

**Document Number:** Nxxxx  
**Date:** 2022-01-11  
**Revises:** None  
**Reply to:** Michael Wong  
Codeplay  
fraggamuffin@gmail.com

# Working Draft, Extensions to C++ for Transactional Memory Version 2

**Note:** this is an early draft. It's known to be incomplet and incorrekt, and it has lots of bad fomattting.

# Contents

<b>Foreword</b>	<b>iii</b>
<b>1 Scope</b>	<b>1</b>
<b>2 Normative references</b>	<b>2</b>
<b>3 Terms and definitions</b>	<b>3</b>
<b>4 General</b>	<b>4</b>
4.1 Implementation compliance . . . . .	4
4.2 Namespaces and headers and modifications to standard classes . . . . .	4
4.3 Feature-testing recommendations (Informative) . . . . .	4
4.4 Future plans (Informative) . . . . .	4
4.5 Acknowledgments . . . . .	4
<b>5 Lexical conventions</b>	<b>6</b>
5.1 Identifiers . . . . .	6
<b>6 Basics</b>	<b>7</b>
<b>8 Statements</b>	<b>9</b>
8.1 Preamble . . . . .	9
8.8 Atomic statement . . . . .	9
<b>16 Library introduction</b>	<b>11</b>

# Foreword

# [foreword]

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see [www.iso.org/directives](http://www.iso.org/directives)).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see [www.iso.org/patents](http://www.iso.org/patents)) or the IEC list of patent declarations received (see <http://patents.iec.ch>).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see [www.iso.org/iso/foreword.html](http://www.iso.org/iso/foreword.html).

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at [www.iso.org/members.html](http://www.iso.org/members.html).

# 1 Scope

[scope]

- <sup>1</sup> This document describes requirements for implementations of an interface that computer programs written in the C++ programming language may use to invoke algorithms with concurrent execution. The algorithms described by this document are realizable across a broad class of computer architectures.
- <sup>2</sup> ISO/IEC 14882:2020 provide important context and specification for this document. This document is written as a set of changes against that specification. Instructions to modify or add paragraphs are written as explicit instructions. Modifications made directly to existing text from ISO/IEC 14882:2020 use underlining to represent added text and ~~strikethrough~~ to represent deleted text.
- <sup>3</sup> This document is non-normative. Some of the functionality described by this document may be considered for standardization in a future version of C++, but it is not currently part of any C++ standard. Some of the functionality in this document may never be standardized, and other functionality may be standardized in a substantially changed form.
- <sup>4</sup> The goal of this document is to build widespread existing practice for concurrency in the C++ standard algorithms library. It gives advice on extensions to those vendors who wish to provide them.

## 2 Normative references

[refs]

<sup>1</sup> The following referenced document is indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

(1.1) — ISO/IEC 14882:2020, *Programming Languages — C++*

<sup>2</sup> ISO/IEC 14882:2020 is herein called the C++ Standard. References to clauses within the C++ Standard are written as “C++20 §3.2”. The library described in C++20 §16-32 is herein called the C++ Standard Library.

### 3 Terms and definitions

**[defs]**

<sup>1</sup> No terms and definitions are listed in this document. ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- (1.1) — IEC Electropedia: available at <https://www.electropedia.org/>
- (1.2) — ISO Online browsing platform: available at <https://www.iso.org/obp>

## 4 General [general]

### 4.1 Implementation compliance [general.compliance]

- <sup>1</sup> Conformance requirements for this document are those defined in C++20 §4.1, as applied to a merged document consisting of C++20 amended by this document.

[*Note 1*: Conformance is defined in terms of the behavior of programs. — *end note*]

### 4.2 Namespaces and headers and modifications to standard classes [general.namespaces]

- <sup>1</sup> Since the extensions described in this technical specification are experimental and not part of the C++ standard library, they are not declared directly within namespace `std`. Unless otherwise specified, all components described in this technical specification either:

- (1.1) — modify an existing interface in the C++ Standard Library in-place,
- (1.2) — are declared in a namespace whose name appends `::experimental::transactional_memory_v2` to a namespace defined in the C++ Standard Library, such as `std`, or
- (1.3) — are declared in a subnamespace of a namespace described in the previous bullet, whose name is not the same as an existing subnamespace of namespace `std`.

- <sup>2</sup> Whenever an unqualified name is used in the specification of a declaration `D`, its meaning is established as-if by performing unqualified name lookup in the context of `D`.

[*Note 1*: Argument-dependent lookup is not performed. — *end note*]

Similarly, the meaning of a *qualified-id* is established as-if by performing qualified name lookup in the context of `D`.

[*Note 2*: Operators in expressions are not so constrained. — *end note*]

### 4.3 Feature-testing recommendations (Informative) [general.feature.test]

- <sup>1</sup> An implementation that provides support for this document should define each feature test macro defined in [Table 1](#) if no associated headers are indicated for that macro, and if associated headers are indicated for a macro, that macro is defined after inclusion of one of the corresponding headers specified in the table.

Table 1: Feature-test macros

Macro name	Value	Header
<code>__cpp_lib_transactional_memory_v2</code>	202110	<code>&lt;experimental/transactional_memory_v2&gt;</code>

### 4.4 Future plans (Informative) [general.plans]

- <sup>1</sup> This section describes tentative plans for future versions of this technical specification and plans for moving content into future versions of the C++ Standard.
- <sup>2</sup> The C++ committee intends to release a new version of this technical specification approximately every few years, containing the `transactional_memory` extensions we hope to add to a near-future version of the C++ Standard. Future versions will define their contents in `std::experimental::transactional_memory_v3`, `std::experimental::transactional_memory_v4`, etc., with the most recent implemented version inlined into `std::experimental`.
- <sup>3</sup> When an extension defined in this or a future version of this technical specification represents enough existing practice, it will be moved into the next version of the C++ Standard by removing the `experimental::transactional_memory_vN` segment of its namespace and by removing the `experimental/` prefix from its header's path.

### 4.5 Acknowledgments [general.ack]

This work is the result of a collaboration of researchers in industry and academia. We wish to thank the original authors of this document, Michael Wong, Hans Boehm, and Michael Spear. We also wish to thank

people who made valuable contributions within and outside these groups, including Jens Maurer, and many others not named here who contributed to the discussion.



## 5 Lexical conventions

**[lex]**

### 5.1 Identifiers

**[lex.name]**

<sup>1</sup> In C++20 §5.11, add [atomic](#) to the table of keywords (Table 5).

## 6 Basics

[basic]

### 6.9.1 Sequential execution

[intro.execution]

Change in C++20 §6.9.1 as indicated:

- 5 A *full-expression* is
- (5.1) — ...
  - (5.2) — an invocation of a destructor generated at the end of the lifetime of an object other than a temporary object (6.7.7) whose lifetime has not been extended, ~~or~~
  - (5.3) — [the start and the end of an atomic block \(8.8 \[stmt.tx\]\)](#), or
  - (5.4) — an expression that is not a subexpression of another expression and that is not otherwise part of a full-expression.

### 6.9.2 Multi-threaded executions and data races

[intro.multithread]

#### 6.9.2.1 Data races

[intro.races]

Change in C++20 §6.9.2.1 as indicated:

- 6 ~~Certain~~ Atomic blocks as well as certain library calls [may synchronize with](#) other [atomic blocks](#) and library calls performed by another thread.

Add a new paragraph after C++20 §6.9.2.1 paragraph 20:

- 21 An atomic block that is not dynamically nested within another atomic block is termed a *transaction*.  
 [Note 1: Due to syntactic constraints, blocks cannot overlap unless one is nested within the other. — *end note*]  
 There is a global total order of execution for all transactions. If, in that total order, a transaction T1 is ordered before a transaction T2, then
- (21.1) — no evaluation in T2 happens before any evaluation in T1 and
  - (21.2) — if T1 and T2 perform conflicting expression evaluations, then the end of T1 synchronizes with the start of T2.
- [Note 2: If the evaluations in T1 and T2 do not conflict, they might be executed concurrently. — *end note*]
- 22 Two actions are *potentially concurrent* if ...

Change in C++20 §6.9.2.1 paragraph 21:

- 21 ...
- [Note 3: It can be shown that programs that correctly use mutexes, [atomic blocks](#), and `memory_order::seq_cst` operations to prevent all data races and use no other synchronization operations behave as if the operations executed by their constituent threads were simply interleaved, with each value computation of an object being taken from the last side effect on that object in that interleaving. This is normally referred to as "sequential consistency". ... — *end note*]

Add a new paragraph after C++20 §6.9.2.1 paragraph 21:

- 22 [Note 4: The following holds for a data-race-free program: If the start of an atomic block T is sequenced before an evaluation A, A is sequenced before the end of T, A strongly happens before some evaluation B, and B is not sequenced before the end of T, then the end of T strongly happens before B. If an evaluation C strongly happens before that evaluation A and C is not sequenced after the start of T, then C strongly happens before the start of T. These properties in turn imply that in any simple interleaved (sequentially consistent) execution, the operations of each atomic block appear to be contiguous in the interleaving. — *end note*]

**6.9.2.2 Forward progress**

[intro.progress]

Change in C++20 §6.9.2.2 as indicated:

- 1 ~~The implementation may assume that any thread will eventually do~~ An *inter-thread side effect*  
is one of the following:
- (1.1) — ~~terminate,~~
  - (1.2) — a call to a library I/O function,
  - (1.3) — an access through a volatile glvalue, or
  - (1.4) — a synchronization operation or an atomic operation ([atomics]).

The implementation may assume that any thread will eventually terminate or evaluate an inter-thread side effect.

[*Note 1*: This is intended to allow compiler transformations such as removal of empty loops, even when termination cannot be proven. — *end note*]

## 8 Statements

[stmt.stmt]

### 8.1 Preamble

[stmt.pre]

- <sup>1</sup> Add a production to the grammar in C++20 §8.1 as indicated:

```

statement:
    labeled-statement
    attribute-specifier-seqopt expression-statement
    attribute-specifier-seqopt compound-statement
    attribute-specifier-seqopt selection-statement
    attribute-specifier-seqopt iteration-statement
    attribute-specifier-seqopt jump-statement
    declaration-statement
    attribute-specifier-seqopt try-block
    atomic-statement

```

Add a new subclause before C++20 §8.8:

### 8.8 Atomic statement

[stmt.tx]

```

atomic-statement:
    atomic do compound-statement

```

- <sup>1</sup> An *atomic-statement* is also called an *atomic block*.

- <sup>2</sup> The start of the atomic block is immediately before the opening { of the *compound-statement*. The end of the atomic block is immediately after the closing } of the *compound-statement*.

[*Note 1*: Thus, variables with automatic storage duration declared in the *compound-statement* are destroyed prior to reaching the end of the atomic block; see C++20 §8.7. — *end note*]

- <sup>3</sup> A goto or switch statement shall not be used to transfer control into an atomic block.

- <sup>4</sup> If the execution of an atomic block evaluates an inter-thread side effect (6.9.2.2) or if an atomic block is exited via an exception, the behavior is undefined.

- <sup>5</sup> *Recommended practice*: In case an atomic block is exited via an exception, the program should be terminated without invoking a terminate handler (C++20 §17.9.5) or destroying any objects with static or thread storage duration (C++20 §6.9.3.4).

- <sup>6</sup> If the execution of an atomic block evaluates any of the following outside of a manifestly constant-evaluated context (C++20 §7.7), the behavior is implementation-defined:

- (6.1) — an *asm-declaration* (C++20 §9.10 [dcl.asm]);
- (6.2) — an invocation of a function other than one of the standard library functions specified in C++20 §16.4.6.17 [atomic.use]), unless the function is inline with a reachable definition;
- (6.3) — a virtual function call (C++20 §7.6.1.3 [expr.call]);
- (6.4) — a function call, unless overload resolution selects
  - (6.4.1) — a named function (C++20 §12.2.2.2.2 [over.call.func]) or
  - (6.4.2) — a function call operator (C++20 §12.2.2.2.3 [over.call.object]), but not a surrogate call function;
- (6.5) — a `co_await` expression (C++20 §7.6.2.3 [expr.await]), a *yield-expression* (C++20 §7.6.17 [expr.yield]), or a `co_return` statement (C++20 §8.7.4 [stmt.return.coroutine]);
- (6.6) — dynamic initialization of a block-scope variable with static storage duration; or
- (6.7) — dynamic initialization of a variable with thread storage duration.

[*Note 2*: The implementation can define that the behavior is undefined in some or all of the cases above. — *end note*]

[*Example 1*:

```
unsigned int f()
{
    static unsigned int i = 0;
    atomic do {
        ++i;
        return i;
    }
}
```

Each invocation of `f` (even when called from several threads simultaneously) retrieves a unique value (ignoring wrap-around). — *end example*]

[*Note 3*: Atomic blocks are likely to perform best where they execute quickly and touch little data. — *end note*]

# 16 Library introduction

[library]

Add a new subclause after C++20 §16.4.6.16:

## 16.4.6.17 Functions usable in an atomic block

[atomic.use]

<sup>1</sup> All library functions may be used in an atomic block (8.8), except

- (1.1) — error category objects ([syserr.errcat.objects])
- (1.2) — time zone database ([time.zone.db])
- (1.3) — clocks ([time.clock])
- (1.4) — **signal** ([support.signal]) and **raise** ([csignal.syn])
- (1.5) — **set\_new\_handler**, **set\_terminate**, **get\_new\_handler**, **get\_terminate** ([handler.functions], [alloc.errors], [exception.syn])
- (1.6) — **system** ([cstdlib.syn])
- (1.7) — startup and termination [support.start.term] except **abort**
- (1.8) — **shared\_ptr** ([util.smartptr.shared]) and **weak\_ptr** ([util.smartptr.weak])
- (1.9) — **synchronized\_pool\_resource** ([mem.res.pool])
- (1.10) — program-wide **memory\_resource** objects ([mem.res.global])
- (1.11) — **setjmp** / **longjmp** ([csetjmp.syn])
- (1.12) — parallel algorithms ([algorithms.parallel])
- (1.13) — **random\_device** ([rand.device])
- (1.14) — **locale** construction ([locale.cons])
- (1.15) — input/output ([input.output])
- (1.16) — atomic operations ([atomics])
- (1.17) — thread support ([thread])