

# A Principled Approach to Open Design Questions for Contracts

Document #: P2932R0  
Date: 2023-09-13  
Project: Programming Language C++  
Audience: SG21 (Contracts)  
Reply-to: Joshua Berne <[jberne4@bloomberg.net](mailto:jberne4@bloomberg.net)>

## Abstract

SG21 has made significant progress toward producing a complete design for a Contracts facility MVP. As work proceeds for completing this feature proposal, open questions must be considered and answered such that the feature will eventually have widespread adoption throughout the entire C++ ecosystem. Fundamental design principles that help guide such decisions for the Contracts facility are presented in this paper as are proposals that address most (if not all) of the open questions affecting the design of the Contracts facility MVP.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Principles</b>	<b>2</b>
2.1	CCAs Check Plain-Language Contracts	3
2.2	Program Semantics Are Independent of Chosen CCA Semantic	3
2.3	Zero Overhead	4
2.4	Choose Ill-Formed to Enable Flexible Evolution	5
<b>3</b>	<b>Proposals</b>	<b>6</b>
3.1	Trivial Special Member Functions	6
3.2	Implicit Lambda Captures	9
3.3	Compile-Time Evaluation	11
3.4	Virtual Functions	16
3.5	Coroutines	18
3.6	Contracts on First Declarations	19
<b>4</b>	<b>Conclusion</b>	<b>21</b>
<b>5</b>	<b>Acknowledgments</b>	<b>22</b>

# Revision History

Revision 0 (September 2023 mailing)

- Original version of the paper for discussion during an SG21 telecon

## 1 Introduction

The Contracts MVP being developed by SG21<sup>1</sup> is progressing smoothly toward a complete proposal according to the agreed-upon schedule.<sup>2</sup> SG21 has reviewed several papers and has reached a general consensus on what will happen when a contract-checking annotation (CCA) is evaluated.

- [P2751R1], “Evaluation of *Checked* Contract-Checking Annotations,” clarified what content will be allowed in the predicate of a CCA and how the results of that evaluation will be interpreted.
- [P2811R7], “Contract-Violation Handlers,” described a link-time replaceable contract-violation handler to allow for customized handling of contract violations at run time.
- [P2834R1], “Semantic Stability Across Contract-Checking Build Modes,” established that having preconditions and postconditions evaluated outside the `noexcept` guarantee of a function that is marked `noexcept` should not be considered.
- [P2877R0], “Contract Build Modes, Semantics, and Implementation Strategies,” made the choice of semantics for a CCA implementation defined and, in particular, showed that such a choice might vary between evaluations and is thus not a compile-time property of a CCA.

Taking the above resolutions together, on top of the existing core MVP proposal, produces an almost complete specification for Contracts as a C++ language feature.<sup>3</sup> Some edge cases and other design points, however, still need to be addressed.<sup>4</sup>

To motivate the choices we intend to make when suggesting answers to these design questions, we will begin by identifying fundamental principles that will codify our design intentions for the C++ Contracts facility. These principles start with basic goals and build up to concrete requirements that all proposals should meet (or maximize) when making engineering decisions for what should become part of C++.

## 2 Principles

Each principle proposed in this paper comes from the author’s experience, follows from SG21’s earlier principles, or aims to maximize some aspect of the quality of the Contracts facility that eventually becomes adopted by the C++ Standard.

---

<sup>1</sup>See [P2521R4] and the forthcoming [P2900R0].

<sup>2</sup>See [P2695R1].

<sup>3</sup>I.e., barring an actual syntax for the feature, which is the subject of a separate, ongoing series of discussions orthogonal to the issues in this paper. For this paper, we will be using the C++20 attribute-like syntax (see [P0542R0]).

<sup>4</sup>See [P2834R1] for the initial explorations of some of these solutions based on the principles described by that paper. See also [P2896R0], which enumerates the open design questions addressed by this paper.

## 2.1 CCAs Check Plain-Language Contracts

When programmers of any language set out to write a subroutine (or function), they have in mind some notion of what will happen when the subroutine is invoked (for a given subset of syntactically valid inputs). The part of that notion to which the programmer is willing to commit as unchanging with changes to the implementation is what we generally refer to as the function’s *contract*. Hence, for any function that can be relied upon by clients, a *contract* must exist between the callers of that function and its implementers. This contract, which we call the *plain-language contract*, exists as a combination of the explicit (through documentation) and implicit (through choice of programming language, coding style, local conventions, etc) agreements to which both parties agree.

A *plain-language contract* can capture a boundless variety of different conditions, requirements, and promises that might exist between callers and callees. Some of these are prone to being validated with a simple C++ expression, such as `x > 0`, while others are completely beyond the scope of the C++ abstract machine, such as “clients of this function have paid their bills recently.”

A *contract-checking annotation* encodes in C++ an algorithm that can detect when a plain-language contract has been violated. When writing a CCA, that new annotation must be checking some aspect of a plain-language contract to detect violations of the agreement that exists between the caller and callee of the function.

Any CCA that identifies a condition as a defect that is not related to its associated *plain-language contract* is clearly doing something above and beyond checking the requirements of that *plain-language contract*. If, for any kind or placement of CCA, there are ways in which that CCA might be evaluated where the predicate might indicate a violation when the CCA’s associated plain-language contract is not violated or not relevant, then a false-positive bug will have been diagnosed, which would be a defect in the structure of the Contracts facility itself.

Therefore, a CCA, when introduced anywhere, must always evaluate in ways that validate its associated plain-language contract.

### Principle 1: CCAs Check a Plain-Language Contract

Each contract-checking annotation placed on the declaration of a function must, when it is evaluated, identify violations of the *plain-language contract* of that function.

## 2.2 Program Semantics Are Independent of Chosen CCA Semantic

The fundamental conceit of a contract-checking facility is that running a program with CCAs checked at run time should tell us something meaningful about what happens when the same source code is evaluated with the CCA *not* checked at run time. If a design choice made by the contract-checking facility leads to the inability to detect bugs in an unchecked build when enabling checking, then the facility will have failed on a fundamental level.

The ramifications of allowing the choice of semantic (i.e., checked or unchecked) to impact the results of the `noexcept` operator were thoroughly explored in [P2834R1]. This paper clearly showed that the seemingly innocuous decision to allow enabled checking to impact the results of the `noexcept` operator could easily lead to situations in which a small bug remains entirely undetectable by the

contract-checking facility because the act of enabling contract checks at run time alters higher-level control flow to avoid the bug entirely.

The adoption of [P2877R0], which allows the selection of CCA semantics to potentially occur at run time, renders compile-time consideration of the semantic of a CCA moot. Still, all design decisions related to runtime-evaluated CCAs must nonetheless conform to this important stability principle, whereby the current semantic has no impact (apart from evaluating the predicate itself) on any aspect of the behavior of the program itself. On the other hand, compile-time evaluation might still occur and might differ based on implementation-defined build options. Such variance at compile time would again open up the same risks for a Contracts design to fail to satisfy its most fundamental purpose.

Therefore, the fundamental principle of [P2834R1] still stands as one that must guide *all* behaviors of CCAs within any C++ contract-checking facility.

#### Principle 2: Program Semantics Are Independent of Chosen CCA Semantics

The semantic with which a CCA is evaluated must not affect the *proximate* compile-time semantics surrounding that annotation.

Importantly, different build options that alter the semantics chosen for contracts must not lead to distinct, well-formed programs having different behaviors, other than the behaviors that are part of the evaluation of CCAs (and potentially their predicates) themselves.

### 2.3 Zero Overhead

While the first major risk to Contracts being a viable contract-checking facility is failing to detect bugs, the second major risk is failure to be adopted. If projects and enterprises commonly disallow the use of Contracts due to fundamental design flaws, such as unavoidable costs introduced even when building with all CCAs *ignored*, then all SG21's effort to design a product will have been for naught. The mere rumor that adding a CCA might increase runtime overhead would be a huge headwind against rapid adoption of the facility. A natural consequence would be that developers think they have compelling reasons to continue using macro-based contract-checking facilities indefinitely.

Macros (such as `assert`) that expand to nothing when the facility is disabled are inherent to most macro-based contract-checking facilities. This characteristic results in there being no tangible runtime overhead to having CCAs in a program's source code, thus minimizing the factors that might discourage ubiquitous use of Contracts.

Expanding to absolutely nothing does bring with it a fundamental problem that experience has taught us we should strive to avoid with the Contracts facility: *bit rot*. When an unchecked build ignores CCA predicates completely (i.e., it treats them as arbitrary *token soup*), failing to maintain CCA predicates when changes are made to the APIs they use can become commonplace. Teams that do not regularly compile their software with Contracts checked can quickly find themselves *unable* to compile checked builds because of the invalidity of poorly maintained CCA predicates. To avoid bit rot, a highly prudent practice is to continue to syntactically check CCA predicates, which invariably brings with it compile-time side effects, such as template instantiations, and some amount of constant evaluation to validate the predicate's semantic correctness.

Other than those compile-time effects of validating a predicate, no other inherent runtime overhead is associated with having a CCA in source code. This property leads to our next fundamental principle.

**Principle 3: Zero Overhead for Ignored Predicates**

The behavior of code *proximate* to a CCA is *as if* the CCA had been commented out, i.e., removed from the source entirely.

By following this principle, we also reduce (to nothing, we hope) the possibility of users wrapping CCA usage in macros that remove them entirely in some builds, thus reducing the chance of having non-ABI-compatible builds when mixing checked and unchecked code in the same application. This benefit hinges on one of the major gains achieved by adopting [P2877R0]: Mixed-semantic builds are, when supported, required to be valid and cannot be ODR violations, thus greatly increasing the ability to manage contract enablement at various granularities.

## 2.4 Choose Ill-Formed to Enable Flexible Evolution

Early in the process of identifying a path for bringing Contracts back into the C++ draft Standard, SG21 solicited and amassed<sup>5</sup> a wide array of use cases for guiding the scope of our next attempt at a viable MVP for a C++ Contracts facility. Despite valiant efforts to refine the semantics and functionality addressed by the MVP, a large number of these use cases will remain unsupported until more work is done and further consensus has been achieved.

When standardizing a solution, we would be wise to leave room for satisfying the use cases that SG21 will address after an initial feature is integrated into C++. To facilitate that evolution, decisions made today must not prevent migration to a more broadly applicable feature in future revisions.

When deciding how to provide this essential freedom, a few *specification tools* are available to us.

- One choice, which was made for quite a few aspects of the C++20 Contracts facility, is to leave the behavior that is not finalized as undefined behavior. This option has the advantage of allowing implementations the flexibility to experiment with various possible solutions along with the severe disadvantage of all those solutions being nonportable and possibly conflicting. SG21 had early consensus to refrain from using undefined behavior in this fashion because doing so is generally perceived to be a source of reduced safety, not increased flexibility or correctness.
- A middle ground rarely chosen is to simply make certain behaviors *implementation defined*. While generally appearing safer than *undefined behavior*, this approach still results in code that is nonportable when different platforms define distinct and incompatible behaviors for the same constructs.
- Another specification tool available to us is to make what we are unsure of ill-formed. That is, given two or more plausible solutions whose behaviors largely but incompletely overlap, rather than making all of that behavior undefined or even implementation defined, we might

---

<sup>5</sup>See [P1995R1].

choose to define the overlapping behavior and (to strongly discourage inconsistency) to leave the remaining, nonoverlapping behavior ill-formed.

- Finally, when we have consensus on a solution but are unclear whether the solution in question will be implementable on all platforms, we might choose to specify the full solution and make it *conditionally supported*. Using this approach, however, effectively introduces multiple *dialects* of the language on different platforms, which is often an approach the WG21 community has avoided.

In general, given the history of SG21’s consensus to focus on the safety and portability of the Contracts feature being developed, the preferable option is to leave room for evolution by making the available design space ill-formed in the current implementation.

#### Principle 4: Make Undecided Behaviors Ill-Formed

To accommodate use cases that are not yet supported by the current contracts proposal, prefer to keep extensibility flexible by declaring unresolved behavior ill-formed, rather than either undefined (“unsafe”) or implementation defined (nonportable).

## 3 Proposals

Each subsection within this section discusses a distinct open question in the specification of the Contracts MVP and provides a proposed resolution based on the principles we have introduced above.

Some proposals are made of multiple constituent parts. The individual parts throughout a section are not numbered, and a final complete, numbered proposal will be made near the end of each section.

### 3.1 Trivial Special Member Functions

A trivial operation — constructor, destructor, or assignment operator — is one that (1) is generated entirely by the compiler and (2) invokes no user-provided code. The copy and move operations, when trivial, become bitwise copyable (and thus *may* be replaced by, e.g., `memmove`). The default constructor and destructor, when trivial, do nothing.<sup>6</sup>

Providing a body — even an empty one — for a special member function results in that function never being trivial. Defaulting a user-declared special member function (via `= default` on the first declaration) will result in it being trivial as long as it doesn’t need to invoke any user-provided code for a base-class or member object.

The definition of the *trivially copyable* trait has evolved over the years, but the main point is that all the default-generated copy and move constructors and copy and move assignment operators are *trivial* (and hence, can treat the object as pure data) as long as at least one of them is not deleted. For a type to be *trivially copyable*, it must also be *trivially destructible*; i.e., the compiler-generated destructor is a no-op.

<sup>6</sup>See also the mention of this issue in [P2521R4], Section 4.4, “Annotations on trivial ops.,” `{con.trv}`.

Now consider that, to minimize runtime overhead and still get substantial coverage, common practice allows for any type that happens to have one or more (programmatically verifiable) class (object) invariants to assert them in the one place where the flow of control must pass for every constructed object: the destructor. This defensive-checking strategy is particularly effective at catching memory overwrites.<sup>7</sup>

Now imagine we have a trivially copyable value type, such as this heavily elided `Date` class:

```
class Date {
    int d_year;    // [ 1 .. 9999 ]
    int d_month;  // [ 1 .. 12   ]
    int d_day;    // [ 1 .. 31   ]
public:
    static bool isValidYMD(int year, int month, int day);
        // Return true if year/month/day represents a valid date.

    Date(int year, int month, int day) [[ pre: isValidYMD(year, month, day) ]];
        // Create a Date object having a valid date.

    int year() const { return d_year; }
        // ...
};
```

In the `Date` class above, the user-declared value constructor creates a valid `Date` object set to the specified `year`, `month`, and `day`, and in a *checked* build, its precondition check invokes the class member function `isValid` to *ensure* (or perhaps just to *observe*) that the date is, in fact, valid. From then on, the only way to change the value of this object is through the use of its compiler-generated assignment operations.<sup>8</sup>

As previously stated, the `Date` class above is trivially copyable, but let's now add a defaulted destructor having a precondition check that validates its invariants:

```
~Date() = default [[ pre: isValidYMD(d_year, d_month, d_day) ]];
    // Destroy this object.
```

Notice that the precondition itself applies the public class member function, `isValidYMD`, to the private data members, e.g., `d_year`, of the class. The Committee has long since agreed<sup>9</sup> that all CCAs are part of a function's implementation; hence, a CCA that is associated with a member function fairly deserves private access to the data members of the class.

---

<sup>7</sup>Feel free to ask the author for stories about the pain that comes from developers who decide to implement types that overwrite their *own* members by calling `memset(this,0,sizeof(this))` in their constructor bodies, under some misguided belief that doing so improved performance and correctness. In particular, if such a class has any member, such as `std::string`, that has a nontrivial default constructor that does more than just zero-initialize, this ill-fated choice leads to painful-to-diagnose problems that such invariant checking can often readily detect.

<sup>8</sup>Recall that, just by declaring any non-special-member constructor, we suppress even the declaration of the default constructor. All five remaining special member functions, however, are generated as usual.

<sup>9</sup>During a Standards Committee meeting in 2015 that involved CCAs, Bjarne Stroustrup expressed the need for a CCA on a member function to have private access to the implementation of the function's class so that the CCA could write efficient preconditions without having to expose extra functions in a public API that were not directly relevant to clients. This discussion led to a paper, [P1289R1], which was considered and achieved consensus in November 2018. This result has remained the SG21 consensus, as indicated in Section 4.3 of [P2521R4].

If CCAs are *ignored*, the destructor does nothing, and because no user-supplied definition is available, one might reasonably presume that this special member function remains *trivial*. In a *checked* build, however, code provided by the user is expected to run when an object of this type is destroyed. In some sense, this `Date` type is *almost*<sup>10</sup> trivially destructible and thus, at least *notionally*, trivially copyable.

Following Principle 2, i.e., we cannot change the semantics of code based on which semantic CCAs are evaluated with, means that we *cannot* take the approach of treating this `Date` class *as if* it were *trivially destructible* and *trivially copyable* in an *unchecked* build but *not* in a *checked* build.

Once again, we are left with two alternatives.

1. Specify that defaulted special member functions that have associated CCAs are never trivial in any build mode. The implication is that merely having an inactive precondition could substantially impact performance,<sup>11</sup> even in an unchecked (e.g., *ignored*) build.
2. Treat contract checks on trivial functions as *skippable* without notice. That is, if the function itself is considered trivial in an *unchecked* build mode, it will report so in *every* build mode. Libraries may circumvent the execution of these special member functions, even in checked build modes. (Note that we are skipping the *check* itself, not just any side effects that executing the check might have produced.) If, however, the compiler would be eliding invocation of the trivial function, it might nonetheless choose to evaluate the CCAs, followed by the trivial operation, depending on the user’s chosen build mode and optimization level.

In the case of the first alternative above, the potential loss in runtime performance from not doing the `memmove` invocation could be substantial, perhaps even dramatic, and would introduce a strong disincentive to adding such defensive checks to otherwise trivial functions, thus violating Principle 3. Hence, we consider the first alternative to be nonviable.

The second alternative is much more consistent with our conclusion with respect to the `noexcept` specifier in that it keeps the two language features — namely triviality and contract checking — maximally orthogonal. This latter approach does come with the risk of perhaps omitting checks that an uninitiated author might have intended the client to run in every case — i.e., including even those cases in which use of an equivalent `memmove` would be substantially more runtime performant. Fortunately, the duly informed author of a CCA for a special member function can always easily opt into this other slower but safer behavior:

```
~Date() [[ pre: isValid(d_year, d_month, d_day)]] { } // empty body
// Destroy this object.
```

---

<sup>10</sup>See [lakos21], Section 2.1.”Generalized Pods,” “Use Cases,” “Skippable destructors (notionally trivially destructible),” pp. 464–470, especially pp. 469–470.

<sup>11</sup>Although a special member function that has an empty body supplied by the user will provide no code and therefore would seem to be no different than the empty body supplied by the compiler, the copy constructor being trivial gives the compiler special permission to bypass calling that copy constructor entirely. This optimization is particularly effective for contiguous sequences of such objects — e.g., `std::vector<my_trivially_copyable_type>` — since repeated calls to the copy constructor can be replaced by a single call to `memmove`. For an object to be considered trivially copyable, however, it must have a trivial destructor. Note that the compiler’s ability to see that the body of a user-supplied destructor is empty doesn’t make that destructor trivial, nor does such compiler ability give license to the *library* to use `memmove` for a type that would otherwise be *trivially copyable* but for its almost trivial destructor.

Simply by providing an empty function body (see the example above), a function that was otherwise *trivial* can easily be made nontrivial in every build mode, thus removing the permission for libraries or the compiler to skip the invocation of this destructor and the evaluation of its associated checks. Consider that the checks are on the destructor, so there's nothing to skip on the copy constructor, and the compiler could easily run the checks on destruction in a checked build even though no code had to run to destroy the object. If the author of the class fails to realize the triviality of the function and, as a result, some check isn't run, no affirmative harm is done since the check was defensive (redundant) anyway and therefore entirely useless in every correctly written program.

Finally, we note that *all* contract checks are presumed to be purely defensive and thus entirely redundant in a defect-free program. Turning off runtime contract checking on otherwise *trivial* functions is, in effect, just one more way for developers to control if and to what extent their programs are checked at run time. Hence, as long as the presence of a CCA doesn't affect the compile-time result of evaluating a trait on a type, we might imagine giving the compiler explicit (e.g., via a compiler switch) freedom to sometimes refrain from invoking such runtime checks on trivial types, perhaps in collaboration with the totality of the (e.g., optimization) build modes.

Putting these suggestions together, we arrive at one proposal regarding the triviality of special member functions with contract-checking annotations.

#### Proposal 1: CCAs Do Not Affect Triviality

A (defaulted) special member function having preconditions or postconditions may still be trivial. Note that trivial special member functions might be replaced (by the language or a library) by bitwise copies or even elided completely, both of which would skip evaluation of any CCAs associated with that function.

### 3.2 Implicit Lambda Captures

As with any other function that may be defined in C++, functions defined using a lambda expression must be annotatable with precondition and postcondition CCAs and must contain assertion CCAs within their bodies. Any issues related to name resolution or appurtenance, such as those raised for trailing return types by [P2036R3], must be resolved by the syntax adopted for Contracts, such as is done in [P2935R0].

#### Proposal 2.1: CCAs Allowed on Lambdas

Precondition and postcondition CCAs may be placed on a lambda expression such that they appertain to the closure object's call operator. Assertion CCAs may appear within the body of a lambda expression.

Similar to affecting the triviality of a special member function, another area in which a CCA might be capable of having an impact on program behavior occurs when the predicate of a CCA within a lambda ODR-uses a local entity. In such cases, that entity might be implicitly captured as part of the generated closure object, thus affecting its size and possibly incurring a large cost to initialize it when that capture is by-value.

In some sense, *not* capturing at all when a CCA will be *ignored* and capturing only when a CCA would have other semantics might be possible. With the adoption of [P2877R0], no mechanism is available to have such properties as what is captured be based on the semantic that a CCA will have when evaluated. Additionally, having such a fundamental property of the closure object as the list of members it contains (and thus its size) be dependent on the chosen semantic of a CCA would be a violation of Principle 2 and thus should not be considered.

While simply allowing a capture might seem like a minimal violation of the zero-overhead principle (Principle 3), consider that the object referenced might far exceed the cost and scope of the lambda itself:

```
std::function<int()> foo(const std::vector<S>& v)
{
    int ndx = pickIndexAtRandom(v);
    return [=]()
        [[ pre : 0 <= ndx && ndx < v.size() ]] // needs to capture v
        {
            return ndx; // Obviously we intend this to capture ndx by value.
        };
}
```

Because the full `v` object must be captured by-value due to referencing it to get its size in the precondition of the lambda, this simple constant-time function becomes linear in the size of `v` due to the capture. That is a subtle and significant performance hit due to the presence of a CCA.

We are left with two options.

1. Allow the implicit capture from a CCA expression, which will happen regardless of the semantic with which the CCA is evaluated.
2. Make an implicit capture due to a CCA's expression ill-formed, preventing subtle costs due to the presence of a CCA.

Choosing the first alternative would violate Principle 3, ensuring that no runtime or object-size overhead is associated with any contract checks compared to simply not having such checks present in the program in the first place. The potentially significant and hidden cost of such captures might discourage the adoption of the Contracts facility in general.

The second alternative — i.e., making ill-formed the ODR-use of a local entity that is not already ODR-used in the body of the lambda apart from any CCA — creates zero overhead when *ignored*, has no effect on object size, and yet clearly informs users when they are referencing a value that is not directly relevant to the body of the lambda. This approach is both consistent with Principle 3 and, if there is any doubt about whether it is the correct long-term direction, is also an application of Principle 4.

The workaround, for those who want to capture that dubious value anyway, is simply to explicitly capture the variable in question or add the appropriate ODR-use directly within the body of the lambda:

```
std::function<int()> foo(const std::vector<S>& v)
{
    int ndx = pickIndexAtRandom(v);
```

```

return [=,v]() // capture v by value
  [[ pre : 0 <= ndx && ndx < v.size() ]] // needs to capture v
{
  return ndx; // Obviously we intend this to capture ndx by value.
};
}

```

Of course, when thinking about the compilation error due to the lack of capture of `v`, a developer will quickly realize that much better alternatives are available: `assert` that `ndx` is in the proper range prior to initializing the lambda or capture only `v.size()` for use in the lambda's precondition.

This reasoning all comes together into one additional proposal regarding lambdas.

#### Proposal 2.2: CCA ODR-Use Does Not Implicitly Capture

Within the expression of a CCA that is attached to or within a lambda, the ODR-use of a local entity does *not* implicitly capture that entity.

The reference to the local entity from the CCA will still continue to attempt to name the member of the closure object instead of the local entity itself. Therefore, if no other factor creates that member of the closure (either an explicit capture or an ODR-use that is *not* in a CCA), a program will be ill-formed.

In comparison to the proposal in [P2890R0], we have complete agreement on the need to support CCAs on lambda functions, and the only difference is whether the exception for captures of local entities from within CCAs should be made.

### 3.3 Compile-Time Evaluation

The specifics of evaluating a CCA at compile time have not been thoroughly pinned down in the MVP and, more importantly, might violate some of our fundamental principles if no additional changes are made.

Currently, after the adoption of [P2877R0], any evaluation of a CCA might or might not evaluate the predicate and detect a violation, and nothing about that proposal was specific to runtime evaluations. Therefore, even during compile-time evaluations, whether a CCA would be checked (i.e., have the *observe* or *enforce* semantic) or not (i.e., have the *ignore* semantic) is implementation defined.

More importantly, the contract-violation handling process involves invoking the global, replaceable contract-violation handler, which clearly can't be known when compiling an individual translation unit (TU).

Egregious differences between the behavior of compile-time evaluation and runtime evaluation are ill advised, so we propose retaining the meanings of the potential semantics of CCA evaluation at compile time and simply changing the effects of violations in a way that has similar spirit but does not allow the same customizability; we change an attempt to invoke the contract-violation handler into one that simply emits a diagnostic, the mechanism we use for the compiler to emit information about a problem to the user.

When a CCA being enforced is violated, we must also decide what happens when the *implementation-defined program termination* occurs, and again we have two potential choices.

1. Make the enclosing constant-evaluated expression ineligible to be a constant expression. This option is the typical specification tool used to say something *cannot* be done at compile time.
2. Make the program ill-formed, with no ability to choose a separate control flow path or to evaluate at run time instead.

With Option 1, one can then decide, based on the evaluation semantic, what to do in other contexts, altering program semantics based on the chosen semantic; doing so would violate Principle 2, and thus we consider this option nonviable. We are thus obliged to choose Option 2, making any constant evaluation that violates a checked CCA ill-formed.

#### Proposal: Compile-Time Contract Violations Emit Diagnostics

When the contract-violation process is evaluated at compile time (i.e., a CCA is violated when being evaluated with the *observe* or *enforce* semantic), a diagnostic will be emitted. When the evaluation is performed with the *enforce* semantic, the program is ill-formed.

The other major aspect of CCA evaluation that we must consider is the evaluation of the predicate itself and whether that must be eligible for compile-time evaluation. More importantly, can that predicate being ineligible for constant evaluation render the enclosing expression ineligible for being a constant expression?

Recall from [P2877R0] that we can effectively consider the evaluation of a CCA with expression  $X$  to be of this form:

```
switch (__current_contract_semantic()) {
case semantic::ignore:
    break;
case semantic::observe:
    if (X) {} else {
        __invoke_violation_handler(contract_info, semantic::observe);
    }
    break;
case semantic::enforce:
    if (X) {} else {
        __invoke_violation_handler(contract_info, semantic::enforce);
        __terminate_on_enforced_violation();
    }
    break;
}
```

The above proposal is equivalent to saying that the intrinsic `__current_contract_semantic()` is a valid core constant expression, and that `__current_contract_semantic` is `constexpr` or `consteval`. As described in [P2877R0], implementations have a wide range of flexibility in what this function may do.

- An implementation that provides a single, global switch to choose the semantics for all CCAs might, for example, when that chosen semantic is enforced, install a version of `semantic`

computation:

```
consteval semantic __current_contract_semantic()
{
    return semantic::enforce;
}
```

- A different configuration might choose to provide a mode in which the semantic computation function that is installed makes different decisions at compile or run time:

```
constexpr semantic __current_contract_semantic()
{
    if consteval {
        // constant-evaluation semantic:
        return semantic::enforce; // or something else, based on compiler flags
    }
    else {
        // runtime semantic
        return semantic::ignore; // or something else, based on compiler flags
    }
}
```

- Other implementations might inspect link-time properties to determine the runtime semantic, provide mechanisms to select the compile-time or runtime semantic based on the location of the CCA being evaluated, or do any number of other operations that produce the behavior such implementations have defined for the selection of CCA semantics.

When determining if an expression containing a CCA can be evaluated at compile time, we will need to identify two things:

1. Is each individual evaluation within the process of evaluating a CCA eligible to be part of a constant expression?
2. If the expression is not eligible, what happens?

The first question we have already answered above for the violation handler and termination: Neither can be performed at compile time, and it is ill-formed if the termination on enforced violation happens at compile time. The full expansion above, however, shows us we must answer questions about additional parts of this expansion as well.

First, is the computation of the semantic, which is the implementation-defined manner in which a semantic is chosen for the evaluation of a CCA, eligible to be done as part of a constant expression? If we were to allow the answer to this question to be *no*, then doing contract checking at compile time would be impossible; the very act of determining the semantic for a CCA would make the enclosing expression no longer a constant expression. This decision would be highly unfortunate, effectively making the use of CCAs completely incompatible with compile-time programming. Therefore, we propose that the selection can be done at compile time.

Proposal: Semantic May Be Selected At Compile Time

The implementation-defined choice of semantic when evaluating a CCA may be evaluated as part of a core constant expression; i.e., selecting the semantic as part of evaluating a CCA does not make the expression containing the CCA ineligible to be a *core constant expression*.

The other question is what happens when the expression,  $x$ , is ineligible to be evaluated at compile time. Here we note that one chosen semantic — the *ignore* semantic — can always make the answer to this question irrelevant. When evaluating a CCA and given our previous proposal, if the *ignore* semantic is chosen, the rest of the evaluation will always be eligible to be part of a constant expression; no other evaluations will be performed. Therefore, because we are striving to maintain Principle 2 (program semantics are independent of chosen CCA semantics), the evaluation of the expression must be just as eligible to be part of a constant expression as the empty branch in the *ignore* case is. Therefore, if the evaluation of the predicate,  $x$ , is *not* eligible to be part of a constant expression, we must make that ill-formed.

Just as we do with contract violations, we should allow *observing* a CCA that fails to be evaluable at compile time to emit a diagnostic while *enforcing* such a CCA should be ill-formed.

Proposal: Use of Nonconstant-Eligible Predicates When Constant Evaluation Is Required Is Ill-Formed

In an expression that is a valid core constant expression, evaluation of a predicate that is not eligible to be a core constant expression emits a diagnostic. If the semantic of the corresponding CCA is *enforce*, the program is ill-formed.

Note the importance of this situation being ill-formed when *enforced*, rather than the expression simply being ineligible to be a constant expression; some manifestly constant-evaluated contexts are also SFINAE contexts and thus would allow control flow to alter based on which semantic is selected for evaluating a CCA.

Consider the following prototypical example of a function with a narrow contract, also a function one might consider beneficial to use at compile time:

```
constexpr double sqrt(double x) [[ pre : x >= 0 ]];
```

Now, interestingly, we can make a concept check that could tell us whether a given expression is a valid Boolean constant expression:

```
#include <type_traits> // for std::bool_constant

template <double x>
concept can_constexpr_sqrt = requires { std::bool_constant<sqrt(x)>(); };
```

Now, code could be written in terms of this concept that would have highly unintuitive behavior since it depends entirely on the chosen semantic of the precondition check on `sqrt` at compile time:

```
template <double x>
void f()
```

```

{
  if constexpr (can_constexpr_sqrt<x>) {
    // #1 --- processing if x >= 0 or preconditions are disabled
  }
  else {
    // #2
  }
}

```

The code block at #2 is reachable only if we were to allow the precondition check failure to be subject to SFINAE and thus let the concept check fail instead of making the program ill-formed when the attempt to form the result `sqrt(-1.0)` was made. With our proposals, the above code would be ill-formed if the precondition’s predicate is evaluated; otherwise, the code would always flow through to the block at #1. This behavior is the same as what we would observe in evaluating this code at run time with no attempt to use it at compile time.

The same reasoning must be applied to potentially constant variables that are not `constexpr` — i.e., reference or non-volatile `const`-qualified variables of integral or enumeration types. These variables are usable in a constant expression if their initializers are core constant expressions, but they are also fine to initialize with noncore constant expressions (and thus dynamically initialize at run time).

Here, following Principle 2 is again important as is keeping any well-formed behavior consistent with the behavior when CCAs are *ignored*. To determine this, an evaluation with all CCAs ignored must first be done to ascertain if the expression, without CCAs, is eligible to be a core constant expression. When this trial evaluation is *not* eligible for the initialization of a potentially constant variable, the CCAs should not matter; the initialization will happen as a dynamic initialization at run time, and any contract violations that might be detected will happen then. A trial evaluation that determines that a CCA is a core constant expression will cause the variable initialization to be a manifestly constant-evaluated context, and thus the rules above will apply.

For such variables, we want to make ill-formed (and not SFINAE-able) any attempt to use them in a manifestly constant-evaluated context.

**Proposal: Nonconstant Eligible Initializers Evaluate CCAs at Run Time**

If the initializer of a `nonconstexpr` potentially-constant variable is not itself a core constant expression when all its CCAs are evaluated with the *ignore* semantic, then that initializer is not a core constant expression; hence, no attempt will be made to perform constant evaluation on the expression again with CCAs potentially having different semantics.

We call out this last proposal because we should not aggressively enforce contracts at compile time that might never be evaluated at run time.

All these proposals related to constant evaluation address individual aspects of the constant evaluation process and how it relates to CCA evaluation. Putting this all together produces one complete proposal that accomplishes all the above goals:

### Proposal 3: Evaluation of CCAs at Compile Time

When determining whether an expression is a core constant expression, first determine if the expression is a core constant expression by evaluating it with all CCAs having the *ignore* semantic.

- If it is a core constant expression or if it is not a core constant expression but is in a manifestly constant-evaluated context, re-evaluate the expression while evaluating the CCAs with semantics chosen in an implementation-defined manner.
  - If evaluating the resulting expression violates an *observed* CCA, emit a diagnostic. Constant evaluation continues as normal after the CCA.
  - If evaluating the result expression violates an *enforced* CCA, the program is ill-formed. Emit a diagnostic.
  - If the expression is not a core constant expression, the program is ill-formed.
- Otherwise, the expression is not a core constant expression.

In comparison to the proposal in [P2894R0], this proposal clarifies that contract violations (and contract predicates that are not evaluable at compile time) during constant evaluation must be hard errors, and the semantics for that proposed here is the result of extensive reflector discussion on the topic. Our understanding is that future revisions of [P2894R0] will propose very similar semantics. The primary difference is the support proposed here to enable a distinction between *observed* and *enforced* CCAs at compile time, where an *observed* CCA produces a diagnostic but is *not* a hard error.

### 3.4 Virtual Functions

With the adoption of [P2954R0] by SG21, CCAs on virtual functions have the following behavior.

- CCAs may be placed on virtual functions that do not override any other virtual functions.
- A virtual function that overrides exactly one with CCAs inherits the CCAs of that overridden function.
- Virtual functions that override multiple functions may do so only if none of those functions have CCAs.

The problem is that inheritance of CCAs from a base class violates Principle 1. By introducing a CCA on a base-class virtual function, all derived class virtual functions implicitly get that same virtual function. The general case for well-defined object-oriented implementations might be okay with this, but in some cases that will misidentify defects in software that are not actually wrong.

In particular, consider a derived class that has a wider contract on a function than its base class:

```
class Car {
    virtual void drive(int speed);
    // The behavior is undefined unless the specified speed
    // is less than 100.
};

class FastCar {
    void drive(int speed) override;
```

```

    // The behavior is undefined unless the specified speed
    // is less than 169.
};

```

This `FastCar` is a completely fine derived class for `Car` and can be used anywhere `Car` can.

Now consider adding a CCA to `Car` to validate the precondition of `drive`:

```

class Car {
    virtual void drive(int speed) [[ pre : speed < 100 ]];
};

```

The default inheritance of this CCA in `FastCar` will cause all code that drives a `FastCar` *fast* to break — even code that knowingly has a `FastCar` and reasonably wants to take advantage of the fact that it has a `FastCar` in hand.

This results in the following (partial) proposal to fix the current SG21 MVP behavior to adhere to our stated principles.

#### Proposal: CCAs Are Not Inherited

A virtual member function that has no precondition or postcondition CCAs on its declaration will have no precondition or postcondition CCAs (and will not inherit those of any functions it may be overriding).

Properly allowing variance of CCAs across class hierarchies requires checking both the CCAs of the virtual function being invoked as well as those of the concrete implementation found through virtual dispatch.<sup>12</sup> This solution might, however, not be ready for SG21 to adopt at this time.

Consider a test function where we use a virtual function via dynamic dispatch through both base-class and derived-class references:

```

void testCars()
{
    FastCar fc;

    Car& cr = fc;
    cr.drive(120); // CCAs of Car::drive and FastCar::drive
                  // should be checked.

    FastCar& fcr = fc;
    fcr.drive(120); // CCAs of FastCar::drive should be checked.
}

```

With the current MVP, the CCAs of `Car::drive` and `FastCar::drive` are the same, so the behavior will be compatible with the above. The current MVP behavior without inherited CCAs does not give a clear answer as to what CCAs would be checked in the above calls.

<sup>12</sup>After the initial MVP achieves consensus, the forthcoming [P2755R0] — or possibly other papers — will propose that the best approach here is to check both the CCAs that are visible based on the static type through which the function is being invoked as well as those of the specific concrete function selected by virtual dispatch, thus guaranteeing that all expectations of both caller and callee are satisfied.

Without a more complete solution, we must then apply Principle 4 and remove the ability to put CCAs on virtual function declarations until that more complete solution is adopted:

#### Proposal 4: No CCAs on Virtual Functions

It is ill-formed to place a `[[pre]]` or `[[post]]` CCA on a virtual function, i.e., a member function marked `virtual` or that overrides a virtual function in a base class.

### 3.5 Coroutines

When users begin to consider what `[[pre]]` and `[[post]]` might mean when applied to a coroutine, a few frequent misunderstandings and open questions would need to be resolved.

- Do preconditions get applied if the body of the coroutine does not begin to execute for a coroutine that starts suspended?
- Are the function parameters named in a precondition referring to the parameters of the function invocation or the copies made within the coroutine frame? What about parameters named by a postcondition?
- Is a return value named by a postcondition a reference to the function's returned value, to a value returned by `co_return`, or to something returned by `co_yield`?

While perhaps clear answers could be provided, they all leave unsatisfied other obvious needs for providing contract checks on coroutines: `pre` and `post` alone do not provide anything like a complete solution for the contract checks on the interface of a coroutine.

Some considerations must also be addressed.

- For coroutine handle types that perhaps are not awaitable and cannot themselves be used with `co_await`, such as `std::generator`, whether a given function is a coroutine is unclear from a client perspective. In this context, `pre` and `post` have clear meanings in terms of the production of that initial returned object. Callers must be able to treat `pre` and `post` in the same fashion independently of whether the function being invoked is a coroutine or not.
- For the implementer of a coroutine, however, many more entry and exit points are generally available between a coroutine and other code; each call to `co_return`, `co_yield`, and even `co_await` is a distinct part of the Promise-type-specific interface the coroutine has with the outside world. Any one of these boundaries might have requirements for correctness that would be beneficial for a user of the coroutine to know and thus could have meaning when placed on the coroutine's declaration.
- When invoking a coroutine, normally a return object is produced and returned to the caller, and little ambiguity arises about when a postcondition would get evaluated for that object. Invoking a coroutine within a `co_await` expression, however, is a fairly involved process, which may involve evaluation before and after suspension and resuming of the call-side coroutine, and when exactly postconditions should be evaluate or which specific options are even implementable becomes unclear.

Should a `co_await` expression result in suspension, arbitrary amounts of other code might

execute prior to resumption, any of which may invalidate a postcondition of the coroutine. Evaluation such as this would break the fundamental intuition many have that a postcondition hold in the code immediately following invocation of a function that had that postcondition.

- Caller-side CCA checks would naturally be able to apply only to the actual function arguments, not the copies within the coroutine frame. This restriction can lead to subtle bugs if properties are checked that do not propagate through such copies.

Possibly worse, forcing a coroutine parameter to be `const` by referring to it from a postcondition not only prevents modification of that parameter within the coroutine body (which is usually a manageable cost), but also causes the initialization of the copy (normally done with an `xvalue` referring to the original function parameter) to be done with a potentially much more expensive *copy* operation, not a *move* operation.

In general, having preconditions and postconditions refer to objects that the body of the function can never actually see is a pitfall that must be considered very carefully.

Currently, no concrete proposal covers the full breadth of the interface a coroutine has with its callers. Without this complete picture, we cannot yet know if `pre` and `post` will have a meaning that is correct and useful to those calling into or implementing coroutines. Possibly `pre` and `post` with their semantics applied caller-side are the optimal solution for coroutines, even with the potential risks. Additionally, coroutines might benefit more from having bespoke *kinds* of CCAs that have different, coroutine-specific semantics (which, due to callers generally being unaware if a function is coroutine, would need to be fully implemented in the coroutine definition) that avoid these pitfalls.

Therefore, we apply Principle 4 to decide that we should disallow *any* CCAs on a coroutine until we have a more complete picture of what we intend to provide.

#### Proposal 5: No CCAs on Coroutines

Specifying a precondition or postcondition on a function that is a coroutine is ill-formed. That is, a function that has a `co_return`, `co_yield`, or `co_await` expression in its definition and also has a `[[pre]]` or `[[post]]` CCA is ill-formed.

Note that Proposal 5 intentionally says nothing about assertions. Contract checks to identify defects can still be added to a coroutine (or within the function so defined by the `Promise` type); they must simply be manually put within the body and cannot have additional control points by which they can be automatically injected.

In comparison to [P2957R0], this paper agrees completely on the treatment of assertion CCAs within a coroutine body. That paper, however, proposes semantics for precondition and postcondition CCAs on coroutines, which we propose disallowing until a more complete and intuitive picture can be provided for the interactions between coroutines and Contracts.

### 3.6 Contracts on First Declarations

SG21 has agreed that, for the MVP, CCAs shall not be repeated on multiple declarations and must always be on the first declaration encountered for a function. This placement reduces the need

for answering questions about behavior if, for instance, a function is invoked in one place without having seen the CCAs on that function yet and in another place after CCAs have been seen.

Note the great benefit to insulating CCAs from either readers of a function's primary declaration or from client translation units; i.e., long-term, allowing CCAs to be repeated on later declarations, omitted from the first declaration and repeated at a later point, or even partially declared on one declaration and then fully defined on a subsequent declaration will be necessary. None of these possibilities would circumvent the one approach being allowed by the MVP: to support CCAs on first declarations only. Therefore, SG21 has followed Principle 4 to leave flexibility for future evolution and must simply define the behavior of the minimal feature we are supporting.

To fully define what we mean by restricting CCAs to be on first declarations only, we need to break down a few points.

- What is a *first* declaration?
- If multiple *first* declarations can appear in the same program, when are such CCAs the same?
- What happens if CCAs are *not* the same?

To allow for functions to be declared in multiple modules or in different header files, modern C++ would determine *first* based on reachability. In general, one declaration (A) is reachable from another (B) if A precedes B within the same TU or if A is declared in a module imported (directly or transitively) by the TU containing B. A declaration D for an entity is a *first* declaration of that entity if no other declarations for D are reachable. Interestingly, by declaring the same function (which must be attached to the global module fragment) in multiple different modules and importing those modules into a third one, multiple declarations can be *first* declarations of that function.

#### Definition: First Declarations

A declaration D for an entity is a *first declaration* of that entity if no other declarations for that entity are reachable from D — i.e., if no such declaration occurs earlier in the same TU and no such declaration is exported or used by a module that has been imported.

Given a clear definition of what is or is not a *first* declaration, we can restate the general rule in the MVP for CCAs being on first declarations only.

#### Proposal 6.1: CCAs on First Declarations Only

A function declaration that is not a first declaration shall not have `[[pre]]` or `[[post]]` CCAs attached to it.

To define if CCAs are the same, we will want to use the one-definition rule (ODR). This rule, however, applies to *definitions*, not declarations. Using part of the approach taken in C++20 Contracts,<sup>13</sup> we can define CCAs on different declarations for the same entity to be the same if they would be the same if placed in a function definition (at the same place where their declaration occurs), excluding potential renaming of function parameters, template parameters, or return value identifiers.

<sup>13</sup>See [N4820] for the draft Standard that included the C++20 Contracts proposal wording.

#### Definition: CCA Sameness

A CCA,  $c_1$ , on a function declaration,  $d_1$ , is the same as a CCA,  $c_2$ , on a function declaration,  $d_2$ , if their predicates,  $p_1$  and  $p_2$ , would satisfy the one-definition rule if placed in function definitions on the declarations  $d_1$  and  $d_2$ , respectively, except for the renaming of parameters, return value identifiers, and template parameters.

With the ability to identify if two CCAs are the same, we can easily extend this definition to two lists of CCAs being the same if all corresponding elements of the lists are the same.

Finally, the question of when to check that CCAs on a function are declared consistently across multiple first declarations must be addressed. Should the function declarations be in distinct TUs with no compile-time awareness of one another, we have little choice but to make a mismatch ill-formed, no diagnostic required (IFNDR).

Whether CCAs match when there are multiple *first* declarations reachable from the same point *could* be checked in two situations:

1. When importing a module that contains a *first* declaration for a function that is already reachable
2. When invoking a function with multiple *first* declarations that are reachable

Both options would require potentially large implementation difficulty, so we recommend leaving detection of mismatched CCAs as a quality of implementation decision; in practice, CCAs should mismatch only when declarations are being copied and pasted inappropriately instead of sourced from a header file with a single, canonical definition, so demanding extraordinary effort to detect such defects does not present as an incredibly high priority.

#### Proposal 6.2: Function CCA Lists Must Be Consistent

The list of CCAs on all first declarations (in all TUs) for a function shall be the same, no diagnostic required.

Because nonfirst declarations may have no CCAs on them, nothing else needs to be said for the Contracts MVP. If CCAs on nonfirst declarations were allowed, they would generally be from different source locations, and validating that the redeclared CCA lists were the same as the original would be straightforward and needed, and mismatches would then be ill-formed, *not* IFNDR.

## 4 Conclusion

SG21 has come a long way while finalizing a design for a Contracts MVP. Even with the small, focused scope of the MVP, open questions remain and must be addressed prior to having Contracts as an adopted feature in ISO C++. This paper has attempted to wrap up the remaining known edge cases that are essential to address to have a complete feature proposal ready to be integrated into the rich, powerful, and sometimes challenging language that C++ is. Should any further questions arise during the adoption of this proposal, we hope these same principles can guide remaining decisions to keep the Contracts facility's design clear and consistent.

As a foundation for reasoning about undecided features for the C++ Contracts facility, and the SG21 MVP in particular, we have presented four principles:

- Principle 1 — CCAs Check a Plain-Language Contract
- Principle 2 — Program Semantics Are Independent of Chosen CCA Semantics
- Principle 3 — Zero Overhead for Ignored Predicates
- Principle 4 — Make Undecided Behaviors Ill-Formed

Based on these principles, the open questions for the Contracts MVP have been addressed by the following proposals.

- Proposal 1 — CCAs Do Not Effect Triviality
- Proposal 2.1 — CCAs Allowed on Lambdas
- Proposal 2.2 — CCA ODR-Use Does Not Implicitly Capture
- Proposal 3 — Evaluation of CCAs at Compile Time
- Proposal 4 — No CCAs on Virtual Functions
- Proposal 5 — No CCAs on Coroutines
- Proposal 6.1 — CCAs on First Declarations Only
- Proposal 6.2 — Function CCA Lists Must Be Consistent

Finally, adoption of each of these reasoned proposals will serve two important objectives.

1. All outstanding design issues presented in [P2896R0] will be resolved.
2. The SG21 Contracts proposal will be a robust, initially useful, easy-to-evolve facility that integrates well with the whole of the C++ language and is ready to adopt and release in a timely fashion.

## 5 Acknowledgments

Thanks to Lori Hughes, Jens Maurer, John Lakos, and Andrzej Krzemiński for reviewing this paper and providing feedback.

## Bibliography

- [lakos21] John Lakos, Vittorio Romeo, Rostislav Khlebnikov, and Alisdair Meredith, *Embracing Modern C++ Safely* (Boston: Addison-Wesley, 2021)
- [N4820] Richard Smith, “Working Draft, Standard for Programming Language C++”, 2019 <http://wg21.link/N4820>
- [P0542R0] G. Dos Reis, J. D. Garcia, J. Lakos, A. Meredith, N. Myers, and B. Stroustrup, “Support for contract based programming in C++”, 2017 <http://wg21.link/P0542R0>

- [P1289R1] J. Daniel Garcia and Ville Voutilainen, “Access control in contract conditions”, 2018  
<http://wg21.link/P1289R1>
- [P1995R1] Joshua Berne, Andrzej Krzemiński, Ryan McDougall, Timur Doumler, and Herb Sutter, “Contracts — Use Cases”, 2020  
<http://wg21.link/P1995R1>
- [P2036R3] Barry Revzin, “Changing scope for lambda trailing-return-type”, 2021  
<http://wg21.link/P2036R3>
- [P2521R4] Andrzej Krzemiński, “Contract support – Record of SG21 consensus”, 2023  
<http://wg21.link/P2521R4>
- [P2695R1] Timur Doumler and John Spicer, “A proposed plan for contracts in C++”, 2023  
<http://wg21.link/P2695R1>
- [P2751R1] Joshua Berne, “Evaluation of *Checked* Contract-Checking Annotations”, 2023  
<http://wg21.link/P2751R1>
- [P2755R0] Joshua Berne, Jake Fevold, and John Lakos, “A Bold Plan for a Complete Contracts Facility (forthcoming)”, 2023
- [P2811R7] Joshua Berne, “Contract-Violation Handlers”, 2023  
<http://wg21.link/P2811R7>
- [P2834R1] Joshua Berne and John Lakos, “Semantic Stability Across Contract-Checking Build Modes”, 2023  
<http://wg21.link/P2834R1>
- [P2877R0] Joshua Berne and Tom Honermann, “Contract Build Modes, Semantics, and Implementation Strategies”, 2023  
<http://wg21.link/P2877R0>
- [P2890R0] Timur Doumler, “Contracts on lambdas”, 2023  
<http://wg21.link/P2890R0>
- [P2894R0] Timur Doumler, “Constant evaluation of Contracts”, 2023  
<http://wg21.link/P2894R0>
- [P2896R0] Timur Doumler, “Outstanding design questions for the Contracts MVP”, 2023  
<http://wg21.link/P2896R0>
- [P2900R0] Joshua Berne, Timur Doumler, and Andrzej Krzemiński, “Contracts for C++”, 2023
- [P2935R0] Joshua Berne, “An Attribute-Like Syntax for Contracts”, 2023  
<http://wg21.link/P2935R0>
- [P2954R0] Ville Voutilainen, “Contracts and virtual functions for the Contracts MVP”, 2023  
<http://wg21.link/P2954R0>
- [P2957R0] Andrzej Krzemiński and Iain Sandoe, “Contracts and coroutines”, 2023  
<http://wg21.link/P2957R0>