

---

---

**Programming Languages — Technical  
Specification for C++ Extensions for  
Parallelism**

*Langages de programmation — Spécification technique pour les  
extensions C++ relatives au parallélisme*

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TS 19570:2018



STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TS 19570:2018



**COPYRIGHT PROTECTED DOCUMENT**

© ISO/IEC 2018

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office  
CP 401 • Ch. de Blandonnet 8  
CH-1214 Vernier, Geneva  
Phone: +41 22 749 01 11  
Fax: +41 22 749 09 47  
Email: [copyright@iso.org](mailto:copyright@iso.org)  
Website: [www.iso.org](http://www.iso.org)

Published in Switzerland

# Contents

Foreword	iv
<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 Namespaces and headers . . . . .	4
4.2 Feature-testing recommendations . . . . .	4
<b>5 Parallel exceptions</b>	<b>5</b>
5.1 Header <code>&lt;experimental/exception_list&gt;</code> synopsis . . . . .	5
<b>6 Execution policies</b>	<b>6</b>
6.1 Header <code>&lt;experimental/execution&gt;</code> synopsis . . . . .	6
6.2 Unsequenced execution policy . . . . .	6
6.3 Vector execution policy . . . . .	6
6.4 Execution policy objects . . . . .	7
<b>7 Parallel algorithms</b>	<b>8</b>
7.1 Wavefront Application . . . . .	8
7.2 Non-Numeric Parallel Algorithms . . . . .	9
<b>8 Task Block</b>	<b>16</b>
8.1 Header <code>&lt;experimental/task_block&gt;</code> synopsis . . . . .	16
8.2 Class <code>task_cancelled_exception</code> . . . . .	16
8.3 Class <code>task_block</code> . . . . .	16
8.4 Function template <code>define_task_block</code> . . . . .	18
8.5 Exception Handling . . . . .	19
<b>9 Data-Parallel Types</b>	<b>20</b>
9.1 General . . . . .	20
9.2 Header <code>&lt;experimental/simd&gt;</code> synopsis . . . . .	20
9.3 <code>simd</code> ABI tags . . . . .	24
9.4 <code>simd</code> type traits . . . . .	26
9.5 Where expression class templates . . . . .	27
9.6 Class template <code>simd</code> . . . . .	30
9.7 <code>simd</code> non-member operations . . . . .	37
9.8 Class template <code>simd_mask</code> . . . . .	42
9.9 Non-member operations . . . . .	46

# Foreword

# [parallel.foreword]

- <sup>1</sup> 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. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.
- <sup>2</sup> 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)).
- <sup>3</sup> 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)).
- <sup>4</sup> Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.
- <sup>5</sup> For an explanation of 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).
- <sup>6</sup> This document was prepared by Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.
- <sup>7</sup> This second edition cancels and replaces the first edition (ISO/IEC 19570:2015) which has been technically revised.
- <sup>8</sup> The main changes compared to the previous edition are as follows:
  - (8.1) — Eliminate previously standardized functionality.
  - (8.2) — Introduce task block.
  - (8.3) — Introduce vector and wavefront policies.
  - (8.4) — Introduce a template library for parallel for loops.
  - (8.5) — Introduce data-parallel vector types.

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).

# Programming Languages — Technical Specification for C++ Extensions for Parallelism

## 1 Scope

[parallel.scope]

- <sup>1</sup> This document describes requirements for implementations of an interface that computer programs written in the C++ programming language can use to invoke algorithms with parallel execution. The algorithms described by this document are realizable across a broad class of computer architectures.
- <sup>2</sup> There is a possibility of a subset of the functionality described by this document being standardized in a future version of C++, but it is not currently part of any C++ standard. There is a possibility of some of the functionality in this document never being standardized, or of it being standardized in a substantially changed form.
- <sup>3</sup> The goal of this document is to build widespread existing practice for parallelism in the C++ programming language. It gives advice on extensions to those vendors who wish to provide them.

## 2 Normative references [parallel.references]

<sup>1</sup> The following documents are referred to in the text in such a way that some or all of their content constitutes requirements 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 14882:2017, *Programming languages — C++*

<sup>2</sup> ISO/IEC 14882:2017 is herein called the C++ Standard. References to clauses within the C++ Standard are written as “C++17 §20”. The library described in ISO/IEC 14882:2017 clauses 20-33 is herein called the *C++ Standard Library*. The C++ Standard Library components described in ISO/IEC 14882:2017 clauses 28, 29.8 and 23.10.10 are herein called the *C++ Standard Algorithms Library*.

<sup>3</sup> Unless otherwise specified, the whole of the C++ Standard’s Library introduction (C++17 §20) is included into this document by reference.

### 3 Terms and definitions [parallel.defns]

<sup>1</sup> For the purposes of this document, the terms, definitions, and symbols given in ISO/IEC 14882:2017 apply.

<sup>2</sup> ISO and IEC maintain terminological databases for use in standardization at the following addresses:

(2.1) — ISO Online browsing platform: available at <https://www.iso.org/obp>

(2.2) — IEC Electropedia: available at <http://www.electropedia.org>

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TS 19570:2018

## 4 General

## [parallel.general]

### 4.1 Namespaces and headers

[parallel.general.namespaces]

- <sup>1</sup> Since the extensions described in this document are experimental and not part of the C++ Standard Library, they should not be declared directly within namespace `std`. Unless otherwise specified, all components described in this document are declared in namespace `std::experimental::parallelism_v2`.

[*Note*: Once standardized, the components described by this document are expected to be promoted to namespace `std`. — *end note*]

- <sup>2</sup> Each header described in this document shall import the contents of `std::experimental::parallelism_v2` into `std::experimental` as if by

```
namespace std::experimental {
    inline namespace parallelism_v2 {}
}
```

- <sup>3</sup> Unless otherwise specified, references to such entities described in this document are assumed to be qualified with `std::experimental::parallelism_v2`, and references to entities described in the C++ Standard Library are assumed to be qualified with `std::`.

- <sup>4</sup> Extensions that are expected to eventually be added to an existing header `<meow>` are provided inside the `<experimental/meow>` header, which shall include the standard contents of `<meow>` as if by

```
#include <meow>
```

### 4.2 Feature-testing recommendations

[parallel.general.features]

- <sup>1</sup> An implementation that provides support for this document shall define the feature test macro(s) in Table 1.

Table 1 — Feature-test macro(s)

Title	Subclause	Macro Name	Value	Header
Task Block	8	<code>__cpp_lib_experimental_parallel_task_block</code>	201711	<code>&lt;experimental/exception_list&gt;</code> <code>&lt;experimental/task_block&gt;</code>
Vector and Wavefront Policies	6.2	<code>__cpp_lib_experimental_execution_vector_policy</code>	201711	<code>&lt;experimental/algorithm&gt;</code> <code>&lt;experimental/execution&gt;</code>
Template Library for Parallel For Loops	7.2.2, 7.2.3, 7.2.4	<code>__cpp_lib_experimental_parallel_for_loop</code>	201711	<code>&lt;experimental/algorithm&gt;</code>
Data-Parallel Vector Types	9	<code>__cpp_lib_experimental_parallel_simd</code>	201803	<code>&lt;experimental/simd&gt;</code>

## 5 Parallel exceptions [parallel.exceptions]

### 5.1 Header <experimental/exception\_list> synopsis [parallel.exceptions.synopsis]

```

namespace std::experimental {
inline namespace parallelism_v2 {

    class exception_list : public exception {
    public:
        using iterator = unspecified;

        size_t size() const noexcept;
        iterator begin() const noexcept;
        iterator end() const noexcept;

        const char* what() const noexcept override;
    };
}
}

```

1 The class `exception_list` owns a sequence of `exception_ptr` objects.

2 `exception_list::iterator` is an iterator which meets the forward iterator requirements and has a value type of `exception_ptr`.

```
size_t size() const noexcept;
```

3 *Returns:* The number of `exception_ptr` objects contained within the `exception_list`.

4 *Complexity:* Constant time.

```
iterator begin() const noexcept;
```

5 *Returns:* An iterator referring to the first `exception_ptr` object returned within the `exception_list`.

```
iterator end() const noexcept;
```

6 *Returns:* An iterator that is past the end of the owned sequence.

```
const char* what() const noexcept override;
```

7 *Returns:* An implementation-defined NTBS.

## 6 Execution policies [parallel.execpol]

### 6.1 Header `<experimental/execution>` synopsis [parallel.execpol.synopsis]

```
#include <execution>

namespace std::experimental {
  inline namespace parallelism_v2 {
    namespace execution {
      // 6.2, Unsequenced execution policy
      class unsequenced_policy;

      // 6.3, Vector execution policy
      class vector_policy;

      // 6.4, Execution policy objects
      inline constexpr unsequenced_policy unseq{ unspecified };
      inline constexpr vector_policy vec{ unspecified };
    }
  }
}
```

### 6.2 Unsequenced execution policy [parallel.execpol.unseq]

```
class unsequenced_policy { unspecified };
```

- <sup>1</sup> The class `unsequenced_policy` is an execution policy type used as a unique type to disambiguate parallel algorithm overloading and indicate that a parallel algorithm's execution may be vectorized, e.g., executed on a single thread using instructions that operate on multiple data items.
- <sup>2</sup> The invocations of element access functions in parallel algorithms invoked with an execution policy of type `unsequenced_policy` are permitted to execute in an unordered fashion in the calling thread, unsequenced with respect to one another within the calling thread. [*Note*: This means that multiple function object invocations may be interleaved on a single thread. — *end note*]
- <sup>3</sup> [*Note*: This overrides the usual guarantee from the C++ Standard, C++17 §4.6 that function executions do not overlap with one another. — *end note*]
- <sup>4</sup> During the execution of a parallel algorithm with the `experimental::execution::unsequenced_policy` policy, if the invocation of an element access function exits via an uncaught exception, `terminate()` will be called.

### 6.3 Vector execution policy [parallel.execpol.vec]

```
class vector_policy { unspecified };
```

- <sup>1</sup> The class `vector_policy` is an execution policy type used as a unique type to disambiguate parallel algorithm overloading and indicate that a parallel algorithm's execution may be vectorized. Additionally, such vectorization will result in an execution that respects the sequencing constraints of wavefront application (7.1). [*Note*: The implementation thus makes stronger guarantees than for `unsequenced_policy`, for example. — *end note*]
- <sup>2</sup> The invocations of element access functions in parallel algorithms invoked with an execution policy of type `vector_policy` are permitted to execute in unordered fashion in the calling thread, unsequenced with respect

to one another within the calling thread, subject to the sequencing constraints of wavefront application (7.1) for the last argument to `for_loop`, `for_loop_n`, `for_loop_strided`, or `for_loop_strided_n`.

- <sup>3</sup> During the execution of a parallel algorithm with the `experimental::execution::vector_policy` policy, if the invocation of an element access function exits via an uncaught exception, `terminate()` will be called.

#### 6.4 Execution policy objects

[parallel.execpol.objects]

```
inline constexpr execution::unsequenced_policy unseq { unspecified };
inline constexpr execution::vector_policy vec { unspecified };
```

- <sup>1</sup> The header `<experimental/execution>` declares a global object associated with each type of execution policy defined by this document.

## 7 Parallel algorithms

[parallel.alg]

### 7.1 Wavefront Application

[parallel.alg.wavefront]

<sup>1</sup> For the purposes of this subclause, an *evaluation* is a value computation or side effect of an expression, or an execution of a statement. Initialization of a temporary object is considered a subexpression of the expression that necessitates the temporary object.

<sup>2</sup> An evaluation *A* *contains* an evaluation *B* if:

- (2.1) — *A* and *B* are not potentially concurrent (C++17 §4.7.1); and
- (2.2) — the start of *A* is the start of *B* or the start of *A* is sequenced before the start of *B*; and
- (2.3) — the completion of *B* is the completion of *A* or the completion of *B* is sequenced before the completion of *A*.

[*Note:* This includes evaluations occurring in function invocations. — *end note*]

<sup>3</sup> An evaluation *A* is *ordered before* an evaluation *B* if *A* is deterministically sequenced before *B*. [*Note:* If *A* is indeterminately sequenced with respect to *B* or *A* and *B* are unsequenced, then *A* is not ordered before *B* and *B* is not ordered before *A*. The ordered before relationship is transitive. — *end note*]

<sup>4</sup> For an evaluation *A* ordered before an evaluation *B*, both contained in the same invocation of an element access function, *A* is a *vertical antecedent* of *B* if:

- (4.1) — there exists an evaluation *S* such that:
  - (4.1.1) — *S* contains *A*, and
  - (4.1.2) — *S* contains all evaluations *C* (if any) such that *A* is ordered before *C* and *C* is ordered before *B*,
  - (4.1.3) — but *S* does not contain *B*, and
- (4.2) — control reached *B* from *A* without executing any of the following:
  - (4.2.1) — a `goto` statement or `asm` declaration that jumps to a statement outside of *S*, or
  - (4.2.2) — a `switch` statement executed within *S* that transfers control into a substatement of a nested selection or iteration statement, or
  - (4.2.3) — a `throw` [*Note:* Even if caught — *end note*], or
  - (4.2.4) — a `longjmp`

[*Note:* Vertical antecedent is an irreflexive, antisymmetric, nontransitive relationship between two evaluations. Informally, *A* is a vertical antecedent of *B* if *A* is sequenced immediately before *B* or *A* is nested zero or more levels within a statement *S* that immediately precedes *B*. — *end note*]

<sup>5</sup> In the following,  $X_i$  and  $X_j$  refer to evaluations of the same expression or statement contained in the application of an element access function corresponding to the  $i^{\text{th}}$  and  $j^{\text{th}}$  elements of the input sequence. [*Note:* There can be several evaluations  $X_k, Y_k$ , etc. of a single expression or statement in application  $k$ , for example, if the expression or statement appears in a loop within the element access function. — *end note*]

<sup>6</sup> *Horizontally matched* is an equivalence relationship between two evaluations of the same expression. An evaluation  $B_i$  is *horizontally matched* with an evaluation  $B_j$  if:

- (6.1) — both are the first evaluations in their respective applications of the element access function, or

- (6.2) — there exist horizontally matched evaluations  $A_i$  and  $A_j$  that are vertical antecedents of evaluations  $B_i$  and  $B_j$ , respectively.

[*Note: Horizontally matched* establishes a theoretical lock-step relationship between evaluations in different applications of an element access function. — *end note*]

- <sup>7</sup> Let  $f$  be a function called for each argument list in a sequence of argument lists. *Wavefront application* of  $f$  requires that evaluation  $A_i$  be sequenced before evaluation  $B_j$  if  $i < j$  and:

- (7.1) —  $A_i$  is sequenced before some evaluation  $B_i$  and  $B_i$  is horizontally matched with  $B_j$ , or

- (7.2) —  $A_i$  is horizontally matched with some evaluation  $A_j$  and  $A_j$  is sequenced before  $B_j$ .

[*Note: Wavefront application* guarantees that parallel applications  $i$  and  $j$  execute such that progress on application  $j$  never gets ahead of application  $i$ . — *end note*] [*Note: The relationships between  $A_i$  and  $B_i$  and between  $A_j$  and  $B_j$  are sequenced before, not vertical antecedent.* — *end note*]

## 7.2 Non-Numeric Parallel Algorithms

[parallel.alg.ops]

### 7.2.1 Header <experimental/algorithm> synopsis

[parallel.alg.ops.synopsis]

```
#include <algorithm>

namespace std::experimental {
  inline namespace parallelism_v2 {
    namespace execution {
      // 7.2.5, No vec
      template<class F>
        auto no_vec(F&& f) noexcept -> decltype(std::forward<F>(f)());

      // 7.2.6, Ordered update class
      template<class T>
        class ordered_update_t;

      // 7.2.7, Ordered update function template
      template<class T>
        ordered_update_t<T> ordered_update(T& ref) noexcept;
    }

    // Exposition only: Suppress template argument deduction.
    template<class T> struct type_identity { using type = T; };
    template<class T> using type_identity_t = typename type_identity<T>::type;

    // 7.2.2, Reductions
    template<class T, class BinaryOperation>
      unspecified reduction(T& var, const T& identity, BinaryOperation combiner);
    template<class T>
      unspecified reduction_plus(T& var);
    template<class T>
      unspecified reduction_multiplies(T& var);
    template<class T>
      unspecified reduction_bit_and(T& var);
    template<class T>
      unspecified reduction_bit_or(T& var);
    template<class T>
      unspecified reduction_bit_xor(T& var);
    template<class T>
      unspecified reduction_min(T& var);
  }
}
```

```

template<class T>
    unspecified reduction_max(T& var);

// 7.2.3, Inductions
template<class T>
    unspecified induction(T&& var);
template<class T, class S>
    unspecified induction(T&& var, S stride);

// 7.2.4, For loop
template<class I, class... Rest>
    void for_loop(type_identity_t<I> start, I finish, Rest&&... rest);
template<class ExecutionPolicy,
         class I, class... Rest>
    void for_loop(ExecutionPolicy&& exec,
                 type_identity_t<I> start, I finish, Rest&&... rest);
template<class I, class S, class... Rest>
    void for_loop_strided(type_identity_t<I> start, I finish,
                        S stride, Rest&&... rest);
template<class ExecutionPolicy,
         class I, class S, class... Rest>
    void for_loop_strided(ExecutionPolicy&& exec,
                        type_identity_t<I> start, I finish,
                        S stride, Rest&&... rest);
template<class I, class Size, class... Rest>
    void for_loop_n(I start, Size n, Rest&&... rest);
template<class ExecutionPolicy,
         class I, class Size, class... Rest>
    void for_loop_n(ExecutionPolicy&& exec,
                  I start, Size n, Rest&&... rest);
template<class I, class Size, class S, class... Rest>
    void for_loop_n_strided(I start, Size n, S stride, Rest&&... rest);
template<class ExecutionPolicy,
         class I, class Size, class S, class... Rest>
    void for_loop_n_strided(ExecutionPolicy&& exec,
                          I start, Size n, S stride, Rest&&... rest);
}
}

```

## 7.2.2 Reductions

[parallel.alg.reductions]

- <sup>1</sup> Each of the function templates in this subclause (7.2.2) returns a *reduction object* of unspecified type having a *reduction value type* and encapsulating a *reduction identity* value for the reduction, a *combiner* function object, and a *live-out object* from which the initial value is obtained and into which the final value is stored.
- <sup>2</sup> An algorithm uses reduction objects by allocating an unspecified number of instances, known as *accumulators*, of the reduction value type. [Note: An implementation can, for example, allocate an accumulator for each thread in its private thread pool. — end note] Each accumulator is initialized with the object's reduction identity, except that the live-out object (which was initialized by the caller) comprises one of the accumulators. The algorithm passes a reference to an accumulator to each application of an element-access function, ensuring that no two concurrently executing invocations share the same accumulator. An accumulator can be shared between two applications that do not execute concurrently, but initialization is performed only once per accumulator.
- <sup>3</sup> Modifications to the accumulator by the application of element access functions accrue as partial results. At some point before the algorithm returns, the partial results are combined, two at a time, using the reduction

object's combiner operation until a single value remains, which is then assigned back to the live-out object. [Note: In order to produce useful results, modifications to the accumulator should be limited to commutative operations closely related to the combiner operation. For example if the combiner is `plus<T>`, incrementing the accumulator would be consistent with the combiner but doubling it or assigning to it would not. — end note]

```
template<class T, class BinaryOperation>
    unspecified reduction(T& var, const T& identity, BinaryOperation combiner);
```

4 *Requires:* T shall meet the requirements of `CopyConstructible` and `MoveAssignable`. The expression `var = combiner(var, var)` shall be well-formed.

5 *Returns:* A reduction object of unspecified type having reduction value type T, reduction identity `identity`, combiner function object `combiner`, and using the object referenced by `var` as its live-out object.

```
template<class T>
    unspecified reduction_plus(T& var);
template<class T>
    unspecified reduction_multiplies(T& var);
template<class T>
    unspecified reduction_bit_and(T& var);
template<class T>
    unspecified reduction_bit_or(T& var);
template<class T>
    unspecified reduction_bit_xor(T& var);
template<class T>
    unspecified reduction_min(T& var);
template<class T>
    unspecified reduction_max(T& var);
```

6 *Requires:* T shall meet the requirements of `CopyConstructible` and `MoveAssignable`.

7 *Returns:* A reduction object of unspecified type having reduction value type T, reduction identity and combiner operation as specified in Table 2 and using the object referenced by `var` as its live-out object.

Table 2 — Reduction identities and combiner operations

Function	Reduction Identity	Combiner Operation
<code>reduction_plus</code>	<code>T()</code>	<code>x + y</code>
<code>reduction_multiplies</code>	<code>T(1)</code>	<code>x * y</code>
<code>reduction_bit_and</code>	<code>(~T())</code>	<code>x &amp; y</code>
<code>reduction_bit_or</code>	<code>T()</code>	<code>x   y</code>
<code>reduction_bit_xor</code>	<code>T()</code>	<code>x ^ y</code>
<code>reduction_min</code>	<code>var</code>	<code>min(x, y)</code>
<code>reduction_max</code>	<code>var</code>	<code>max(x, y)</code>

[Example: The following code updates each element of `y` and sets `s` to the sum of the squares.

```
extern int n;
extern float x[], y[], a;
float s = 0;
for_loop(execution::vec, 0, n,
    reduction(s, 0.0f, plus<>()),
    [&](int i, float& accum) {
        y[i] += a*x[i];
```

```

    accum += y[i]*y[i];
  }
);

```

— end example]

### 7.2.3 Inductions

[parallel.alg.inductions]

- 1 Each of the function templates in this subclause return an *induction object* of unspecified type having an *induction value type* and encapsulating an initial value  $i$  of that type and, optionally, a *stride*.
- 2 For each element in the input range, an algorithm over input sequence  $S$  computes an *induction value* from an induction variable and ordinal position  $p$  within  $S$  by the formula  $i + p * stride$  if a stride was specified or  $i + p$  otherwise. This induction value is passed to the element access function.
- 3 An induction object may refer to a *live-out* object to hold the final value of the induction sequence. When the algorithm using the induction object completes, the live-out object is assigned the value  $i + n * stride$ , where  $n$  is the number of elements in the input range.

```

template<class T>
    unspecified induction(T&& var);
template<class T, class S>
    unspecified induction(T&& var, S stride);

```

- 4 *Returns:* An induction object with induction value type `remove_cv_t<remove_reference_t<T>>`, initial value `var`, and (if specified) stride `stride`. If  $T$  is an lvalue reference to non-const type, then the object referenced by `var` becomes the live-out object for the induction object; otherwise there is no live-out object.

### 7.2.4 For loop

[parallel.alg.forloop]

```

template<class I, class... Rest>
    void for_loop(type_identity_t<I> start, I finish, Rest&&... rest);
template<class ExecutionPolicy, class I, class... Rest>
    void for_loop(ExecutionPolicy&& exec, type_identity_t<I> start, I finish, Rest&&... rest);

template<class I, class S, class... Rest>
    void for_loop_strided(type_identity_t<I> start, I finish, S stride, Rest&&... rest);
template<class ExecutionPolicy, class I, class S, class... Rest>
    void for_loop_strided(ExecutionPolicy&& exec, type_identity_t<I> start, I finish, S stride,
        Rest&&... rest);

template<class I, class Size, class... Rest>
    void for_loop_n(I start, Size n, Rest&&... rest);
template<class ExecutionPolicy, class I, class Size, class... Rest>
    void for_loop_n(ExecutionPolicy&& exec, I start, Size n, Rest&&... rest);

template<class I, class Size, class S, class... Rest>
    void for_loop_n_strided(I start, Size n, S stride, Rest&&... rest);
template<class ExecutionPolicy, class I, class Size, class S, class... Rest>
    void for_loop_n_strided(ExecutionPolicy&& exec, I start, Size n, S stride, Rest&&... rest);

```

- 1 *Requires:* For the overloads with an `ExecutionPolicy`,  $I$  shall be an integral type or meet the requirements of a forward iterator type; otherwise,  $I$  shall be an integral type or meet the requirements of an input iterator type. `Size` shall be an integral type and  $n$  shall be non-negative.  $S$  shall have integral type and `stride` shall have non-zero value. `stride` shall be negative only if  $I$  has integral type or meets the requirements of a bidirectional iterator. The `rest` parameter pack shall have at least one

element, comprising objects returned by invocations of `reduction` (7.2.2) and/or `induction` (7.2.3) function templates followed by exactly one invocable element-access function, *f*. For the overloads with an `ExecutionPolicy`, *f* shall meet the requirements of `CopyConstructible`; otherwise, *f* shall meet the requirements of `MoveConstructible`.

2 *Effects*: Applies *f* to each element in the *input sequence*, as described below, with additional arguments corresponding to the reductions and inductions in the `rest` parameter pack. The length of the input sequence is:

- (2.1) — `n`, if specified,
- (2.2) — otherwise `finish - start` if neither `n` nor `stride` is specified,
- (2.3) — otherwise  $1 + (\text{finish} - \text{start} - 1) / \text{stride}$  if `stride` is positive,
- (2.4) — otherwise  $1 + (\text{start} - \text{finish} - 1) / -\text{stride}$ .

The first element in the input sequence is `start`. Each subsequent element is generated by adding `stride` to the previous element, if `stride` is specified, otherwise by incrementing the previous element. [*Note*: As described in the C++ Standard, C++17 §28.1, arithmetic on non-random-access iterators is performed using `advance` and `distance`. — *end note*] [*Note*: The order of the elements of the input sequence is important for determining ordinal position of an application of *f*, even though the applications themselves may be unordered. — *end note*]

The first argument to *f* is an element from the input sequence. [*Note*: If *I* is an iterator type, the iterators in the input sequence are not dereferenced before being passed to *f*. — *end note*] For each member of the `rest` parameter pack excluding *f*, an additional argument is passed to each application of *f* as follows:

- (2.5) — If the pack member is an object returned by a call to a reduction function listed in subclause 7.2.2, then the additional argument is a reference to an accumulator of that reduction object.
- (2.6) — If the pack member is an object returned by a call to `induction`, then the additional argument is the induction value for that induction object corresponding to the position of the application of *f* in the input sequence.

3 *Complexity*: Applies *f* exactly once for each element of the input sequence.

4 *Remarks*: If *f* returns a result, the result is ignored.

## 7.2.5 No vec

[`parallel.alg.novec`]

```
template<class F>
  auto no_vec(F&& f) noexcept -> decltype(std::forward<F>(f)());
```

1 *Effects*: Evaluates `std::forward<F>(f)()`. When invoked within an element access function in a parallel algorithm using `vector_policy`, if two calls to `no_vec` are horizontally matched within a wavefront application of an element access function over input sequence *S*, then the execution of `f` in the application for one element in *S* is sequenced before the execution of `f` in the application for a subsequent element in *S*; otherwise, there is no effect on sequencing.

2 *Returns*: The result of `f`.

3 *Notes*: If `f` exits via an exception, then `terminate` will be called, consistent with all other potentially-throwing operations invoked with `vector_policy` execution.

[*Example*:

```
extern float y[];
extern int* p;
for_loop(vec, 0, n, [&](int i) {
  y[i] += y[i+1];
  if (y[i] < 0) {
```

```

        no_vec([]{
            *p++ = i;
        });
    }
};

```

The updates `*p++ = i` will occur in the same order as if the policy were `seq.` — *end example*]

## 7.2.6 Ordered update class

[parallel.alg.ordupdate.class]

```

template<class T>
class ordered_update_t {
    T& ref_;           // exposition only
public:
    ordered_update_t(T& loc) noexcept
        : ref_(loc) {}
    ordered_update_t(const ordered_update_t&) = delete;
    ordered_update_t& operator=(const ordered_update_t&) = delete;

    template <class U>
        auto operator=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ = std::move(rhs); }); };
    template <class U>
        auto operator+=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ += std::move(rhs); }); };
    template <class U>
        auto operator-=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ -= std::move(rhs); }); };
    template <class U>
        auto operator*=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ *= std::move(rhs); }); };
    template <class U>
        auto operator/=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ /= std::move(rhs); }); };
    template <class U>
        auto operator%=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ %= std::move(rhs); }); };
    template <class U>
        auto operator>>=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ >>= std::move(rhs); }); };
    template <class U>
        auto operator<<=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ <<= std::move(rhs); }); };
    template <class U>
        auto operator&=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ &= std::move(rhs); }); };
    template <class U>
        auto operator^=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ ^= std::move(rhs); }); };
    template <class U>
        auto operator|=(U rhs) const noexcept
            { return no_vec([&]{ return ref_ |= std::move(rhs); }); };

    auto operator++() const noexcept
        { return no_vec([&]{ return ++ref_; }); };
    auto operator++(int) const noexcept

```

```

    { return no_vec([&]{ return ref_++; }); }
    auto operator--() const noexcept
    { return no_vec([&]{ return --ref_; }); }
    auto operator--(int) const noexcept
    { return no_vec([&]{ return ref_--; }); }
};

```

- <sup>1</sup> An object of type `ordered_update_t<T>` is a proxy for an object of type `T` intended to be used within a parallel application of an element access function using a policy object of type `vector_policy`. Simple increments, assignments, and compound assignments to the object are forwarded to the proxied object, but are sequenced as though executed within a `no_vec` invocation. [*Note*: The return-value deduction of the forwarded operations results in these operations returning by value, not reference. This formulation prevents accidental collisions on accesses to the return value. — *end note*]

### 7.2.7 Ordered update function template

[parallel.alg.ordupdate.func]

```

template<T>
    ordered_update_t<T> ordered_update(T& loc) noexcept;

```

- <sup>1</sup> *Returns*: { `loc` }.

## 8 Task Block

[parallel.taskblock]

### 8.1 Header <experimental/task\_block> synopsis [parallel.taskblock.synopsis]

```

namespace std::experimental {
inline namespace parallelism_v2 {
    class task_cancelled_exception;

    class task_block;

    template<class F>
        void define_task_block(F&& f);

    template<class f>
        void define_task_block_restore_thread(F&& f);
}
}

```

### 8.2 Class task\_cancelled\_exception [parallel.taskblock.task\_cancelled\_exception]

```

namespace std::experimental {
inline namespace parallelism_v2 {

    class task_cancelled_exception : public exception
    {
    public:
        task_cancelled_exception() noexcept;
        virtual const char* what() const noexcept override;
    };
}
}

```

- <sup>1</sup> The class `task_cancelled_exception` defines the type of objects thrown by `task_block::run` or `task_block::wait` if they detect that an exception is pending within the current parallel block. See 8.5, below.

```
virtual const char* what() const noexcept;
```

- <sup>2</sup> *Returns:* An implementation-defined NTBS.

### 8.3 Class task\_block [parallel.taskblock.class]

```

namespace std::experimental {
inline namespace parallelism_v2 {

    class task_block
    {
    private:
        ~task_block();

    public:
        task_block(const task_block&) = delete;
        task_block& operator=(const task_block&) = delete;
        void operator&() const = delete;
    };
}
}

```

```

    template<class F>
    void run(F&& f);

    void wait();
};
}
}

```

- 1 The class `task_block` defines an interface for forking and joining parallel tasks. The `define_task_block` and `define_task_block_restore_thread` function templates create an object of type `task_block` and pass a reference to that object to a user-provided function object.
- 2 An object of class `task_block` cannot be constructed, destroyed, copied, or moved except by the implementation of the task block library. Taking the address of a `task_block` object via `operator&` is ill-formed. Obtaining its address by any other means (including `addressof`) results in a pointer with an unspecified value; dereferencing such a pointer results in undefined behavior.
- 3 A `task_block` is *active* if it was created by the *nearest enclosing task block*, where *task block* refers to an invocation of `define_task_block` or `define_task_block_restore_thread` and *nearest enclosing* means the most recent invocation that has not yet completed. Code designated for execution in another thread by means other than the facilities in this subclause (e.g., using `thread` or `async`) are not enclosed in the task block and a `task_block` passed to (or captured by) such code is not active within that code. Performing any operation on a `task_block` that is not active results in undefined behavior.
- 4 When the argument to `task_block::run` is called, no `task_block` is active, not even the `task_block` on which `run` was called. (The function object should not, therefore, capture a `task_block` from the surrounding block.)

[Example:

```

define_task_block([&](auto& tb) {
    tb.run([&]{
        tb.run([] { f(); }); // Error: tb is not active within run
        define_task_block([&](auto& tb2) { // Define new task block
            tb2.run(f);
            ...
        });
    });
    ...
});

```

— end example]

[Note: Implementations are encouraged to diagnose the above error at translation time. — end note]

```

template<class F> void run(F&& f);

```

- 5 *Requires:* `F` shall be `MoveConstructible`. `DECAY_COPY(std::forward<F>(f))()` shall be a valid expression.
- 6 *Requires:* `*this` shall be the active `task_block`.
- 7 *Effects:* Evaluates `DECAY_COPY(std::forward<F>(f))()`, where `DECAY_COPY(std::forward<F>(f))` is evaluated synchronously within the current thread. The call to the resulting copy of the function object is permitted to run on an unspecified thread created by the implementation in an unordered fashion relative to the sequence of operations following the call to `run(f)` (the continuation), or indeterminately sequenced within the same thread as the continuation. The call to `run` synchronizes with the call to the function object. The completion of the call to the function object synchronizes with

the next invocation of `wait` on the same `task_block` or completion of the nearest enclosing task block (i.e., the `define_task_block` or `define_task_block_restore_thread` that created this `task_block`).

8 *Throws:* `task_cancelled_exception`, as described in 8.5.

9 *Remarks:* The `run` function may return on a thread other than the one on which it was called; in such cases, completion of the call to `run` synchronizes with the continuation. [*Note:* The return from `run` is ordered similarly to an ordinary function call in a single thread. — *end note*]

10 *Remarks:* The invocation of the user-supplied function object `f` may be immediate or may be delayed until compute resources are available. `run` might or might not return before the invocation of `f` completes.

```
void wait();
```

11 *Preconditions:* `*this` shall be the active `task_block`.

12 *Effects:* Blocks until the tasks spawned using this `task_block` have completed.

13 *Throws:* `task_cancelled_exception`, as described in 8.5.

14 *Postconditions:* All tasks spawned by the nearest enclosing task block have completed.

15 *Remarks:* The `wait` function may return on a thread other than the one on which it was called; in such cases, completion of the call to `wait` synchronizes with subsequent operations. [*Note:* The return from `wait` is ordered similarly to an ordinary function call in a single thread. — *end note*]

[*Example:*

```
define_task_block([&](auto& tb) {
    tb.run([&]{ process(a, w, x); }); // Process a[w] through a[x]
    if (y < x) tb.wait();           // Wait if overlap between [w,x] and [y,z]
    process(a, y, z);              // Process a[y] through a[z]
});
```

— *end example*]

#### 8.4 Function template `define_task_block` [parallel.taskblock.define\_task\_block]

```
template<class F> void define_task_block(F&& f);
template<class F> void define_task_block_restore_thread(F&& f);
```

1 *Requires:* Given an lvalue `tb` of type `task_block`, the expression `f(tb)` shall be well-formed.

2 *Effects:* Constructs a `task_block` `tb` and calls `f(tb)`.

3 *Throws:* `exception_list`, as specified in 8.5.

4 *Postconditions:* All tasks spawned from `f` have finished execution.

5 *Remarks:* The `define_task_block` function may return on a thread other than the one on which it was called unless there are no task blocks active on entry to `define_task_block` (see 8.3), in which case the function returns on the original thread. When `define_task_block` returns on a different thread, it synchronizes with operations following the call. [*Note:* The return from `define_task_block` is ordered similarly to an ordinary function call in a single thread. — *end note*] The `define_task_block_restore_thread` function always returns on the same thread as the one on which it was called.

6 *Notes:* It is expected (but not mandated) that `f` will (directly or indirectly) call `tb.run(function-object)`.

## 8.5 Exception Handling

[parallel.taskblock.exceptions]

- <sup>1</sup> Every `task_block` has an associated exception list. When the task block starts, its associated exception list is empty.
- <sup>2</sup> When an exception is thrown from the user-provided function object passed to `define_task_block` or `define_task_block_restore_thread`, it is added to the exception list for that task block. Similarly, when an exception is thrown from the user-provided function object passed into `task_block::run`, the exception object is added to the exception list associated with the nearest enclosing task block. In both cases, an implementation may discard any pending tasks that have not yet been invoked. Tasks that are already in progress are not interrupted except at a call to `task_block::run` or `task_block::wait` as described below.
- <sup>3</sup> If the implementation is able to detect that an exception has been thrown by another task within the same nearest enclosing task block, then `task_block::run` or `task_block::wait` may throw `task_canceled_exception`; these instances of `task_canceled_exception` are not added to the exception list of the corresponding task block.
- <sup>4</sup> When a task block finishes with a non-empty exception list, the exceptions are aggregated into an `exception_list` object, which is then thrown from the task block.
- <sup>5</sup> The order of the exceptions in the `exception_list` object is unspecified.

## 9 Data-Parallel Types [parallel.simd]

### 9.1 General [parallel.simd.general]

- <sup>1</sup> The data-parallel library consists of data-parallel types and operations on these types. A data-parallel type consists of elements of an underlying arithmetic type, called the *element type*. The number of elements is a constant for each data-parallel type and called the *width* of that type.
- <sup>2</sup> Throughout this Clause, the term *data-parallel type* refers to all *supported* (9.6.1) specializations of the `simd` and `simd_mask` class templates. A *data-parallel object* is an object of *data-parallel type*.
- <sup>3</sup> An *element-wise operation* applies a specified operation to the elements of one or more data-parallel objects. Each such application is unsequenced with respect to the others. A *unary element-wise operation* is an element-wise operation that applies a unary operation to each element of a data-parallel object. A *binary element-wise operation* is an element-wise operation that applies a binary operation to corresponding elements of two data-parallel objects.
- <sup>4</sup> Throughout this Clause, the set of *vectorizable types* for a data-parallel type comprises all cv-unqualified arithmetic types other than `bool`.
- <sup>5</sup> [*Note: The intent is to support acceleration through data-parallel execution resources, such as SIMD registers and instructions or execution units driven by a common instruction decoder. If such execution resources are unavailable, the interfaces support a transparent fallback to sequential execution. — end note*]

### 9.2 Header `<experimental/simd>` synopsis [parallel.simd.synopsis]

```

namespace std::experimental {
inline namespace parallelism_v2 {
  namespace simd_abi {
    using scalar = see below;
    template<int N> using fixed_size = see below;
    template<class T> inline constexpr int max_fixed_size = implementation-defined;
    template<class T> using compatible = implementation-defined;
    template<class T> using native = implementation-defined;

    template<class T, size_t N, class... Abis> struct deduce { using type = see below; };
    template<class T, size_t N, class... Abis> using deduce_t =
      typename deduce<T, N, Abis...>::type;
  }

  struct element_aligned_tag {};
  struct vector_aligned_tag {};
  template<size_t> struct overaligned_tag {};
  inline constexpr element_aligned_tag element_aligned{};
  inline constexpr vector_aligned_tag vector_aligned{};
  template<size_t N> inline constexpr overaligned_tag<N> overaligned{};

  // 9.4, simd type traits
  template<class T> struct is_abi_tag;
  template<class T> inline constexpr bool is_abi_tag_v = is_abi_tag<T>::value;

  template<class T> struct is_simd;
  template<class T> inline constexpr bool is_simd_v = is_simd<T>::value;

```

```

template<class T> struct is_simd_mask;
template<class T> inline constexpr bool is_simd_mask_v = is_simd_mask<T>::value;

template<class T> struct is_simd_flag_type;
template<class T> inline constexpr bool is_simd_flag_type_v =
    is_simd_flag_type<T>::value;

template<class T, class Abi = simd_abi::compatible<T>> struct simd_size;
template<class T, class Abi = simd_abi::compatible<T>>
    inline constexpr size_t simd_size_v = simd_size<T,Abi>::value;

template<class T, class U = typename T::value_type> struct memory_alignment;
template<class T, class U = typename T::value_type>
    inline constexpr size_t memory_alignment_v = memory_alignment<T,U>::value;

template<class T, class V> struct rebind_simd { using type = see below; };
template<class T, class V> using rebind_simd_t = typename rebind_simd<T, V>::type;
template<int N, class V> struct resize_simd { using type = see below; };
template<int N, class V> using resize_simd_t = typename resize_simd<N, V>::type;

// 9.6, Class template simd
template<class T, class Abi = simd_abi::compatible<T>> class simd;
template<class T> using native_simd = simd<T, simd_abi::native<T>>;
template<class T, int N> using fixed_size_simd = simd<T, simd_abi::fixed_size<N>>;

// 9.8, Class template simd_mask
template<class T, class Abi = simd_abi::compatible<T>> class simd_mask;
template<class T> using native_simd_mask = simd_mask<T, simd_abi::native<T>>;
template<class T, int N> using fixed_size_simd_mask =
    simd_mask<T, simd_abi::fixed_size<N>>;

// 9.7.5, Casts
template<class T, class U, class Abi> see below simd_cast(const simd<U, Abi>&) noexcept;
template<class T, class U, class Abi> see below static_simd_cast(const simd<U, Abi>&) noexcept;

template<class T, class Abi>
    fixed_size_simd<T, simd_size_v<T, Abi>>
        to_fixed_size(const simd<T, Abi>&) noexcept;
template<class T, class Abi>
    fixed_size_simd_mask<T, simd_size_v<T, Abi>>
        to_fixed_size(const simd_mask<T, Abi>&) noexcept;
template<class T, int N>
    native_simd<T> to_native(const fixed_size_simd<T, N>&) noexcept;
template<class T, int N>
    native_simd_mask<T> to_native(const fixed_size_simd_mask<T, N>&) noexcept;
template<class T, int N>
    simd<T> to_compatible(const fixed_size_simd<T, N>&) noexcept;
template<class T, int N>
    simd_mask<T> to_compatible(const fixed_size_simd_mask<T, N>&) noexcept;

template<size_t... Sizes, class T, class Abi>
    tuple<simd<T, simd_abi::deduce_t<T, Sizes>>...>
        split(const simd<T, Abi>&) noexcept;
template<size_t... Sizes, class T, class Abi>
    tuple<simd_mask<T, simd_mask_abi::deduce_t<T, Sizes>>...>

```

```

    split(const simd_mask<T, Abi>&) noexcept;
template<class V, class Abi>
    array<V, simd_size_v<typename V::value_type, Abi> / V::size()>
        split(const simd<typename V::value_type, Abi>&) noexcept;
template<class V, class Abi>
    array<V, simd_size_v<typename V::simd_type::value_type, Abi> / V::size()>
        split(const simd_mask<typename V::simd_type::value_type, Abi>&) noexcept;

template<size_t N, class T, class A>
    array<resize_simd<simd_size_v<T, A> / N, simd<T, A>>, N>
        split_by(const simd<T, A>& x) noexcept;
template<size_t N, class T, class A>
    array<resize_simd<simd_size_v<T, A> / N, simd_mask<T, A>>, N>
        split_by(const simd_mask<T, A>& x) noexcept;

template<class T, class... Abis>
    simd<T, simd_abi::deduce_t<T, (simd_size_v<T, Abis> + ...)>>
        concat(const simd<T, Abis>&...) noexcept;
template<class T, class... Abis>
    simd_mask<T, simd_abi::deduce_t<T, (simd_size_v<T, Abis> + ...)>>
        concat(const simd_mask<T, Abis>&...) noexcept;

template<class T, class Abi, size_t N>
    resize_simd<simd_size_v<T, Abi> * N, simd<T, Abi>>
        concat(const array<simd<T, Abi>, N>& arr) noexcept;
template<class T, class Abi, size_t N>
    resize_simd<simd_size_v<T, Abi> * N, simd_mask<T, Abi>>
        concat(const array<simd_mask<T, Abi>, N>& arr) noexcept;

// 9.9.4, Reductions
template<class T, class Abi> bool all_of(const simd_mask<T, Abi>&) noexcept;
template<class T, class Abi> bool any_of(const simd_mask<T, Abi>&) noexcept;
template<class T, class Abi> bool none_of(const simd_mask<T, Abi>&) noexcept;
template<class T, class Abi> bool some_of(const simd_mask<T, Abi>&) noexcept;
template<class T, class Abi> int popcount(const simd_mask<T, Abi>&) noexcept;
template<class T, class Abi> int find_first_set(const simd_mask<T, Abi>&);
template<class T, class Abi> int find_last_set(const simd_mask<T, Abi>&);

bool all_of(T) noexcept;
bool any_of(T) noexcept;
bool none_of(T) noexcept;
bool some_of(T) noexcept;
int popcount(T) noexcept;
int find_first_set(T);
int find_last_set(T);

// 9.5, Where expression class templates
template<class M, class T> class const_where_expression;
template<class M, class T> class where_expression;

// 9.9.5, Where functions
template<class T, class Abi>
    where_expression<simd_mask<T, Abi>, simd<T, Abi>>
        where(const typename simd<T, Abi>::mask_type&, simd<T, Abi>&) noexcept;

```

```

template<class T, class Abi>
  const_where_expression<simd_mask<T, Abi>, simd<T, Abi>>
    where(const typename simd<T, Abi>::mask_type&, const simd<T, Abi>&) noexcept;

template<class T, class Abi>
  where_expression<simd_mask<T, Abi>, simd_mask<T, Abi>>
    where(const type_identity_t<simd_mask<T, Abi>>&, simd_mask<T, Abi>&) noexcept;

template<class T, class Abi>
  const_where_expression<simd_mask<T, Abi>, simd_mask<T, Abi>>
    where(const type_identity_t<simd_mask<T, Abi>>&, const simd_mask<T, Abi>&) noexcept;

template<class T>
  where_expression<bool, T>
    where(see below k, T& d) noexcept;

template<class T>
  const_where_expression<bool, T>
    where(see below k, const T& d) noexcept;

// 9.7.4, Reductions
template<class T, class Abi, class BinaryOperation = plus<>>
  T reduce(const simd<T, Abi>&,
           BinaryOperation = {});

template<class M, class V, class BinaryOperation>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               typename V::value_type identity_element,
                               BinaryOperation binary_op);

template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               plus<> binary_op = {}) noexcept;

template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               multiplies<> binary_op) noexcept;

template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               bit_and<> binary_op) noexcept;

template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               bit_or<> binary_op) noexcept;

template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               bit_xor<> binary_op) noexcept;

template<class T, class Abi>
  T hmin(const simd<T, abi>&) noexcept;
template<class M, class V>
  typename V::value_type hmin(const const_where_expression<M, V>&) noexcept;
template<class T, class Abi>
  T hmax(const simd<T, abi>&) noexcept;
template<class M, class V>
  typename V::value_type hmax(const const_where_expression<M, V>&) noexcept;

// 9.7.6, Algorithms

```

```

template<class T, class Abi>
  simd<T, Abi>
  min(const simd<T, Abi>& a, const simd<T, Abi>& b) noexcept;
template<class T, class Abi>
  simd<T, Abi>
  max(const simd<T, Abi>& a, const simd<T, Abi>& b) noexcept;
template<class T, class Abi>
  pair<simd<T, Abi>, simd<T, Abi>>
  minmax(const simd<T, Abi>& a, const simd<T, Abi>& b) noexcept;
template<class T, class Abi>
  simd<T, Abi>
  clamp(const simd<T, Abi>& v,
        const simd<T, Abi>& lo,
        const simd<T, Abi>& hi);
}
}

```

- <sup>1</sup> The header `<experimental/simd>` defines class templates, tag types, trait types, and function templates for element-wise operations on data-parallel objects.

### 9.3 simd ABI tags

[parallel.simd.abi]

```

namespace simd_abi {
  using scalar = see below;
  template<int N> using fixed_size = see below;
  template<class T> inline constexpr int max_fixed_size = implementation-defined;
  template<class T> using compatible = implementation-defined;
  template<class T> using native = implementation-defined;
}

```

- <sup>1</sup> An *ABI tag* is a type in the `std::experimental::parallelism_v2::simd_abi` namespace that indicates a choice of size and binary representation for objects of data-parallel type. [*Note: The intent is for the size and binary representation to depend on the target architecture. — end note*] The ABI tag, together with a given element type implies a number of elements. ABI tag types are used as the second template argument to `simd` and `simd_mask`.
- <sup>2</sup> [*Note: The ABI tag is orthogonal to selecting the machine instruction set. The selected machine instruction set limits the usable ABI tag types, though (see 9.6.1). The ABI tags enable users to safely pass objects of data-parallel type between translation unit boundaries (e.g. function calls or I/O). — end note*]
- <sup>3</sup> `scalar` is an alias for an unspecified ABI tag that is different from `fixed_size<1>`. Use of the `scalar` tag type requires data-parallel types to store a single element (i.e., `simd<T, simd_abi::scalar>::size()` returns 1).
- <sup>4</sup> The value of `max_fixed_size<T>` is at least 32.
- <sup>5</sup> `fixed_size<N>` is an alias for an unspecified ABI tag. `fixed_size` does not introduce a non-deduced context. Use of the `simd_abi::fixed_size<N>` tag type requires data-parallel types to store `N` elements (i.e. `simd<T, simd_abi::fixed_size<N>>::size()` is `N`). `simd<T, fixed_size<N>>` and `simd_mask<T, fixed_size<N>>` with `N > 0` and `N <= max_fixed_size<T>` shall be supported. Additionally, for every supported `simd<T, Abi>` (see 9.6.1), where `Abi` is an ABI tag that is not a specialization of `simd_abi::fixed_size`, `N == simd<T, Abi>::size()` shall be supported.
- <sup>6</sup> [*Note: It is unspecified whether `simd<T, fixed_size<T, fixed_size<N>>` with `N > max_fixed_size<T>` is supported. The value of `max_fixed_size<T>` can depend on compiler flags and can change between different compiler versions. — end note*]
- <sup>7</sup> [*Note: An implementation can forego ABI compatibility between differently compiled translation units*

for `simd` and `simd_mask` specializations using the same `simd_abi::fixed_size<N>` tag. Otherwise, the efficiency of `simd<T, Abi>` is likely to be better than for `simd<T, fixed_size<simd_size_v<T, Abi>>` (with `Abi` not a specialization of `simd_abi::fixed_size`). — *end note*

<sup>8</sup> An implementation may define additional *extended ABI tag* types in the `std::experimental::parallelism_v2::simd_abi` namespace, to support other forms of data-parallel computation.

<sup>9</sup> `compatible<T>` is an implementation-defined alias for an ABI tag. [*Note*: The intent is to use the ABI tag producing the most efficient data-parallel execution for the element type `T` that ensures ABI compatibility between translation units on the target architecture. — *end note*] [*Example*: Consider a target architecture supporting the extended ABI tags `__simd128` and `__simd256`, where the `__simd256` type requires an optional ISA extension on said architecture. Also, the target architecture does not support `long double` with either ABI tag. The implementation therefore defines `compatible<T>` as an alias for:

(9.1) — `scalar` if `T` is the same type as `long double`, and

(9.2) — `__simd128` otherwise.

— *end example*]

<sup>10</sup> `native<T>` is an implementation-defined alias for an ABI tag. [*Note*: The intent is to use the ABI tag producing the most efficient data-parallel execution for the element type `T` that is supported on the currently targeted system. For target architectures without ISA extensions, the `native<T>` and `compatible<T>` aliases will likely be the same. For target architectures with ISA extensions, compiler flags may influence the `native<T>` alias while `compatible<T>` will be the same independent of such flags. — *end note*] [*Example*: Consider a target architecture supporting the extended ABI tags `__simd128` and `__simd256`, where hardware support for `__simd256` only exists for floating-point types. The implementation therefore defines `native<T>` as an alias for

(10.1) — `__simd256` if `T` is a floating-point type, and

(10.2) — `__simd128` otherwise.

— *end example*]

```
template<T, size_t N, class... Abis> struct deduce { using type = see below; };
```

<sup>11</sup> The member type shall be present if and only if

(11.1) — `T` is a vectorizable type, and

(11.2) — `simd_abi::fixed_size<N>` is supported (see 9.3), and

(11.3) — every type in the `Abis` pack is an ABI tag.

<sup>12</sup> Where present, the member typedef `type` shall name an ABI tag type that satisfies

(12.1) — `simd_size<T, type> == N`, and

(12.2) — `simd<T, type>` is default constructible (see 9.6.1).

If `N` is 1, the member typedef `type` is `simd_abi::scalar`. Otherwise, if there are multiple ABI tag types that satisfy the constraints, the member typedef `type` is implementation-defined. [*Note*: It is expected that extended ABI tags can produce better optimizations and thus are preferred over `simd_abi::fixed_size<N>`. Implementations can base the choice on `Abis`, but can also ignore the `Abis` arguments. — *end note*]

<sup>13</sup> The behavior of a program that adds specializations for `deduce` is undefined.

## 9.4 simd type traits

[parallel.simd.traits]

```
template<class T> struct is_abi_tag { see below };
```

1 The type `is_abi_tag<T>` is a `UnaryTypeTrait` with a base characteristic of `true_type` if `T` is a standard or extended ABI tag, and `false_type` otherwise.

2 The behavior of a program that adds specializations for `is_abi_tag` is undefined.

```
template<class T> struct is_simd { see below };
```

3 The type `is_simd<T>` is a `UnaryTypeTrait` with a base characteristic of `true_type` if `T` is a specialization of the `simd` class template, and `false_type` otherwise.

4 The behavior of a program that adds specializations for `is_simd` is undefined.

```
template<class T> struct is_simd_mask { see below };
```

5 The type `is_simd_mask<T>` is a `UnaryTypeTrait` with a base characteristic of `true_type` if `T` is a specialization of the `simd_mask` class template, and `false_type` otherwise.

6 The behavior of a program that adds specializations for `is_simd_mask` is undefined.

```
template<class T> struct is_simd_flag_type { see below };
```

7 The type `is_simd_flag_type<class T>` is a `UnaryTypeTrait` with a base characteristic of `true_type` if `T` is one of

(7.1) — `element_aligned_tag`, or

(7.2) — `vector_aligned_tag`, or

(7.3) — `overaligned_tag<N>` with  $N > 0$  and  $N$  an integral power of two,

and `false_type` otherwise.

8 The behavior of a program that adds specializations for `is_simd_flag_type` is undefined.

```
template<class T, class Abi = simd_abi::compatible<T>> struct simd_size { see below };
```

9 `simd_size<T, Abi>` shall have a member `value` if and only if

(9.1) — `T` is a vectorizable type, and

(9.2) — `is_abi_tag_v<Abi>` is true.

[*Note:* The rules are different from those in (9.6.1). — *end note*]

10 If `value` is present, the type `simd_size<T, Abi>` is a `BinaryTypeTrait` with a base characteristic of `integral_constant<size_t, N>` with  $N$  equal to the number of elements in a `simd<T, Abi>` object. [*Note:* If `simd<T, Abi>` is not supported for the currently targeted system, `simd_size<T, Abi>::value` produces the value `simd<T, Abi>::size()` would return if it were supported. — *end note*]

11 The behavior of a program that adds specializations for `simd_size` is undefined.

```
template<class T, class U = typename T::value_type> struct memory_alignment { see below };
```

12 `memory_alignment<T, U>` shall have a member `value` if and only if

(12.1) — `is_simd_mask_v<T>` is true and `U` is `bool`, or

(12.2) — `is_simd_v<T>` is true and `U` is a vectorizable type.

13 If `value` is present, the type `memory_alignment<T, U>` is a `BinaryTypeTrait` with a base characteristic of `integral_constant<size_t, N>` for some implementation-defined `N` (see 9.6.5 and 9.8.4). [*Note: value identifies the alignment restrictions on pointers used for (converting) loads and stores for the give type T on arrays of type U. — end note*]

14 The behavior of a program that adds specializations for `memory_alignment` is undefined.

```
template<class T, class V> struct rebind_simd { using type = see below; };
```

15 The member type is present if and only if

(15.1) — `V` is either `simd<U, Abi0>` or `simd_mask<U, Abi0>`, where `U` and `Abi0` are deduced from `V`, and

(15.2) — `T` is a vectorizable type, and

(15.3) — `simd_abi::deduce<T, simd_size_v<U, Abi0>, Abi0>` has a member type `type`.

16 Let `Abi1` denote the type `deduce_t<T, simd_size_v<U, Abi0>, Abi0>`. Where present, the member typedef `type` names `simd<T, Abi1>` if `V` is `simd<U, Abi0>` or `simd_mask<T, Abi1>` if `V` is `simd_mask<U, Abi0>`.

```
template<int N, class V> struct resize_simd { using type = see below; };
```

17 The member type is present if and only if

(17.1) — `V` is either `simd<T, Abi0>` or `simd_mask<T, Abi0>`, where `T` and `Abi0` are deduced from `V`, and

(17.2) — `simd_abi::deduce<T, N, Abi0>` has a member type `type`.

18 Let `Abi1` denote the type `deduce_t<T, N, Abi0>`. Where present, the member typedef `type` names `simd<T, Abi1>` if `V` is `simd<T, Abi0>` or `simd_mask<T, Abi1>` if `V` is `simd_mask<T, Abi0>`.

## 9.5 Where expression class templates

[parallel.simd.whereexpr]

```
template<class M, class T> class const_where_expression {
    const M mask;    // exposition only
    T& data;         // exposition only

public:
    const_where_expression(const const_where_expression&) = delete;
    const_where_expression& operator=(const const_where_expression&) = delete;

    T operator-() const && noexcept;
    T operator+() const && noexcept;
    T operator~() const && noexcept;

    template<class U, class Flags> void copy_to(U* mem, Flags f) const &&;
};

template<class M, class T>
class where_expression : public const_where_expression<M, T> {
public:
    template<class U> void operator=(U&& x) && noexcept;
    template<class U> void operator+=(U&& x) && noexcept;
    template<class U> void operator-=(U&& x) && noexcept;
    template<class U> void operator*=(U&& x) && noexcept;
    template<class U> void operator/=(U&& x) && noexcept;
    template<class U> void operator%=(U&& x) && noexcept;
    template<class U> void operator&=(U&& x) && noexcept;
};
```

```

template<class U> void operator|=(U&& x) && noexcept;
template<class U> void operator^=(U&& x) && noexcept;
template<class U> void operator<<=(U&& x) && noexcept;
template<class U> void operator>>=(U&& x) && noexcept;

void operator++() && noexcept;
void operator++(int) && noexcept;
void operator--() && noexcept;
void operator--(int) && noexcept;

template<class U, class Flags> void copy_from(const U* mem, Flags) &&;
};

```

- 1 The class templates `const_where_expression` and `where_expression` abstract the notion of selecting elements of a given object of arithmetic or data-parallel type.
- 2 The first templates argument `M` shall be cv-unqualified `bool` or a cv-unqualified `simd_mask` specialization.
- 3 If `M` is `bool`, `T` shall be a cv-unqualified arithmetic type. Otherwise, `T` shall either be `M` or `typename M::simd_type`.
- 4 In this subclause, if `M` is `bool`, `data[0]` is used interchangeably for `data`, `mask[0]` is used interchangeably for `mask`, and `M::size()` is used interchangeably for 1.
- 5 The *selected indices* signify the integers  $i \in \{j \in \mathbb{N} | j < M::size() \wedge \text{mask}[j]\}$ . The *selected elements* signify the elements `data[i]` for all selected indices  $i$ .
- 6 In this subclause, the type `value_type` is an alias for `T` if `M` is `bool`, or an alias for `typename T::value_type` if `is_simd_mask_v<M>` is true.
- 7 [*Note: The where functions 9.9.5 initialize mask with the first argument to where and data with the second argument to where. — end note*]

```

T operator-() const && noexcept;
T operator+() const && noexcept;
T operator~() const && noexcept;

```

- 8 *Returns:* A copy of `data` with the indicated unary operator applied to all selected elements.

```

template<class U, class Flags> void copy_to(U* mem, Flags) const &&;

```

- 9 *Requires:*

- (9.1) — If `M` is not `bool`, the largest selected index is less than the number of values pointed to by `mem`.
- (9.2) — If the template parameter `Flags` is `vector_aligned_tag`, `mem` shall point to storage aligned by `memory_alignment_v<T, U>`.
- (9.3) — If the template parameter `Flags` is `overaligned_tag<N>`, `mem` shall point to storage aligned by `N`.
- (9.4) — If the template parameter `Flags` is `element_aligned_tag`, `mem` shall point to storage aligned by `alignof(U)`.

- 10 *Effects:* Copies the selected elements as if `mem[i] = static_cast<U>(data[i])` for all selected indices  $i$ .

- 11 *Throws:* Nothing.

- 12 *Remarks:* This function shall not participate in overload resolution unless

- (12.1) — `is_simd_flag_type_v<Flags>` is true, and
- (12.2) — either

- (12.2.1) — U is bool and value\_type is bool, or  
 (12.2.2) — U is a vectorizable type and value\_type is not bool.

```
template<class U> void operator=(U&& x) && noexcept;
```

- 13 *Effects:* Replaces `data[i]` with `static_cast<T>(std::forward<U>(x)) [i]` for all selected indices *i*.  
 14 *Remarks:* This operator shall not participate in overload resolution unless U is convertible to T.

```
template<class U> void operator+=(U&& x) && noexcept;
template<class U> void operator-=(U&& x) && noexcept;
template<class U> void operator*=(U&& x) && noexcept;
template<class U> void operator/=(U&& x) && noexcept;
template<class U> void operator%=(U&& x) && noexcept;
template<class U> void operator&=(U&& x) && noexcept;
template<class U> void operator|=(U&& x) && noexcept;
template<class U> void operator^=(U&& x) && noexcept;
template<class U> void operator<<=(U&& x) && noexcept;
template<class U> void operator>>=(U&& x) && noexcept;
```

- 15 *Effects:* Replaces `data[i]` with `static_cast<T>(data @ std::forward<U>(x)) [i]` (where @ denotes the indicated operator) for all selected indices *i*.  
 16 *Remarks:* Each of these operators shall not participate in overload resolution unless the return type of `data @ std::forward<U>(x)` is convertible to T. It is unspecified whether the binary operator, implied by the compound assignment operator, is executed on all elements or only on the selected elements.

```
void operator++() && noexcept;
void operator++(int) && noexcept;
void operator--() && noexcept;
void operator--(int) && noexcept;
```

- 17 *Effects:* Applies the indicated operator to the selected elements.  
 18 *Remarks:* Each of these operators shall not participate in overload resolution unless the indicated operator can be applied to objects of type T.

```
template<class U, class Flags> void copy_from(const U* mem, Flags) &&;
```

- 19 *Requires:*

- (19.1) — If `is_simd_flag_type_v<U>` is true, for all selected indices *i*, *i* shall be less than the number of values pointed to by `mem`.  
 (19.2) — If the template parameter `Flags` is `vector_aligned_tag`, `mem` shall point to storage aligned by `memory_alignment_v<T, U>`.  
 (19.3) — If the template parameter `Flags` is `overaligned_tag<N>`, `mem` shall point to storage aligned by *N*.  
 (19.4) — If the template parameter `Flags` is `element_aligned_tag`, `mem` shall point to storage aligned by `alignof(U)`.

- 20 *Effects:* Replaces the selected elements as if `data[i] = static_cast<value_type>(mem[i])` for all selected indices *i*.

- 21 *Throws:* Nothing.

- 22 *Remarks:* This function shall not participate in overload resolution unless

- (22.1) — `is_simd_flag_type_v<Flags>` is true, and  
 (22.2) — either  
 (22.2.1) — U is bool and value\_type is bool, or  
 (22.2.2) — U is a vectorizable type and value\_type is not bool.

## 9.6 Class template simd

[parallel.simd.class]

## 9.6.1 Class template simd overview

[parallel.simd.overview]

```

template<class T, class Abi> class simd {
public:
    using value_type = T;
    using reference = see below;
    using mask_type = simd_mask<T, Abi>;
    using abi_type = Abi;

    static constexpr size_t size() noexcept;

    simd() noexcept = default;

    // 9.6.4, simd constructors
    template<class U> simd(U&& value) noexcept;
    template<class U> simd(const simd<U, simd_abi::fixed_size<size()>>&) noexcept;
    template<class G> explicit simd(G&& gen) noexcept;
    template<class U, class Flags> simd(const U* mem, Flags f);

    // 9.6.5, simd copy functions
    template<class U, class Flags> copy_from(const U* mem, Flags f);
    template<class U, class Flags> copy_to(U* mem, Flags f);

    // 9.6.6, simd subscript operators
    reference operator[](size_t);
    value_type operator[](size_t) const;

    // 9.6.7, simd unary operators
    simd& operator++() noexcept;
    simd operator++(int) noexcept;
    simd& operator--() noexcept;
    simd operator--(int) noexcept;
    mask_type operator!() const noexcept;
    simd operator~() const noexcept;
    simd operator+() const noexcept;
    simd operator-() const noexcept;

    // 9.7.1, simd binary operators
    friend simd operator+(const simd&, const simd&) noexcept;
    friend simd operator-(const simd&, const simd&) noexcept;
    friend simd operator*(const simd&, const simd&) noexcept;
    friend simd operator/(const simd&, const simd&) noexcept;
    friend simd operator%(const simd&, const simd&) noexcept;
    friend simd operator&(const simd&, const simd&) noexcept;
    friend simd operator|(const simd&, const simd&) noexcept;
    friend simd operator^(const simd&, const simd&) noexcept;
    friend simd operator<<(const simd&, const simd&) noexcept;
    friend simd operator>>(const simd&, const simd&) noexcept;
    friend simd operator<<(const simd&, int) noexcept;
    friend simd operator>>(const simd&, int) noexcept;

    // 9.7.2, simd compound assignment
    friend simd& operator+=(simd&, const simd&) noexcept;
    friend simd& operator-=(simd&, const simd&) noexcept;

```

```

friend simd& operator*=(simd&, const simd&) noexcept;
friend simd& operator/=(simd&, const simd&) noexcept;
friend simd& operator%=(simd&, const simd&) noexcept;
friend simd& operator&=(simd&, const simd&) noexcept;
friend simd& operator|=(simd&, const simd&) noexcept;
friend simd& operator^=(simd&, const simd&) noexcept;
friend simd& operator<<=(simd&, const simd&) noexcept;
friend simd& operator>>=(simd&, const simd&) noexcept;
friend simd& operator<<=(simd&, int) noexcept;
friend simd& operator>>=(simd&, int) noexcept;

// 9.7.3, simd compare operators
friend mask_type operator==(const simd&, const simd&) noexcept;
friend mask_type operator!=(const simd&, const simd&) noexcept;
friend mask_type operator>=(const simd&, const simd&) noexcept;
friend mask_type operator<=(const simd&, const simd&) noexcept;
friend mask_type operator>(const simd&, const simd&) noexcept;
friend mask_type operator<(const simd&, const simd&) noexcept;
};

```

- <sup>1</sup> The class template `simd` is a data-parallel type. The width of a given `simd` specialization is a constant expression, determined by the template parameters.
- <sup>2</sup> Every specialization of `simd` shall be a complete type. The specialization `simd<T, Abi>` is supported if `T` is a vectorizable type and
- (2.1) — `Abi` is `simd_abi::scalar`, or
  - (2.2) — `Abi` is `simd_abi::fixed_size<N>`, with `N` constrained as defined in 9.3.

If `Abi` is an extended ABI tag, it is implementation-defined whether `simd<T, Abi>` is supported. [*Note: The intent is for implementations to decide on the basis of the currently targeted system. — end note*]

If `simd<T, Abi>` is not supported, the specialization shall have a deleted default constructor, deleted destructor, deleted copy constructor, and deleted copy assignment. Otherwise, the following are true:

- (2.3) — `is_nothrow_move_constructible_v<simd<T, Abi>>`, and
- (2.4) — `is_nothrow_move_assignable_v<simd<T, Abi>>`, and
- (2.5) — `is_nothrow_default_constructible_v<simd<T, Abi>>`.

[*Example: Consider an implementation that defines the extended ABI tags `__simd_x` and `__gpu_y`. When the compiler is invoked to translate to a machine that has support for the `__simd_x` ABI tag for all arithmetic types other than `long double` and no support for the `__gpu_y` ABI tag, then:*

- (2.6) — `simd<T, simd_abi::__gpu_y>` is not supported for any `T` and has a deleted constructor.
- (2.7) — `simd<long double, simd_abi::__simd_x>` is not supported and has a deleted constructor.
- (2.8) — `simd<double, simd_abi::__simd_x>` is supported.
- (2.9) — `simd<long double, simd_abi::scalar>` is supported.

— *end example*]

- 3 Default initialization performs no initialization of the elements; value-initialization initializes each element with T(). [*Note:* Thus, default initialization leaves the elements in an indeterminate state. — *end note*]
- 4 Implementations should enable explicit conversion from and to implementation-defined types. This adds one or more of the following declarations to class `simd`:

```
explicit operator implementation-defined() const;
explicit simd(const implementation-defined& init);
```

[*Example:* Consider an implementation that supports the type `__vec4f` and the function `__vec4f __vec4f _addsub(__vec4f, __vec4f)` for the currently targeted system. A user may require the use of `_vec4f_addsub` for maximum performance and thus writes:

```
using V = simd<float, simd_abi::__simd128>;
V addsub(V a, V b) {
    return static_cast<V>(_vec4f_addsub(static_cast<__vec4f>(a), static_cast<__vec4f>(b)));
}
```

— *end example*]

### 9.6.2 `simd` width

[parallel.simd.width]

```
static constexpr size_t size() noexcept;
```

- 1 *Returns:* The width of `simd<T, Abi>`.

### 9.6.3 Element references

[parallel.simd.reference]

- 1 A `reference` is an object that refers to an element in a `simd` or `simd_mask` object. `reference::value_type` is the same type as `simd::value_type` or `simd_mask::value_type`, respectively.
- 2 Class `reference` is for exposition only. An implementation is permitted to provide equivalent functionality without providing a class with this name.

```
class reference // exposition only
{
public:
    reference() = delete;
    reference(const reference&) = delete;

    operator value_type() const noexcept;

    template<class U> reference operator=(U&& x) && noexcept;

    template<class U> reference operator+=(U&& x) && noexcept;
    template<class U> reference operator-=(U&& x) && noexcept;
    template<class U> reference operator*=(U&& x) && noexcept;
    template<class U> reference operator/=(U&& x) && noexcept;
    template<class U> reference operator%=(U&& x) && noexcept;
    template<class U> reference operator|=(U&& x) && noexcept;
    template<class U> reference operator&=(U&& x) && noexcept;
    template<class U> reference operator^=(U&& x) && noexcept;
    template<class U> reference operator<<=(U&& x) && noexcept;
    template<class U> reference operator>>=(U&& x) && noexcept;

    reference operator++() && noexcept;
    value_type operator++(int) && noexcept;
    reference operator--() && noexcept;
```

```
value_type operator--(int) && noexcept;
```

```
friend void swap(reference&& a, reference&& b) noexcept;
friend void swap(value_type&& a, reference&& b) noexcept;
friend void swap(reference&& a, value_type&& b) noexcept;
};
```

```
operator value_type() const noexcept;
```

3 *Returns:* The value of the element referred to by *\*this*.

```
template<class U> reference operator=(U&& x) && noexcept;
```

4 *Effects:* Replaces the referred to element in `simd` or `simd_mask` with `static_cast<value_type>(std::forward<U>(x))`.

5 *Returns:* A copy of *\*this*.

6 *Remarks:* This function shall not participate in overload resolution unless `declval<value_type>() = std::forward<U>(x)` is well-formed.

```
template<class U> reference operator+=(U&& x) && noexcept;
template<class U> reference operator-=(U&& x) && noexcept;
template<class U> reference operator*=(U&& x) && noexcept;
template<class U> reference operator/=(U&& x) && noexcept;
template<class U> reference operator%=(U&& x) && noexcept;
template<class U> reference operator|=(U&& x) && noexcept;
template<class U> reference operator&=(U&& x) && noexcept;
template<class U> reference operator^=(U&& x) && noexcept;
template<class U> reference operator<<=(U&& x) && noexcept;
template<class U> reference operator>>=(U&& x) && noexcept;
```

7 *Effects:* Applies the indicated compound operator to the referred to element in `simd` or `simd_mask` and `std::forward<U>(x)`.

8 *Returns:* A copy of *\*this*.

9 *Remarks:* This function shall not participate in overload resolution unless `declval<value_type>() @= std::forward<U>(x)` (where `@=` denotes the indicated compound assignment operator) is well-formed.

```
reference operator++() && noexcept;
reference operator--() && noexcept;
```

10 *Effects:* Applies the indicated operator to the referred to element in `simd` or `simd_mask`.

11 *Returns:* A copy of *\*this*.

12 *Remarks:* This function shall not participate in overload resolution unless the indicated operator can be applied to objects of type `value_type`.

```
value_type operator++(int) && noexcept;
value_type operator--(int) && noexcept;
```

13 *Effects:* Applies the indicated operator to the referred to element in `simd` or `simd_mask`.

14 *Returns:* A copy of the referred to element before applying the indicated operator.

15 *Remarks:* This function shall not participate in overload resolution unless the indicated operator can be applied to objects of type `value_type`.

```
friend void swap(reference&& a, reference&& b) noexcept;
friend void swap(value_type& a, reference&& b) noexcept;
friend void swap(reference&& a, value_type& b) noexcept;
```

16 *Effects:* Exchanges the values `a` and `b` refer to.

#### 9.6.4 simd constructors

[parallel.simd.ctor]

```
template<class U> simd(U&&) noexcept;
```

1 *Effects:* Constructs an object with each element initialized to the value of the argument after conversion to `value_type`.

2 *Remarks:* Let `From` denote the type `remove_cv_t<remove_reference_t<U>>`. This constructor shall not participate in overload resolution unless:

- (2.1) — `From` is a vectorizable type and every possibly value of `From` can be represented with type `value_type`, or
- (2.2) — `From` is not an arithmetic type and is implicitly convertible to `value_type`, or
- (2.3) — `From` is `int`, or
- (2.4) — `From` is `unsigned int` and `value_type` is an unsigned integral type.

```
template<class U> simd(const simd<U, simd_abi::fixed_size<size()>>& x) noexcept;
```

3 *Effects:* Constructs an object where the  $i^{\text{th}}$  element equals `static_cast<T>(x[i])` for all  $i$  in the range of `[0, size())`.

4 *Remarks:* This constructor shall not participate in overload resolution unless

- (4.1) — `abi_type` is `simd_abi::fixed_size<size()>`, and
- (4.2) — every possible value of `U` can be represented with type `value_type`, and
- (4.3) — if both `U` and `value_type` are integral, the integer conversion rank (C++17 §7.15) of `value_type` is greater than the integer conversion rank of `U`.

```
template<class G> simd(G&& gen) noexcept;
```

5 *Effects:* Constructs an object where the  $i^{\text{th}}$  element is initialized to `gen(integral_constant<size_t, i>())`.

6 *Remarks:* This constructor shall not participate in overload resolution unless `simd(gen(integral_constant<size_t, i>()))` is well-formed for all  $i$  in the range of `[0, size())`.

7 The calls to `gen` are unsequenced with respect to each other. Vectorization-unsafe standard library functions may not be invoked by `gen` (C++17 §28.4.3).

```
template<class U, class Flags> simd(const U* mem, Flags);
```

8 *Requires:*

- (8.1) `[mem, mem + size())` is a valid range.
- (8.2) — If the template parameter `Flags` is `vector_aligned_tag`, `mem` shall point to storage aligned by `memory_alignment_v<simd, U>`.
- (8.3) — If the template parameter `Flags` is `overaligned_tag<N>`, `mem` shall point to storage aligned by `N`.
- (8.4) — If the template parameter `Flags` is `element_aligned_tag`, `mem` shall point to storage aligned by `alignof(U)`.

9 *Effects:* Constructs an object where the  $i^{\text{th}}$  element is initialized to `static_cast<T>(mem[i])` for all  $i$  in the range of `[0, size())`.

10 *Remarks:* This constructor shall not participate in overload resolution unless

(10.1) — `is_simd_flag_type_v<Flags>` is true, and

(10.2) — `U` is a vectorizable type.

### 9.6.5 simd copy functions

[parallel.simd.copy]

```
template<class U, class Flags> void copy_from(const U* mem, Flags);
```

1 *Requires:*

(1.1) — `[mem, mem + size())` is a valid range.

(1.2) — If the template parameter `Flags` is `vector_aligned_tag`, `mem` shall point to storage aligned by `memory_alignment_v<simd, U>`.

(1.3) — If the template parameter `Flags` is `overaligned_tag<N>`, `mem` shall point to storage aligned by `N`.

(1.4) — If the template parameter `Flags` is `element_aligned_tag`, `mem` shall point to storage aligned by `alignof(U)`.

2 *Effects:* Replaces the elements of the `simd` object such that the  $i^{\text{th}}$  element is assigned with `static_cast<T>(mem[i])` for all  $i$  in the range of `[0, size())`.

3 *Remarks:* This function shall not participate in overload resolution unless

(3.1) — `is_simd_flag_type_v<Flags>` is true, and

(3.2) — `U` is a vectorizable type.

```
template<class U, class Flags> void copy_to(U* mem, Flags) const;
```

4 *Requires:*

(4.1) — `[mem, mem + size())` is a valid range.

(4.2) — If the template parameter `Flags` is `vector_aligned_tag`, `mem` shall point to storage aligned by `memory_alignment_v<simd, U>`.

(4.3) — If the template parameter `Flags` is `overaligned_tag<N>`, `mem` shall point to storage aligned by `N`.

(4.4) — If the template parameter `Flags` is `element_aligned_tag`, `mem` shall point to storage aligned by `alignof(U)`.

5 *Effects:* Copies all `simd` elements as if `mem[i] = static_cast<U>(operator[](i))` for all  $i$  in the range of `[0, size())`.

6 *Remarks:* This function shall not participate in overload resolution unless

(6.1) — `is_simd_flag_type_v<Flags>` is true, and

(6.2) — `U` is a vectorizable type.

**9.6.6 simd subscript operators**

[parallel.simd.subscr]

reference operator[](size\_t i);

1 *Requires:*  $i < \text{size}()$ .2 *Returns:* A reference (see 9.6.3) referring to the  $i^{\text{th}}$  element.3 *Throws:* Nothing.

value\_type operator[](size\_t i) const;

4 *Requires:*  $i < \text{size}()$ .5 *Returns:* The value of the  $i^{\text{th}}$  element.6 *Throws:* Nothing.**9.6.7 simd unary operators**

[parallel.simd.unary]

1 Effects in this subclause are applied as unary element-wise operations.

simd&amp; operator++() noexcept;

2 *Effects:* Increments every element by one.3 *Returns:* \*this.

simd operator++(int) noexcept;

4 *Effects:* Increments every element by one.5 *Returns:* A copy of \*this before incrementing.

simd&amp; operator--() noexcept;

6 *Effects:* Decrements every element by one.7 *Returns:* \*this.

simd operator--(int) noexcept;

8 *Effects:* Decrements every element by one.9 *Returns:* A copy of \*this before decrementing.

mask\_type operator!() const noexcept;

10 *Returns:* A simd\_mask object with the  $i^{\text{th}}$  element set to !operator[]( $i$ ) for all  $i$  in the range of  $[0, \text{size}())$ .

simd operator~() const noexcept;

11 *Returns:* A simd object where each bit is the inverse of the corresponding bit in \*this.12 *Remarks:* This operator shall not participate in overload resolution unless T is an integral type.

simd operator+() const noexcept;

13 *Returns:* \*this.

simd operator-() const noexcept;

14 *Returns:* A simd object where the  $i^{\text{th}}$  element is initialized to -operator[]( $i$ ) for all  $i$  in the range of  $[0, \text{size}())$ .

**9.7 simd non-member operations**

[parallel.simd.nonmembers]

**9.7.1 simd binary operators**

[parallel.simd.binary]

```
friend simd operator+(const simd& lhs, const simd& rhs) noexcept;
friend simd operator-(const simd& lhs, const simd& rhs) noexcept;
friend simd operator*(const simd& lhs, const simd& rhs) noexcept;
friend simd operator/(const simd& lhs, const simd& rhs) noexcept;
friend simd operator%(const simd& lhs, const simd& rhs) noexcept;
friend simd operator&(const simd& lhs, const simd& rhs) noexcept;
friend simd operator|(const simd& lhs, const simd& rhs) noexcept;
friend simd operator^(const simd& lhs, const simd& rhs) noexcept;
friend simd operator<<(const simd& lhs, const simd& rhs) noexcept;
friend simd operator>>(const simd& lhs, const simd& rhs) noexcept;
```

1 *Returns:* A simd object initialized with the results of applying the indicated operator to lhs and rhs as a binary element-wise operation.

2 *Remarks:* Each of these operators shall not participate in overload resolution unless the indicated operator can be applied to objects of type value\_type.

```
friend simd operator<<(const simd& v, int n) noexcept;
friend simd operator>>(const simd& v, int n) noexcept;
```

3 *Returns:* A simd object where the  $i^{\text{th}}$  element is initialized to the result of applying the indicated operator to  $v[i]$  and  $n$  for all  $i$  in the range of  $[0, \text{size}())$ .

4 *Remarks:* These operators shall not participate in overload resolution unless the indicated operator can be applied to objects of type value\_type.

**9.7.2 simd compound assignment**

[parallel.simd.cassign]

```
friend simd& operator+=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator-=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator*=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator/=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator%=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator&=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator|=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator^=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator<<=(simd& lhs, const simd& rhs) noexcept;
friend simd& operator>>=(simd& lhs, const simd& rhs) noexcept;
```

1 *Effects:* These operators apply the indicated operator to lhs and rhs as an element-wise operation.

2 *Returns:* lhs

3 *Remarks:* These operators shall not participate in overload resolution unless the indicated operator can be applied to objects of type value\_type.

```
friend simd& operator<<=(simd& lhs, int n) noexcept;
friend simd& operator>>=(simd& lhs, int n) noexcept;
```

4 *Effects:* Equivalent to: `return operator@=(lhs, simd(n));`

5 *Remarks:* These operators shall not participate in overload resolution unless the indicated operator can be applied to objects of type value\_type.

**9.7.3 simd compare operators**

[parallel.simd.comparison]

```
friend mask_type operator==(const simd& lhs, const simd& rhs) noexcept;
friend mask_type operator!=(const simd& lhs, const simd& rhs) noexcept;
friend mask_type operator>=(const simd& lhs, const simd& rhs) noexcept;
friend mask_type operator<=(const simd& lhs, const simd& rhs) noexcept;
friend mask_type operator>(const simd& lhs, const simd& rhs) noexcept;
friend mask_type operator<(const simd& lhs, const simd& rhs) noexcept;
```

- 1 *Returns:* A `simd_mask` object initialized with the results of applying the indicated operator to `lhs` and `rhs` as a binary element-wise operation.

**9.7.4 Reductions**

[parallel.simd.reductions]

- 1 In this subclause, `BinaryOperation` shall be a binary element-wise operation.

```
template<class T, class Abi, class BinaryOperation = plus<>>
  T reduce(const simd<T, Abi>& x, BinaryOperation binary_op = {});
```

- 2 *Requires:* `binary_op` shall be callable with two arguments of type `T` returning `T`, or callable with two arguments of type `simd<T, A1>` returning `simd<T, A1>` for every `A1` that is an ABI tag type.
- 3 *Returns:* `GENERALIZED_SUM(binary_op, x.data[i], ...)` for all  $i$  in the range of  $[0, \text{size}())$  (C++17 §29.2).
- 4 *Throws:* Any exception thrown from `binary_op`.

```
template<class M, class V, class BinaryOperation>
  typename V::value_type reduce(const const_where_expression<M, V>& x,
                               typename V::value_type identity_element,
                               BinaryOperation binary_op = {});
```

- 5 *Requires:* `binary_op` shall be callable with two arguments of type `T` returning `T`, or callable with two arguments of type `simd<T, A1>` returning `simd<T, A1>` for every `A1` that is an ABI tag type. The results of `binary_op(identity_element, x)` and `binary_op(x, identity_element)` shall be equal to `x` for all finite values `x` representable by `V::value_type`.
- 6 *Returns:* If `none_of(x.mask)`, returns `identity_element`. Otherwise, returns `GENERALIZED_SUM(binary_op, x.data[i], ...)` for all selected indices  $i$ .
- 7 *Throws:* Any exception thrown from `binary_op`.

```
template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x, plus<> binary_op) noexcept;
```

- 8 *Returns:* If `none_of(x.mask)`, returns 0. Otherwise, returns `GENERALIZED_SUM(binary_op, x.data[i], ...)` for all selected indices  $i$ .

```
template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x, multiplies<> binary_op) noexcept;
```

- 9 *Returns:* If `none_of(x.mask)`, returns 1. Otherwise, returns `GENERALIZED_SUM(binary_op, x.data[i], ...)` for all selected indices  $i$ .

```
template<class M, class V>
  typename V::value_type reduce(const const_where_expression<M, V>& x, bit_and<> binary_op) noexcept;
```

- 10 *Requires:* `is_integral_v<V::value_type>` is true.
- 11 *Returns:* If `none_of(x.mask)`, returns `~V::value_type()`. Otherwise, returns `GENERALIZED_SUM(binary_op, x.data[i], ...)` for all selected indices  $i$ .

```

template<class M, class V>
    typename V::value_type reduce(const const_where_expression<M, V>& x, bit_or<> binary_op) noexcept;
template<class M, class V>
    typename V::value_type reduce(const const_where_expression<M, V>& x, bit_xor<> binary_op) noexcept;

```

12 *Requires:* `is_integral_v<V::value_type>` is true.

13 *Returns:* If `none_of(x.mask)`, returns 0. Otherwise, returns *GENERALIZED\_SUM*(`binary_op`, `x.data[i]`, ...) for all selected indices *i*.

```

template<class T, class Abi> T hmin(const simd<T, Abi>& x) noexcept;

```

14 *Returns:* The value of an element `x[j]` for which `x[j] <= x[i]` for all *i* in the range of `[0, size())`.

```

template<class M, class V> typename V::value_type hmin(const const_where_expression<M, V>& x) noexcept;

```

15 *Returns:* If `none_of(x.mask)`, the return value is `numeric_limits<V::value_type>::max()`. Otherwise, returns the value of an element `x.data[j]` for which `x.mask[j] == true` and `x.data[j] <= x.data[i]` for all selected indices *i*.

```

template<class T, class Abi> T hmax(const simd<T, Abi>& x) noexcept;

```

16 *Returns:* The value of an element `x[j]` for which `x[j] >= x[i]` for all *i* in the range of `[0, size())`.

```

template<class M, class V> typename V::value_type hmax(const const_where_expression<M, V>& x) noexcept;

```

17 *Returns:* If `none_of(x.mask)`, the return value is `numeric_limits<V::value_type>::lowest()`. Otherwise, returns the value of an element `x.data[j]` for which `x.mask[j] == true` and `x.data[j] >= x.data[i]` for all selected indices *i*.

### 9.7.5 Casts

[parallel.simd.casts]

```

template<class T, class U, class Abi> see below simd_cast(const simd<U, Abi>& x) noexcept;

```

1 Let `To` denote `T::value_type` if `is_simd_v<T>` is true, or `T` otherwise.

2 *Returns:* A `simd` object with the *i*<sup>th</sup> element initialized to `static_cast<To>(x[i])` for all *i* in the range of `[0, size())`.

3 *Remarks:* The function shall not participate in overload resolution unless

(3.1) — every possible value of type `U` can be represented with type `To`, and

(3.2) — either

(3.2.1) — `is_simd_v<T>` is false, or

(3.2.2) — `T::size() == simd<U, Abi>::size()` is true.

4 The return type is

(4.1) — `T` if `is_simd_v<T>` is true;

(4.2) — otherwise, `simd<T, Abi>` if `U` is the same type as `T`;

(4.3) otherwise, `simd<T, simd_abi::fixed_size<simd<U, Abi>::size()>>`

```

template<class T, class U, class Abi> see below static_simd_cast(const simd<U, Abi>& x) noexcept;

```

5 Let `To` denote `T::value_type` if `is_simd_v<T>` is true or `T` otherwise.

6 *Returns:* A `simd` object with the *i*<sup>th</sup> element initialized to `static_cast<To>(x[i])` for all *i* in the range of `[0, size())`.

7 *Remarks:* The function shall not participate in overload resolution unless either