

N3196: Omnibus Memory Model and Atomics Paper

ISO/IEC JTC1 SC22 WG21 N3196 = 10-0186 - 2010-11-11

Paul E. McKenney, paulmck@linux.vnet.ibm.com
Mark Batty, mjb220@cl.cam.ac.uk
Clark Nelson, clark.nelson@intel.com
Hans Boehm, hans.boehm@hp.com
Anthony Williams, anthony@justsoftwaresolutions.co.uk
Scott Owens, Scott.Owens@cl.cam.ac.uk
Susmit Sarkar, susmit.sarkar@cl.cam.ac.uk
Peter Sewell, Peter.Sewell@cl.cam.ac.uk
Tjark Weber, tw333@cam.ac.uk
Michael Wong, michaelw@ca.ibm.com
Lawrence Crowl, crowl@google.com
Benjamin Kosnik, bkoz@redhat.com

Introduction

Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber recently analyzed a formalized variant of the C++ memory model, which uncovered a number of potential issues discussed in an email thread entitled “Some more memory model issues from Mark Batty” initiated by Hans Boehm on June 13, 2010, in another email thread entitled “Further C++ concurrency discussion” initiated by Mark Batty on July 28, 2010, and in discussions in the Concurrency Working Group in Rapperswil. This paper summarizes the ensuing discussion, calling out changes that appear uncontroversial and summarizing positions on contended issues. Some of these issues have been captured as national-body comments, and the disposition of any remaining issues is to be determined.

Please note that only those issues related to the memory model and atomics are included in this paper.

More details on the work leading up to this paper may be found [here](#) and [here](#).

This paper is based on [N3125](#) and [N3136](#), and on ensuing discussion in the Concurrency Working Group in Batavia.

Editorial Issues Discussed in June Email Thread

GB 5: Inter-thread-happens-before is not acyclic [Clark]

Sections:

1.9

Comment:

The evaluation of function arguments are now indeterminately sequenced, rather than left completely unspecified, as part of the new language describing the memory model. A clearer example of unspecified behavior should be used here.

Proposal:

Make the editorial change.

Resolution:

Done

Non-Controversial Issues Discussed in June Email Thread

CA 8, GB 10: Inter-thread-happens-before is not acyclic [Clark]

Sections:

1.10p11

Comment:

The following litmus test is generally agreed to be disallowed, however, Batty et al. uncovered a case where the standard does not forbid it, as shown by the following example from their paper:

Thread 0	Thread 1
<code>r1 = x.load(memory_order_consume);</code>	<code>r2 = y.load(memory_order_consume);</code>
<code>y.store(1, memory_order_release);</code>	<code>x.store(1, memory_order_release);</code>

The standard permits the counter-intuitive outcome `x == 1 && y == 1`, an outcome that cannot occur on any hardware platform that we are aware of.

Proposal:

In 1.10p11, add wording prohibiting cycles in the happens-before relation.

Resolution:

Adopt the proposal.

CA 9: Imposed happens-before edges should be synchronizes-with [Clark]

Sections:

1.10p7, 27.2.3p2, 29.3p1, 30.3.1.2p6, 30.3.1.5p7, 30.6.4p7, 30.6.9p5, and 30.6.10.1p23

Comment:

The happens-before relation is not transitive, and so it is not appropriate to specify happens-before for library functions that are intended to impose ordering because happens-before cannot always be extended using a trailing sequenced-before relation. Therefore, synchronizes-with should be used in place of happens-before for this purpose.

Proposal:

In 1.10p7, move the normative language about release synchronizing with acquire to the library (29.3p1) where it belongs, replacing the 1.10p7 use with an example. In the other sections, replace “happens before” with “synchronizes with”.

Resolution:

Adopt the proposal.

CA 11: “Subsequent” in vsse de finition

Comment

Batty et al. propose removing the word “subsequent” from 1.10p12 (presumably instead meaning 1.10p13), stating that this will clarify the definition.

Discussion

This change has interesting consequences. The current wording is as follows, with the word being proposed for removal so marked:

The *visible sequence of side effects* on an atomic object M , with respect to a value computation B of M , is a maximal contiguous sub-sequence of side effects in the modification order of M , where the first side effect is visible with respect to B , and for every ~~subsequent~~ side effect, it is not the case that B happens before it. The value of an atomic object M , as determined by evaluation B , shall be the value stored by some operation in the visible sequence of M with respect to B . Furthermore, if a value computation A of an atomic object M happens before a value computation B of M , and the value computed by A corresponds to the value stored by side effect X , then the

value computed by B shall either equal the value computed by A, or be the value stored by side effect Y, where Y follows X in the modification order of M. [Note: This effectively disallows compiler reordering of atomic operations to a single object, even if both operations are “relaxed” loads. This effectively makes the “cache coherence” guarantee provided by most hardware available to C++ atomic operations. — end note] [Note: The visible sequence depends on the “happens before” relation, which depends on the values observed by loads of atomics, which we are restricting here. The intended reading is that there must exist an association of atomic loads with modifications they observe that, together with suitably chosen modification orders and the “happens before” relation derived as described above, satisfy the resulting constraints as imposed here. — end note]

The effect of the current wording is as follows:

1. The last side-effect in the modification order of M that happens before value computation B is the visible side effect. Call it V.
2. The first side-effect in the modification order of M such that B happen before it will be called I. I and all subsequent side-effects are *not* in the visible sequence of side effects.
3. The word “subsequent” adds the constraint that no side effect in the modification order of M that precedes V can be part of the visible sequence of side effects.

Does some hardware actually operate in this fashion, so that a value computation might return some value preceding the last side-effect in the modification order of M that happens-before that value computation?

Resolution:

Adopt the changes proposed for CA 8.

CA 12: The use of maximal in the definition of r lease sequence [Paul]

Sections:

1.10p6

Comment:

Batty et al. describe an interpretation of 1.10p6 that would only require that release sequences be extended back to the first release operation in a given thread out of a sequence of release operations on a given object.

This interpretation can be considered perverse in light of the wording of 1.10p7, however, the suggested modification is consistent with the intent.

Proposal:

Update 1.10p6 to explicitly call out the head of the release sequence.

Resolution:

Adopt the proposal.

CA 13: Wording of the read-read coherence condition [Paul]

Section:

1.10p13

Comment:

Batty et al. suggest that the following wording from 1.10p13:

Furthermore, if a value computation A of an atomic object M happens before a value computation B of M, and the value computed by A corresponds to the value stored by side effect X, then the value computed by B shall either equal the value computed by A, or be the value stored by side effect Y, where Y follows X in the modification order of M.

be changed to the following:

Furthermore, if a value computation A of an atomic object M happens before a value computation B of M, and A takes its value from the side effect X, then the value computed by B shall either be the value stored by X, or the value stored by a side effect Y, where Y follows X in the modification order of M.

Proposal:

Resolve this with the separate read-read, read-write, write-read, and write-write coherence requirements following 1.10p13 that were motivated by CA 18, CA 19, CA 20, GB 11, and GB 12.

Resolution:

Adopt the proposal.

CA 14: Initialization of atomics

Sections:

1.10p4

Comment:

Batty et al. suggest adding the following non-normative note to 1.10p4:

[Note: There may be non-atomic writes to atomic objects, for example on initialization and re-initialization. — end note]

Discussion

There was some dissatisfaction with this approach expressed in the June email thread. It is quite possible specifying this would be encroaching on the prerogatives of implementors, who are in any case free to perform operations non-atomically when permitted by the as-if rule. Implementors may also perform initializations atomically, again, when permitted by the as-if rule.

Resolution

Adopt the update proposed by US 168.

CA 15: Intra-thread dependency-order ed-before [Paul]

Section:

1.10p9

Comment:

Batty et al. note that, unlike synchronizes-with, the dependency-ordered before relation can operate within a thread. This was not the intent. Instead, intra-thread operations are covered by the rules applying to execution of a single thread.

Proposal:

Update 1.10p9 to make the inter-thread nature of the dependency-ordered before relation explicit.

Resolution:

Adopt the proposal.

CA 18, CA 19, CA 20, GB 11, GB 12: Consequences of lack of well-defined coherence order

Section:

1.10p13

Comment:

These national-body comments called out a number of potential anomalies in the memory model that could result given the current parlous state of the working draft's specification of cache coherence.

Proposal:

Update 1.10p13 and add several subsequent paragraphs to explicitly define requirements for cache coherence.

Resolution:

Adopt the proposal.

CA 22, GB 15: Control Dependencies for Atomics [Paul]

Sections:

N/A

Comment:

“Control dependencies for atomics

Given the examples of compilers interchanging data and control dependencies, and that control dependencies are architecturally respected on Power/ARM for load->store (and on Power for load->load with a relatively cheap isync), we're not sure why carries-a-dependency-to does not include control dependencies between atomics.”

Proposal:

Please clarify.

Discussion

At the time that the memory model was formulated, there was considerable uncertainty as to what architectures respect control dependencies, and to what extent. It appears that this uncertainty is being cleared up, and our hope is that it will be ripe for standardization in a later TR.

Resolution

Not a Defect Future

US 10: Overlapping Atomics [Lawrence]

Sections:

1.10p14

Comment:

The definition of a data race does not take into account two overlapping atomic operations.

Proposal:

Augment the first sentence: The execution of a program contains a data race if it contains two conflicting actions in different threads, at least one of which is not atomic (or both are atomic and operate on overlapping, but not-identical, memory locations), and neither happens before the other.

Discussion

The premise is incorrect; atomic objects may not overlap. The type argument to the `atomic` template must be a trivially-copyable type (29.5.3p1) and atomic objects are not trivially copyable. The atomic types provide no means to obtain a reference to internal members; all atomic operations are copy-in/copy-out. In short, any attempt to generate a pair of atomic variables whose memory overlaps results in undefined behavior.

Resolution

Not a Defect

US 12: N3074 [Paul]

Sections:

1.10p14

Comment:

Adopt [N3074](#).

Proposal:

N3074's proposed change to 1.10p2 has been adopted. However, its proposed change to 1.10p14 has not and needs to be adopted.

Resolution:

Adopt the proposal.

US 168, US 171: Initializing Atomics [Lawrence]

Sections:

29.6p4, 29.6p7

Comment:

The definition of the default constructor needs exposition.

Proposal:

Add a new paragraph: `A::A() = default`; Effects: Leaves the atomic object in an uninitialized state. [Note: These semantics ensure compatibility with C. --end note]

Number

US 171

Sections:

29.6p6

Comment:

The `atomic_init` definition "Non-atomically assigns the value" is not quite correct, as the `atomic_init` purpose is initialization.

Proposal:

Change "Non-atomically assigns the value desired to `*object`." with "Initializes `*object` with value desired". Add the note: "[Note: This function should only be applied to objects that have been default constructed. These semantics ensure compatibility with C. --end note]"

Discussion

Adopt as recommended, but with more clarity.

Resolution

After 29.6p4 [atomics.types.operations], add a new function description for `A::A()`.

Edit 29.6 [atomics.types.operations] paragraph 7 to call out initialization and then move it to just after the paragraph inserted above.

US 38: Generalized Infinite Loops [Clark]

Sections:

1.10p16, 6.5p5

Comment:

The statement that certain infinite loops may be assumed to terminate should also apply to go-to loops and possibly infinite recursion. We expect that compiler analyses that would take advantage of this can often no longer identify the origin of such a loop.

Proposal:

Move paragraph 6.5p5 to follow 1.10p16, but generalize it to all code, not just for loops.

Resolution:

Adopt the proposal.

GB 8, US 9, US 11: Mutexes versus Locks [Lawrence]

Sections:

1.10p4, 1.10p7

Comment:

The text says that the library "provides ... operations on locks". It should say "operations on mutexes", since it is the mutexes that provide the synchronization. A lock is just an abstract concept (though the library types `unique_lock` and `lock_guard` model ownership of locks) and as such cannot have operations performed on it. This mistake is carried through in the notes in that paragraph and in 1.10p7.

Proposal:

Change 1.10p4 as follows:

"The library defines a number of atomic operations (Clause 29) and operations on mutexes (Clause 30) that are specially identified as synchronization operations. These operations play a special role in making assignments in one thread visible to another. A synchronization operation on one or more memory locations is either a consume operation, an acquire operation, a release operation, or both an acquire and release operation. A synchronization operation without an associated memory location is a fence and can be either an acquire fence, a release fence, or both an acquire and release fence. In addition, there are relaxed atomic operations, which are not synchronization operations, and atomic read-modify-write operations, which have special characteristics. [Note: For example, a call that acquires a lock on a mutex will perform an acquire operation on the locations comprising the mutex. Correspondingly, a call that releases the same lock will perform a release operation on those same locations. Informally, performing a release operation on A forces prior side effects on other memory locations to become visible to other threads that later perform a consume or an acquire operation on A. "Relaxed" atomic operations are not synchronization operations even though, like synchronization operations, they cannot contribute to data races. -- end note]"

Change 1.10p7 as follows:

"Certain library calls synchronize with other library calls performed by another thread. In particular, an atomic operation A that performs a release operation on an atomic object M synchronizes with an atomic operation B that performs an acquire operation on M and reads a value written by any side effect in the release sequence headed by A. [Note: Except in the specified cases, reading a later value does not necessarily ensure visibility as described below. Such a requirement would sometimes interfere with efficient implementation. -- end note] [Note: The specifications of the synchronization operations define when one reads the value written by another. For atomic objects, the definition is clear. All operations on a given mutex occur in a single total order. Each lock acquisition "reads the value written" by the last lock release on the same mutex. -- end note]"

Number

US 9

Sections:

1.10p4

Comment:

The "operations on locks" do not provide synchronization, as locks are defined in Clause 30.

Proposal:

Change "operations on locks" to "locking operations". (Covered by GB 8.)

Number

US 11

Sections:

1.10p7

Comment:

There is some confusion between locks and mutexes.

Proposal:

Change "lock" when used as a noun to "mutex". (Covered by GB 8.)

Discussion

Adopt the wording of GB 8, but with additional use of "mutex" for clarity.

Resolution

Adopt the proposal

CH 2: Observable Behavior of Atomics [Lawrence]

Sections:

1.9 and 1.10

Comment:

It's not clear whether relaxed atomic operations are observable behaviour.

Proposal:

Clarify it.

Discussion

Normatively, the behavior is well-defined by 1.9p8. If the atomic object is volatile, then all operations on it are observable, otherwise not. Note that "observable" means "observable outside of the program."

Resolution

Not a Defect.

Controversial Issues Discussed in June Email Thread

CA 17: 1.10p12 phrasing

The last note of 1.10p12 refers to data races "as defined here". Batty et al. recommend that this change to "as defined below". Given that data races are defined in 1.10p14, it is easy to argue for "below", however, it is equally easy to argue that the scope of "here" is the whole of 1.10.

Wording

1.10p4

National-body comment: GB 8, US 9, US 11

Changes:

The library defines a number of atomic operations (Clause 29) and operations on ~~locks~~ mutexes (Clause 30) that are specially identified as synchronization operations. These operations play a special role in making assignments in one thread visible to another. A synchronization operation on one or more memory locations is either a consume operation, an acquire operation, a release operation, or both an acquire and release operation. A synchronization operation without an associated memory location is a fence and can be either an acquire fence, a release fence, or both an acquire and release fence. In addition, there are relaxed atomic operations, which are not synchronization operations, and atomic read-modify-write operations, which have special

characteristics. [Note: For example, a call that acquires a lock mutex will perform an acquire operation on the locations comprising the lock mutex. Correspondingly, a call that releases the same lock mutex will perform a release operation on those same locations. Informally, performing a release operation on *A* forces prior side effects on other memory locations to become visible to other threads that later perform a consume or an acquire operation on *A*. "Relaxed" atomic operations are not synchronization operations even though, like synchronization operations, they cannot contribute to data races. —end note]

1.10p6

National-body comment: CA 12

Changes:

A release sequence from a release operation *A* on an atomic object *M* is a maximal contiguous sub-sequence of side effects in the modification order of *M*, where the first operation is a release *A*, and every subsequent operation

- is performed by the same thread that performed the release *A*, or
- is an atomic read-modify-write operation.

1.10p7

National-body comments: CA 9, GB 8, US 9, US 11

Changes:

Certain library calls *synchronize with* other library calls performed by another thread. ~~In particular, an atomic operation *A* that performs a release operation on an atomic object *M* synchronizes with an atomic operation *B* that performs an acquire operation on *M* and reads a value written by any side effect in the release sequence headed by *A*.~~ [Example: An atomic store-release synchronizes with a load-acquire that takes its value from the store. (29.3 atomics.order) — end example] [Note: Except in the specified cases, reading a later value does not necessarily ensure visibility as described below. Such a requirement would sometimes interfere with efficient implementation. —end note] [Note: The specifications of the synchronization operations define when one reads the value written by another. For atomic objects, the definition is clear. All operations on a given lock mutex occur in a single total order. Each mutex lock acquisition “reads the value written” by the last mutex lock release on the same mutex. —end note]

1.10p9

National-body comment: CA 15

Changes:

An evaluation *A* is dependency-ordered before an evaluation *B* if

- *A* performs a release operation on an atomic object *M*, and on another thread, *B* performs a consume operation on *M* and reads a value written by any side effect in the release sequence headed by *A*, or
- for some evaluation *X*, *A* is dependency-ordered before *X* and *X* carries a dependency to *B*.

1.10p11

National-body comments: CA 8, GB 10

Changes:

An evaluation *A* happens before an evaluation *B* if:

- *A* is sequenced before *B*, or
- *A* inter-thread happens before *B*.

The implementation shall ensure that no program execution demonstrates a cycle in the "happens before" relation. [Note: This cycle would otherwise be possible only through the use of consume operations. — end note]

1.10.p13

National-body comments: CA 11, CA 13, CA 18, CA 19, CA 20, GB 11, GB 12

Changes:

The *visible sequence of side effects* on an atomic object *M*, with respect to a value computation *B* of *M*, is a maximal contiguous sub-sequence of side effects in the modification order of *M*, where the first side effect is visible with respect to *B*, and for every ~~subsequent~~ side effect, it is not the case that *B* happens before it. The value of an atomic object *M*, as determined by evaluation *B*, shall be the value stored by some operation in the a visible sequence of *M* with respect to *B*. [Note: It can be shown that the visible sequence of side effects of a value computation is unique given the coherence requirements below. — end note] Furthermore, if a value computation *A* of an atomic object *M* happens before a value computation *B* of *M*, and the value computed by *A* corresponds to the value stored by side effect *X*, then the value computed by *B* shall either equal the value computed by *A*, or be the value stored by side effect *Y*, where *Y* follows *X* in the modification order of *M*. ~~[Note: This effectively disallows compiler reordering of atomic operations to a single object, even if both operations are "relaxed" loads. This effectively makes the "cache coherence" guarantee provided by most hardware available to C++ atomic operations. — end note] [Note: The visible sequence depends on the "happens before" relation, which depends on the values observed by loads of atomics, which is restricted here. The intended reading is that there must exist an association of atomic loads with modifications they observe that, together with suitably chosen modification orders and the "happens before" relation derived as described above, satisfy the resulting constraints as imposed here. — end note]~~

New Paragraphs Following 1.10p13

National-body comments: CA 13, CA 18, CA 19, CA 20, GB 11, GB 12

Changes:

If an operation *A* that modifies an atomic object *M* happens before an operation *B* that modifies *M*, then *A* shall be earlier than *B* in the modification order of *M*. [Note: This requirement is known as write-write coherence. — end note]

If a value computation *A* of an atomic object *M* happens before a value computation *B* of *M*, and *A* takes its value from a side effect *X* on *M*, then the value computed by *B* shall either be the value stored by *X*, or the value stored by a side effect *Y* on *M*, where *Y* follows *X* in the modification order of *M*. [Note: This requirement is known as read-read coherence. — end note]

If a value computation *A* of an atomic object *M* happens before an operation *B* on *M*, then *A* shall take its value from a side effect *X* on *M*, where *X* precedes *B* in the modification order of *M*. [Note: This requirement is known as read-write coherence. — end note]

If a side effect *X* on an atomic object *M* happens before a value computation *B* of *M*, then the evaluation *B* shall take its value from *X* or from a side effect *Y* that follows *X* in the modification

order of M. [Note: This requirement is known as write-read coherence. — end note]

[Note: These four coherence requirements effectively disallow compiler reordering of atomic operations to a single object, even if both operations are relaxed loads. This effectively makes the “cache coherence” guarantee provided by most hardware available to C++ atomic operations. — end note]

[Note: The visible sequence of side effects depends on the “happens before” relation, which depends on the values observed by loads of atomics, which we are restricting here. The intended reading is that there must exist an association of atomic loads with modifications they observe that, together with suitably chosen modification orders and the “happens before” relation derived as described above, satisfy the resulting constraints as imposed here. — end note]

1.10p14

National-body comment: US 12

Changes:

The execution of a program contains a *data race* if it contains two conflicting actions in different threads, at least one of which is not atomic, and neither happens before the other. Any such data race results in undefined behavior. [Note: It can be shown that programs that correctly use ~~simple locks mutexes and memory_order_seq_cst operations~~ to prevent all data races and ~~that~~ use no other synchronization operations behave as ~~though the executions of~~ if the operations executed by their constituent threads were simply interleaved, with each ~~observed value-value computation~~ of an object being taken from the last value assigned ~~last side effect on that object~~ in that interleaving. This is normally referred to as “sequential consistency”. However, this applies only to ~~data-race-free~~ programs, and ~~data-race-free~~ programs cannot observe most program transformations that do not change single-threaded program semantics. In fact, most single-threaded program transformations continue to be allowed, since any program that behaves differently as a result must perform an undefined operation. — end note]

Paragraph Following 1.10p16

National-body comment: US 38

Changes:

The implementation is allowed to assume that any thread will eventually do one of the following:

- terminate,
- make a call to a library I/O function,
- access or modify a volatile object, or
- perform a synchronization operation or an atomic operation.

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

6.5p5

National-body comment: US 38

Deletion:

~~A loop that, outside of the *for-init-statement* in the case of a *for*-statement,~~

- ~~makes no calls to library I/O functions, and~~

- ~~does not access or modify volatile objects, and~~
- ~~performs no synchronization operations (1.10) or atomic operations (Clause 29)~~

~~may be assumed by the implementation to terminate. [Note: This is intended to allow compiler transformations, such as removal of empty loops, even when termination cannot be proven. — end note]~~

27.2.3p2

National-body comment: CA 9

Changes:

If one thread makes a library call *a* that writes a value to a stream and, as a result, another thread reads this value from the stream through a library call *b* such that this does not result in a data race, then *a*'s write happens before *b*'s read.

New Paragraph Following 29.3p1

National-body comment: CA 9

Changes:

An atomic operation *A* that performs a release operation on an atomic object *M* synchronizes with an atomic operation *B* that performs an acquire operation on *M* and takes its value from any side effect in the release sequence headed by *A*.

New Paragraph Following 29.6p4

National-body comment: US 168, US 171

Changes:

A::*A*() = default;

Effects: Leaves the atomic object in an uninitialized state. [Note: These semantics ensure compatibility with C. —end note]

29.6p7

National-body comment: US 168, US 171

Change as follow, then move `atomic_init()` to follow the new paragraph added after 29.6p4:

Effects: Non-atomically ~~assigns the value desired to *object-~~initializes *object with value desired. This function shall only be applied to objects that have been default constructed, and then only once. [Note: These semantics ensure compatibility with C. —end note] [Note: Concurrent access from another thread, even via an atomic operation, constitutes a data race. —end note]

30.3.1.2p6

National-body comment: CA 9

Changes:

Synchronization: The completion of the invocation of the constructor ~~happens before~~ synchronizes with the beginning of the invocation of the copy of *f*.

30.3.1.5p7

National-body comment: CA 9

Changes:

Synchronization: The completion of the thread represented by **this* ~~happens before~~ synchronizes with (1.10) the corresponding successful *join()* returns. [Note: Operations on **this* are not synchronized. — end note]

30.6.4p7

National-body comment: CA 9

Changes:

Calls to functions that successfully set the stored result of an associated asynchronous state synchronize with (1.10) calls to functions successfully detecting the ready state resulting from that setting. The storage of the result (whether normal or exceptional) into the associated asynchronous state ~~happens before~~ synchronizes with (1.10) that state is set to ready the successful return from a call to a waiting function on the associated asynchronous state.

30.6.9p5

National-body comment: CA 9

Changes:

Synchronization: the invocation of *async* ~~happens before~~ synchronizes with (1.10) the invocation of *f*. [Note: this statement applies even when the corresponding *future* object is moved to another thread. — end note] If the invocation is not deferred, a call to a waiting function on an asynchronous return object that shares the associated asynchronous state created by this *async* call shall block until the associated thread has completed. If the invocation is not deferred, the *join()* on the created thread ~~happens before~~ synchronizes with (1.10) the first function that successfully detects the ready status of the associated asynchronous state returns or before the function that gives up the last reference to the associated asynchronous state returns, whichever happens first. If the invocation is deferred, the completion of the invocation of the deferred function ~~happens before~~ synchronizes with the the successful return from a call to a waiting function on the associated asynchronous state. ~~calls to the waiting functions return.~~

30.6.10.1p23

National-body comment: CA 9

Changes:

Synchronization: a successful call to *operator()* synchronizes with (1.10) a call to any member function of a *future*, *shared_future*, or *atomic_future* object that shares the associated asynchronous state of **this*. The completion of the invocation of the stored task and the storage of the result (whether normal or exceptional) into the associated asynchronous state ~~happens before~~ synchronizes with (1.10) the successful return from any member function that detects that the state is set to ready. [Note: *operator()* synchronizes and serializes with other functions through the associated asynchronous state. — end note]