.

The operand of the unary operator `await` can be any type `S<<T>>` (“`S` eventually holding a `T`”) with meets the following conditions:

1. `S` has a parameter-less function `get` which will either eventually produce a value of type `T`, or throw an exception.
2. `S` has a function `then` accepting a single-parameter function object taking a parameter `S<<T>>`, `S<<T>>&`, or `const S<<T>>`. Once passed to `then`, the value held by the parameter must be immediately available for retrieval by a call to `get`.
3. Optionally, if `S<<T>>` has a `bool`-returning function `is_ready()` indicating whether a value is already held, the implementation of `await` can be made more efficient.

What motivates #1 is that the implementation of `await` needs a well-defined way to retrieve the value or the exception eventually held by the operand. #2 is what allows the implementation to avoid blocking, as `then` serves to signal the availability of a held value or exception. #3 allows an implementation to avoid pausing the resumable function when a value is already available in the operand.

Similarly, the set of allowable return types of resumable functions may be expanded to include any type `S<<T>>` that meets the following conditions:

1. `S` contains a `typedef ... _promise_type;` which indicates a type used to fill the placeholder returned by the resumable function when it reaches a return statement or propagates an exception.
2. `_promise_type` must have at least a function `set_value(T)` and a function `set_exception(exception_ptr)`.
3. There must be an implicit or explicit conversion from `S::promise_type` to `S`.

#1 is motivated by the need to identify the type that should be used to set the value or exception held by the returned `S<<T>>` object. #2 identifies what is needed from `S<<T>>::promise_type` in order to

actually set the value or exception, and #3 establishes the canonical means of constructing an instance of `S<<T>>` from the corresponding `S<<T>>::promise_type` instance.

5.2 Generator Functions

While we are solely proposing adding the capabilities of asynchronous waiting at this point, it is useful to demonstrate that the resumable function concept may be expanded to include generator functions, which require the same kind of interruption / resumption of function logic as asynchronous wait. Such a demonstration should allow the reader to see how resumable functions are more broadly useful than what is covered by the scenario(s) in section 1.

Consider the need to sequentially process a sequence of values, integers for the sake of simplicity. We do not have the need to access elements of the sequence in random order, or to modify the sequence in place. The logic of producing the sequence we shall call the *generator*, and the logic processing it is the *processor*.

With existing means, we could utilize a vector to represent the sequence and pass it from the generator to the processor, but it would mean allocating memory for the whole sequence before starting the traversal. With a big sequence, that could be a lot of memory, possibly unbounded. To avoid this, we could use two threads and rely on some form of bounded queue to pass data from the generator to the processor. However, that means introducing concurrency to achieve something inherently serial.

What we desire is the ability to produce each value of the sequence when it's needed but no sooner, i.e. lazily. Generator functions, a concept available in languages like C# and Python, address this need elegantly and efficiently. Writing general lazy sequence generator is surprisingly hard without language support and surprisingly simple with it. With a couple of minor modifications to the resumable functions definition, they could be available in C++, too.

Consider a function `lazy_tform()`, which takes an iterator pair and a function, applies the function to each element in the input sequence, and lazily produces a sequence. This is very similar to `std::transform()`, except for the lazy part – whereas `transform()` will generate (write to the output iterator it is passed) the entire output sequence before returning, `lazy_tform()` will only generate one value at a time, and only when asked to.

```
template<typename Iter>
sequence<int> lazy_tform(Iter beg, Iter end, std::function<int(int)> func) resumable
{
    for (auto iter = beg; iter != end; ++iter)
    {
        yield func(*iter);
    }
}
```

In this example, we have introduced two new concepts: the `sequence<T>` type, and the 'yield' statement. `sequence<T>` is a collection type that supports only iteration, the iterator type it provides is an input iterator.

A caller of `lazy_tform()` may use the sequence as such (the use of old-style for-statement is for the purpose of being very explicit about what happens):

```
auto rng = range(5, 15);

auto squares = lazy_tform(rng.begin(), rng.end(), [](int x) { return x*2; });

for (auto iter = squares.begin(); iter != squares.end(); ++iter)
{
    std::cout << " Next: " << *iter << std::endl;
}
```

Here, `range()` is assumed to be a function generating some collection of integers. It could, for example, be a resumable function generating the sequence of numbers between its two arguments:

```
sequence<int> range(int low, int high) resumable
{
    for (int i = low; i <= high; ++i)
    {
        yield i;
    }
}
```

If so, we have effectively chained two sequence generators together – the one generating a 1-step sequence from 5 to 15, and the one generating the sequence of squares of its input sequence. It is important to realize that the generator of the first sequence, the `range()` function, is only generating data as it is being asked for by the second generator, the `lazy_tform()` function, which is only doing so as its caller is advancing the iterator using `operator++`.

What happens when `lazy_tform()` is called?

1. The function is called and immediately returns a `sequence<int>` instance before reaching the first statement in the function body.
2. When `squares.begin()` is called, the sequence creates an iterator, and executes the resumable function to the first `yield`.
3. When `operator*` is called on the iterator, the value generated by the most recently executed `yield` statement is returned.
4. When `operator++` is called on the iterator, the resumable function is resumed and executed to the next `yield` statement.
5. Once the function is advanced to the point where it returns (or falls off the end), the iterator reaches its end position and the sequence is terminated.

Note that *there is no concurrency involved*, everything happens on the thread of the caller. The functionality of generator functions is entirely unrelated to concurrency, it is about on-demand generation of sequence values.

With resumable functions thus expanded to include generator functions, it is necessary to ponder what happens when we want to mix `yield` and `await` in the same function – one is used to pause the function while generating data, one is used to pause the function while waiting for data to be available.

The answer is, unfortunately, not as simple as the resumable function eventually producing a sequence (`future<sequence<T>>`) or producing a sequence of eventual values (`sequence<future<T>>`). Since the generation of any value in the output sequence may be interleaved with `await`, it is the `begin` and `operator++` functions that need to produce eventual values.

Such a resumable function would look like this (imagining our reading data from a distant table):

```
async_sequence<table_row> get_cloud_table_rows(int low, int high) resumable
{
    for (int i = low; i <= high; ++i)
    {
        auto table_row = await table_client.get_row(i);
        yield table_row;
    }
}
```

And the caller would look like this, using `explicit` for syntax:

```
auto rows = get_cloud_table_rows(100,150);
for (auto iter = await rows.begin(); iter != rows.end(); await ++iter)
{
    auto row = *iter;
    ...
}
```

Since the existing language support for range-based `for` loops does not have any knowledge of asynchronous iterators (they don't currently exist), *a new version of range-based for would have to be added*. The specific syntax is not obvious, but something like this might work:

```
for await (auto row : get_cloud_table_rows(100,150))
{
    ...
}
```

Going back to the points made at the beginning of this section, inclusion of this discussion on generator functions here was for the purpose of demonstrating that resumable functions is a general enough concept to be extended to other uses. It is our intent to bring a proposal for adding generator functionality to C++ in the near future.