

Postconditions odr-using a parameter modified in an overriding function

Timur Doumler (papers@timur.audio)
Joshua Berne (jberne4@bloomberg.net)

Document #: P3484R1
Date: 2024-11-07
Project: Programming Language C++
Audience: SG21, EWG

Abstract

This paper considers the case where an overridden function odr-uses a non-reference function parameter in its postcondition assertion, and then an overriding function drops `const` on the declaration of that parameter, rendering the postcondition assertion in the overridden function essentially meaningless. We propose several alternatives for how to address this problem in the Contracts MVP [[P2900R10](#)].

This paper is the first part of a trilogy of papers dealing with known issues in the Contracts MVP [[P2900R10](#)] regarding postconditions odr-using non-reference function parameters:

- [[D3484R1](#)] Postconditions odr-using a parameter modified in an overriding function;
- [[D3487R0](#)] Postconditions odr-using a parameter that may be passed in registers;
- [[D3489R0](#)] Postconditions odr-using a parameter of dependent type.

These issues should be considered together, and ideally resolved in a consistent way.

1 Background

[[P2900R10](#)] Section 3.4.4 specifies that if a non-reference function parameter is odr-used in a postcondition assertion, that function parameter must be `const` on *all* declarations of that function, otherwise the program is ill-formed. Importantly, this requirement on all declarations of the function includes its *defining* declaration; thus, the implementation of the function can be assumed to not modify the value of that parameter.

The rationale for this rule is that allowing the implementation to modify a parameter that is later used to check the postcondition would render that check meaningless. Consider the following function declaration:

```
// Returns: a number guaranteed to be greater or equal to the number passed in  
int f(const int i) post (r: r >= i);
```

Now, if we could drop the `const` on the defining declaration of `f`, we could write a dummy implementation of `f` as follows:

```
int f(int i) { // i is not const
    i = 0;
    return i;
}
```

Such an implementation blatantly fails to satisfy the postcondition of `f`, and is therefore incorrect. However, this defect cannot be caught by checking the postcondition assertion of `f`, because its predicate is now `0 >= 0` which always evaluates to `true`:

```
void test() {
    int j = f(3); // no contract violation detected
    // precondition does not hold — j is now 0, which is smaller than 3!
}
```

In the above scenario, the implementation of `f` explicitly modifies the parameter. However, an even more dangerous variant of this scenario is when such modifications happen *implicitly*. In particular, if the parameter is not `const` on the defining declaration of `f`, it may be returned by value, which will perform an implicit move, i.e. the postcondition assertion would observe the parameter object in a moved-from state.

For these reasons, [P2900R10] makes the above implementation of `f` ill-formed: `const` cannot be dropped from any function parameter odr-used in a postcondition assertion.

Further, [P2900R10] Section 3.4.4 specifies that if a non-reference function parameter is odr-used in a postcondition assertion, and that function is implemented as a coroutine, the program is ill-formed, *even if* all such function parameters have been declared `const` by the user.

The rationale for this rule is that a coroutine will move-from its function parameters to initialise the parameter copies in the coroutine frame, and therefore the function parameters of a coroutine are effectively not `const`, even if declared as such by the user (see [P3387R0] for discussion). This case is thus notionally similar to the previous case where `const` was dropped from the parameter declaration in the implementation.

Together, these rules provide a static guarantee that, if a parameter is odr-used in a postcondition assertion, this parameter object *will not be modified* between the call to a function and the evaluation of its postcondition assertions. An additional benefit of this design is that a static analysis tool can reason about the value of a function parameter in a postcondition without having to take into account the function definition, which will often be either too complicated to analyse or unavailable at the call site.

Reference parameters are excluded from the above restrictions because references refer to objects declared elsewhere, and the value of those objects when the function call completes are still relevant and available to the caller because those objects can still have outside references known to the caller. There is no expectation that the value will remain unchanged after the function body has executed, and many functions that pass an object through a modifiable reference do so with the exact intention of modifying that object; therefore, such parameters would not be used in a postcondition assertion with that expectation in mind.

2 The problem

Let us now slightly modify the example above by making `f` a virtual function:

```
struct Base {
    virtual int f(const int i) post (r: r >= i);
};
```

Note that if we override a virtual function, C++ allows dropping `const` from the parameter declaration in the overriding function, and [P2900R10] currently does not have any provision to make such an override ill-formed:

```
struct Derived : Base {
    int f(int i) override; // OK
};
```

This means that we can implement `Derived::f` such that it modifies the value of the parameter:

```
int Derived::f(int i) {
    i = 0;
    return i;
};
```

[P2900R10] Section 3.5.3 specifies the semantics of precondition and postcondition assertions on virtual functions: in a virtual function call, the function contract assertions of both the statically called function `Base::f` and the final overrider `Derived::f` are checked (see [P3097R0] for discussion). However, if we now call `Derived::f` through a reference to `Base`, we are in for an unpleasant surprise:

```
void test(Base& b) {
    int j = b.f(3); // no contract violation detected
    // precondition does not hold — j is now 0, which is smaller than 3!
}

int main() {
    Derived d;
    test(d);
}
```

In the program above, even if the postcondition assertion of `Base::f` is checked, the fact that the implementation of `Derived::f` does *not* satisfy the postcondition of `Base::f` is not caught, because the parameter `i` has been modified in `Derived::f`, rendering the postcondition assertion of `Base::f` meaningless. This is exactly the scenario that the specification in [P2900R10] Section 3.4.4 seeks to avoid; nevertheless, the code above is well-formed according to the current specification. The more dangerous variant of this scenario that a parameter might be modified *implicitly*, for example moved-from when returned by value, is also possible and well-formed.

Unlike other cases that may lead to a `const` parameter being modified in the function body, namely when `const` is dropped on a subsequent declaration of the *same* function, or when the function is implemented as a coroutine (which are both ill-formed in [P2900R10]), in this case `Base::f` does not know whether there are any functions overriding it, or how the corresponding parameter in those overriding functions is declared.

3 Possible solutions

We are aware of four possible options for dealing with this problem. These are, from most to least restrictive:

- V1. Disallow odr-using any non-reference function parameter in a postcondition assertion that applies to a virtual function, regardless of whether that parameter is declared `const`, unless that function is marked `final` or is a member function of a class marked `final`.
- V2. Require that if a non-reference parameter is odr-used in a postcondition assertion on a virtual function, that parameter must also be declared `const` in every declaration of every overriding function.

- V3. Allow overriding functions to drop `const` from a non-reference function parameter, with no special provision, i.e., “you get what you get”: in the virtual function call above, which invokes `Derived::f`, the postcondition check on `Base::f` succeeds even though `Derived::f` does not satisfy the postcondition of `Base::f`.
- V4. Allow overriding functions to drop `const` from a non-reference function parameter, but make it undefined behaviour to actually modify a parameter object in an overriding function if that parameter is a non-reference parameter declared `const` in an overridden function.

We enumerated the options with a “V” prefix (for “virtual”), to distinguish them from the options from [D3487R0] that have an “R” prefix (for “registers”) and the options from [D3489R0] that have a “D” prefix (for “dependent”).

Below we discuss the tradeoffs of each option.

Option V1

Option V1 would catch the bug in the example above and is the most conservative choice. This choice is consistent with the choice we made for postcondition assertions on coroutines in [P2900R10]: if a function odr-uses a non-reference parameter in its postcondition, and that parameter is declared `const`, but there might be some other reason why the implementation of the function may modify that parameter anyway, the program is ill-formed. One such reason is that the function is a coroutine and thus the parameter will be modified by the underlying coroutine machinery. Another such reason is that the function is a non-`final` virtual function and thus an overriding function could modify the parameter.

Just like in the coroutine case, for non-`final` virtual functions the workaround would be to use a postcondition capture, once this post-MVP feature becomes available (see [P3098R0]):

```
struct Base {
    virtual int f(const int i) post (r: r >= i);    // error: cannot odr-use parameter i
};

struct Base {
    virtual int f(const int i) post [i] (r: r >= i); // OK: explicitly capturing i by copy
}
```

However, Option V1 has several downsides.

First, having to capture any parameter in order to odr-use it in the postcondition assertion means that we have to pay the cost of the copy. For coroutines, making that copy is the *only* way to get access to the pre-moved-from value of a parameter in the postcondition; on the other hand, for virtual functions, the copy will be unnecessary in most cases, i.e. we would be paying for the freedom to modify that parameter in an override, but we will most likely not make use of that freedom.

Further, unlike coroutines, any modifications to the parameters must be done explicitly in an overriding function and they will not happen implicitly as part of non-obvious language machinery. This suggests that there is a much lower risk of the user accidentally getting it wrong. For coroutines, there is no way the user could write an implementation that avoids the parameter modification; on the other hand, for virtual functions, there is a very simple way: just do not drop `const` from the parameter declaration on overriding functions, and do not modify that parameter in the function’s implementation.

Finally, in today’s C++ ecosystem, virtual functions are more pervasive than coroutines, and postcondition assertions have more known use cases for virtual functions than for coroutines (the usefulness of postcondition assertions on coroutines is fundamentally limited due to the nature of

coroutines in C++). Entirely disallowing the ability to odr-use parameters in the postcondition assertion of a non-`final` virtual function in the first version of Contracts for C++ that we ship could noticeably hamper the usability of the feature. Shipping postcondition captures as proposed in [P3098R0] in the same version could somewhat mitigate but not fully remove the friction.

Overall, Option V1 might therefore be a disproportionately harsh measure for a relatively rare problem.

Option V2

Option V2 would also catch the bug in the example above and is a less restrictive choice than Option V1: it would make only cases ill-formed where the parameter can actually be modified (when the overriding function actually drops the `const`). This option makes the behaviour for overriding functions consistent with the behaviour of subsequent declarations of the *same* function: if we redeclare a function, *or* an overload of that function, and drop `const` on a parameter in that declaration, the program is ill-formed. Option V2 thus seems more appealing than Option V1.

However, the tradeoff is that Option V2 could also lead to remote code breakage, which directly violates Design Principle 15 of [P2900R10], “No Client-Side Language Break”. In particular, adding a postcondition to a virtual function that odr-uses a `const` parameter would remotely break any client code that overrides that function and yet has not added the `const`, or any override that is implemented as a coroutine.

We anticipate that particularly the former case would be a frequent problem. Today, hardly anyone adds `const` to the declaration of a non-reference parameter as it has essentially no meaning; with [P2900R10], anyone who wishes to odr-use a parameter in a postcondition will *have* to add the `const`, breaking any overriding functions.

The crucial difference to the coroutine case is that providing a definition for a given function that makes the function a coroutine is local, not remote code, whereas overriding a function can happen in an entirely different component of the program. The possibility of such remote code breakage due to the introduction of Contracts could hamper their adoption and make releasing low-level libraries with newly introduced function-contract assertions vastly more difficult.

Option V3

Option V3 is the status quo¹ in [P2900R10]. It avoids the usage limitations imposed by Option V1 and the remote code breakage imposed by Option V2. However, this option too has tradeoffs.

One downside of Option V3 is that, unlike Options V1 and V2, it does not actually catch the bug. As a partial remedy, an implementation could easily issue a *warning* if an overriding function drops the `const` on a parameter odr-used in the postcondition assertion of an overridden function. The crucial difference to the coroutine case is that the implementation *knows* that the function is virtual at the point of declaration. While we cannot normatively mandate such a warning, we can add a non-normative recommended practice note to the wording that such a warning be issued.

Another downside of Option V3 compared to Options V1 and V2 is that it would mean abandoning the static guarantee of the current [P2900R10] design that, if a parameter is odr-used in a postcondition assertion, that parameter object *will not be modified* between the call to a function and the evaluation of its postcondition assertions.

¹Since the current specification in that paper does not contain any special rules for parameter declarations on overriding functions, dropping `const` in such declarations is currently allowed in [P2900R10] and “you get what you get”.

It follows that a static analysis tool can no longer reason about the value of a function parameter in a postcondition, at least not without taking into account which function will be selected by virtual dispatch and what the body of that function does (both are typically unknowable at the call site). This greatly reduces the amount of useful information that a static analysis tool could extract from a postcondition assertion at the call site of a virtual function.

Option V4

Option V4 is the option that was chosen in the past by C++2a Contracts [P0542R5]. However, we can immediately discard it as unviable. Not only would this option fail to catch the bug in the example above, but it is also actively user-hostile: the mere act of adding a postcondition assertion to a program would have the potential to introduce new undefined behaviour, even if its predicate were written correctly. This directly violates Design Principle 13 of [P2900R10], “Explicitly Define All New Behaviour”.

4 Proposal

We believe that Options V1 — V3 are all worth considering and the tradeoffs of each option are sufficiently clear. We therefore propose all three options to determine which option has more consensus in SG21.

Note that choosing Option V1 would leave the door open to adopting Option V2 or Option V3 without breaking changes at some point in the future, while Option V2 could only be evolved towards Option V3 but not Option V1, and Option V3 could not be evolved towards either of the other options.

Note further that [D3487R0] — which deals with a different problem regarding parameters in postconditions — proposes the more extreme options R1 (removing postcondition assertions from [P2900R10] altogether), R2 (disallowing odr-use of *any* parameters in the postcondition assertions of *any* functions), and R3 (disallowing odr-use of *non-reference* parameters in the postcondition assertions of *any* functions). Options R1 — R3 from [D3487R0] would remove the issue discussed in this paper and should therefore be considered alongside Options V1 — V3 as possible solutions.

5 Wording

The proposed wording changes are relative to the wording proposed in [P2900R10].

Option V1

Modify [dcl.contract.func] as follows:

If the predicate of a postcondition assertion of a function odr-uses ([basic.def.odr]) a non-reference parameter of that function, that parameter shall be declared `const` and shall not have array or function type; if the function is virtual it shall be marked with the *virt-specifier* `final` (see [class.virtual]) or it shall be a member function of a class with the *class-virt-specifier* `final` (see [class.pre]). [Note: This requirement applies even to declarations that do not specify the *postcondition-specifier*. Arrays and functions are still usable when declared with the equivalent pointer types ([dcl.fct]). — end note] [Example: [...] — end example]

Option V2

Modify [dcl.contract.func] as follows:

If the predicate of a postcondition assertion of a function f odr-uses ([basic.def.odr]) a non-reference parameter of ~~that function~~ f , that parameter and the corresponding parameter on any functions that override f shall be declared `const` and shall not have array or function type. [*Note*: This requirement applies even to declarations that do not specify the *postcondition-specifier*. Arrays and functions are still usable when declared with the equivalent pointer types ([dcl.fct]). — *end note*] [*Example*: [...] — *end example*]

Option V3

Modify [dcl.contract.func] as follows:

If the predicate of a postcondition assertion of a function odr-uses ([basic.def.odr]) a non-reference parameter of that function, that parameter shall be declared `const` and shall not have array or function type. [*Note*: This requirement applies even to declarations that do not specify the *postcondition-specifier*. Arrays and functions are still usable when declared with the equivalent pointer types ([dcl.fct]). — *end note*] [*Example*: [...] — *end example*]

Recommended practice: Implementations should issue a diagnostic when an overriding function omits `const` from the declaration of a non-reference parameter whose corresponding parameter in an overridden function is odr-used in a postcondition assertion of that overridden function.

Acknowledgements

Thanks to John Lakos and Oliver Rosten for their review of the paper.

Revision history

- **R0**, 2024-10-31: Initial version as presented to SG21
- **R1**, 2024-11-04: Incorporated SG21 feedback, various other fixes

Bibliography

- [D3484R1] Timur Doumler and Joshua Berne. Postconditions odr-using a parameter modified in an overriding function. <https://wg21.link/d3484r1>, 2024-11-01.
- [D3487R0] Timur Doumler and Joshua Berne. Postconditions odr-using a parameter that may be passed in registers. <https://wg21.link/d3487r0>, 2024-11-01.
- [D3489R0] Timur Doumler and Joshua Berne. Postconditions odr-using a parameter of dependent type. <https://wg21.link/d3489r0>, 2024-11-01.
- [P0542R5] G. Dos Reis, J. D. Garcia, J. Lakos, A. Meredith, N. Myers, and B. Stroustrup. Support for contract based programming in C++. <https://wg21.link/p0542r5>, 2018-06-08.
- [P2900R10] Joshua Berne, Timur Doumler, and Andrzej Krzemiński. Contracts for C++. <https://wg21.link/p2900r10>, 2024-10-12.

- [P3097R0] Timur Doumler, Joshua Berne, and Gašper Ažman. Contracts for C++: Support for Virtual Functions. <https://wg21.link/p3097r0>, 2024-04-15.
- [P3098R0] Timur Doumler, Gašper Ažman, and Joshua Berne. Contracts for C++: Postcondition captures. <https://wg21.link/p3098r0>, 2024-10-14.
- [P3387R0] Timur Doumler, Joshua Berne, Iain Sandoe, and Peter Bindels. Contract assertions on coroutines. <https://wg21.link/p3387r0>, 2024-10-09.