

Document No: WG21 N3289  
INCITS / PL22.16 11-0059  
Date: 2011-04-07  
Project: Programming Language C++  
References: WG21 N3092, SC 22 N4512: ISO/IEC FCD 14882  
Reply to: Barry Hedquist  
INCITS/PL22.16 IR  
Email: [beh@peren.com](mailto:beh@peren.com)

## **ISO/IEC FCD 14882, C++0X Responses to National Body Comments**

Attached are the ISO/IEC JTC1 SC22/WG21 responses to the National Body Comments submitted to JTC1 SC22 in response to the SC22 Letter Ballot for ISO/IEC FCD 14882, Final Committee Draft of the revision of ISO/IEC 14882:2003, aka C++0X.

Comments that were originally submitted without numbering were numbered manually in the exact order of the NB's official ballot response. The comments were then organized per the hierarchy of the balloted document, SC22 N4512 (WG21 N3092). No editing of any kind was done on any of the comments. The responses reflect the consensus position of SC22/WG21 for each comment.

Where applicable, pointers to publically available documents are used to provide additional details for the response.

CWG xxx – refers to Core Working Group Issue xxx  
LWG xxx – refers to Library Working Group Issue xxx  
Nxxxx – refers to SC22/WG21 paper Nxxxx

The following individuals are to be commended for their outstanding work in compiling these responses:

William 'Mike' Miller - Edison Design Group  
Beman Dawes - BoostPro  
Pete Becker - Project Editor, Roundhouse Consulting, Ltd.  
Alisdair Meredith - LWG Chair, Bloomberg

# ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2  | (3)   | 4                                       | 5   | (6)   | (7)   |
|-----------------|--|---|---|---|---|---|
| MB <sup>1</sup> | Clause No./<br>Subclause No./<br>Annex<br>(e.g. 3.1) | Paragraph/<br>Figure/Table/<br>Note<br>(e.g. Table 1) | Type<br>of<br>com-<br>ment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations<br>on each comment submitted                     |
| ITTF<br>01      | General  |   | ed                                      | The ISO/IEC headers and footers should appear at the top and bottom of every page.  | Insert the ISO/IEC headers and footers at the top and bottom of every page.   | ACCEPTED  |
| US<br>01        | 1-30   |   | ge                                      | It appears that the C committee could possibly make some changes for C1X that we may still want to follow for C++ to avoid gratuitous incompatibilities.                      | Make any appropriate changes to avoid unnecessary incompatibilities with C1X resulting from changes to the WG14 C standard draft.   | ACCEPTED  |
| US<br>02        | 1-30   |   | ge                                      | The active issues identified in the CWG and LWG issues lists as if the date that the FCD was published (N3083 and N3087) must be addressed and appropriate action taken.      | Appropriate action would include making changes to the FCD, identifying an issue as not requiring a change to the FCD, or deferring an issue to a later draft or a later standard.  | ACCEPTED  |
| DE<br>1         | 1 through 15   |   | te                                      | Consider applying the resolutions of the active core issues in Ready status (see WG21 N3083).   |   | ACCEPTED  |
| CH<br>1         | all  |   | ge/te                                   | The issues on the issues lists (WG21 N3083 and N3087) shall be addressed before the standard becomes final.   |   | ACCEPTED  |
| US<br>03        | 1 - 29   |   | te                                      | The threading model does not make basic guarantees needed to write correct programs. We should not repeat the error that POSIX made in early standards (and later corrected). | Add requirements that all no-blocked threads will (however slowly) make progress and that all visible side effects will (eventually) be seen by other threads. Possibly use the word "should" if an absolute requirement is impossible. | ACCEPTED<br>See paper <a href="#">N3209</a>                               |
| US<br>04        | all  | all   | ed                                      | Many identifiers are hyphenated and broken across line boundaries. As a consequence, searching for the identifier fails.  | Protect all identifiers against hyphenation.  | ACCEPTED  |
| US<br>05        | all  | all   | ed                                      | The word "object" often copies as "ob ject".  |   | REJECTED<br>Some PDF viewers do this. Don't know why, or how to avoid it. |
| US<br>06        | various  | various   | ed                                      | ~ (U+007E) is sometimes replaced with ~ (U+223C), causing cut and paste from the standard to fail, notably in 2.14.5.   | Use U+007E consistently   | ACCEPTED  |
| US<br>07        | various  | various   | ed                                      | ' (U+0027) is consistently replaced with ' (U+2019), causing cut and paste to fail. This is also an issue with  | Use U+0027 consistently in code samples (i.e. monospace font)   | ACCEPTED  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | the 1998 standard.  |  |   |
| GB 1            | 1.1   | 2  | Ed                           | The C99 standard supports inline functions, so this should not be listed as a distinguishing feature of C++.  | strike "inline functions" from the list of C++ 'additional facilities'   | ACCEPTED  |
| GB 2            | 1.2   |  | Ge                           | In [intro.refs] the spec references ISO/IEC 9945:2003 even though a later revision, ISO/IEC 9945:2008 has already been released:<br><a href="http://www.iso.org/iso/catalogue_detail.htm?csnumber=50516">http://www.iso.org/iso/catalogue_detail.htm?csnumber=50516</a> | The section should be updated to reference the latest version.<br>In addition, since POSIX is a registered trademark of the IEEE, the spec should use the registered trademark (R) symbol wherever it references it. Alternatively, it can refer to ISO/IEC 9945.  | REJECTED<br>The current regex specification reflects the 2003 version of POSIX; changing the reference to the newer revision could have normative impact. There was no consensus for this change. |
| ITTF 02         | 1.2   |  | ed                           | The introductory text to the Normative references is not correct.   | Delete the current introductory text to the Normative references and replace with the following:<br><br>"The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies." | ACCEPTED  |
| ITTF 03         | 1.3   |  | ed                           | The title to the subclause does not accurately reflect the content.   | Change the title to "Terms and definitions".   | ACCEPTED  |
| ITTF 04         | 1.3   |  | ed                           | 3.1 of the ISO/IEC Directives, Part 2, states that the following introductory wording shall be used where all terms and definitions are given the document itself:<br><br>"For the purposes of this document, the following terms and definitions apply."               | Change the introductory text to:<br><br>"For the purposes of this document, the following terms and definitions apply."  | ACCEPTED  |
| ITTF            | 1.3   |  | ed                           | D.1.5.3 of the ISO/IEC Directives, Part 2 states that the form of a definition shall be such that it can replace the  | Delete the definite or indefinite article at the   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| 05              |   |  |                              | term in context.   | beginning of each definition.<br>Redraft definitions 1.3.11 and 1.3.13 so that they can replace the term in context (i.e. they should not be more than one sentence). |  |
| ITTF 06         | 1.3   |  | ed                           | D.3.1 of the ISO/IEC Directives, Part 2, states that the definition shall not be followed by a full stop.  | Remove the full stops at the end of the definitions.  | ACCEPTED   |
| ITTF 07         | 1.3   |  | ed                           | D.3.9 of the ISO/IEC Directives, Part 2, provides examples on how to present examples and notes to terms and definitions. The examples and notes to the terms and definitions are not presented in accordance with D.3.9 of the ISO/IEC Directives, Part 2.        | Redraft the notes and examples to the terms and definitions in accordance with D.3.9 of the ISO/IEC Directives, Part 2.   | REJECTED<br><br>The current style has been accepted for the previous two C++ Standards. It is designed for large, complex documents; the ISO rules are impractical in a document of this size. |
| GB 3            | 1.3   | 1  | Ed                           | The library stretches out to clause 30 now, and the terms in 17.3 should cover them too.   | Update reference to clause 27 to say clause 30.   | ACCEPTED   |
| JP 15           | 1.3   | 1  | E                            | There is a description, "17.3 defines additional terms that are used only in Clauses 17 through 27 and Annex D.", but the terms defined in 17.3 are also used in Clauses 28 through 30, which are added recently. So the scope should be expanded to include them. | 17.3 defines additional terms that are used only in Clauses 17 through 30 and Annex D.  | ACCEPTED   |
| GB 4            | 1.3.10                                      |  | Ed                           | The phrase "catch clause" in the 2003 standard has become "catch Clause"   | Change back the "catch clause"  | ACCEPTED   |
| RU 1            | 1.7   | p.5, line 5 from end                         | ed                           | Reference missed   | Insert reference "(2.3)" after "basic execution character set"  | ACCEPTED   |
| GB 5            | 1.9   | 3  | Ed                           | The evaluation of function arguments are now indeterminately sequenced, rather than left completely unspecified, as part of the new language describing the  | [Need to identify a better example to propose]  | ACCEPTED   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | memory model. A clearer example of unspecified behavior should be used here.  |  |   |
| GB 6            | 1.9   | 4  | Ed                           | There are core issues surrounding the undefined behavior of dereferencing a null pointer. It appears the intent is that dereferencing *is* well defined, but using the result of the dereference will yield undefined behavior. This topic is too confused to be the reference example of undefined behavior, or should be stated more precisely if it is to be retained. | [Identify a better example of undefined behavior to propose]   | ACCEPTED  |
| CH 2            | 1.9 and 1.10                                |  | te                           | It's not clear whether relaxed atomic operations are observable behaviour.  | Clarify it.  | REJECTED<br><br>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."<br><br>See paper <a href="#">N3196</a> |
| GB 7            | 1.9.6                                       | p6   | Te                           | From a naive, first-time reader's point of view, 1.9.6 made me double take, as it seemed it could imply any signal would leave the program in an unspecified state after completing. I believe I understand the intent, but to clarify it, I'd change it to make it clear that the unspecified state applies only for the duration of the signal handler.                 | Change:<br><br>6 When the processing of the abstract machine is interrupted by receipt of a signal, the values of objects which are neither<br>— of type volatile std::sig_atomic_t nor<br>— lock-free atomic objects (29.4)<br>are unspecified, and the value of any object not in either of these two categories that is modified by the handler becomes undefined.<br><br>to:<br><br>6 When the processing of the abstract machine is | ACCEPTED<br><br>The reference should be 1.9 paragraph 6, not 1.9.6 (which does not exist).  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              |   | interrupted by receipt of a signal, the values of objects which are neither<br>— of type volatile std::sig_atomic_t nor<br>— lock-free atomic objects (29.4)<br>are unspecified <b>for the duration of the signal handler</b> , and the value of any object not in either of these two categories that is modified by the handler becomes undefined. |  |
| US 08           | 1.9   | footnote 7                                   | te                           | The footnote "Overloaded operators are never assumed to be associative or commutative." is either meaningless or overly restrictive.  | Change the footnote to "Overloaded operators are assumed to be non-associative and non-commutative until proven otherwise."  | REJECTED<br><br>The statement involved is non-normative and is, in general, correct. An implementation can treat overloaded operators as associative or commutative only under the "as-if" rule, so the statement is clear enough. |
| CA 23           | 1.10, 29                                    | 1.10, 29                                     | Te                           | C1x has added new atomics<br>C1x has added new atomics syntax, and in some cases new semantics and operations. C++0x needs to consider aligning with the new C1x atomics                    | Add back compatibility between C++0x and C1x atomics   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a>   |
| CA 14           | 1.10p4                                      | 1.10p4                                       | ed                           | Initialisation of atomics<br>Add the following to 1.10p4:<br><br>[ Note: There may be non-atomic writes to atomic objects, for example on initialization and re-initialization. - end note] | Add the following to 1.10p4:<br>[ Note: There may be non-atomic writes to atomic objects, for example on initialization and reinitialization. - end note]  | ACCEPTED with MODIFICATIONS<br><br>The suggested update from US 168 was adopted.<br><br>See paper <a href="#">N3196</a>  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>Rationale: We believe the intent is that for any atomic there is a distinguished initialisation write, but that this need not happens-before all the other operations on that atomic - specifically so that the initialisation write might be non-atomic and hence give rise to a data race, and hence undefined behaviour, in examples such as this (from Hans):</p> <pre> atomic&lt; atomic&lt;int&gt; * &gt; p f() { atomic&lt;int&gt;x;            W_na x   p.store(&amp;x,mo_rlx);     W_rlx p=&amp;x }                     </pre> <p>(where na is nonatomic and rlx is relaxed). We suspect also that no other mixed atomic/nonatomic access to the same location is intended to be permitted. The possibility of non-atomic writes on atomic objects is not mentioned in 1.10, and (before talking with Hans) we didn't realise it was intended, so we suggest adding the note above to clarify things.</p> |   |  |
| CA 12           | 1.10p6                                      | 1.10p6                                       | te                           | <p>The use of maximal in the definition of release sequence (proposed edit seems reasonable to Clark)</p> <p>We suggest that 1.10p6 be changed to:</p> <p>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, and every subsequent operation</p>   | <p>We suggest that 1.10p6 be changed to:</p> <p>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, and every subsequent operation</p> | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>- is performed by the same thread that performed the release, or</p> <p>- is an atomic read-modify-write operation.</p> <p>Rationale: The current wording of the standard suggests that release sequences are maximal with respect to sequence inclusion, i.e. that if there are two release operations in the modification order,</p> <pre> mod      mod rel1-----&gt;rel2-----&gt;w                     </pre> <p>then [rel1;rel2;w] is the only release sequence, as the other candidate [rel2;w] is included in it. This interpretation precludes synchronizing with releases which have other releases sequenced-before them. We believe that the intention is actually to define the maximal release sequence from a particular release operation, which would admit both [rel1;rel2;w] and [rel2;w].</p> | <p>- is performed by the same thread that performed the release, or</p> <p>- is an atomic read-modify-write operation.</p> |  |
| US 09           | 1.10  | para 4                                       | te                           | The "operations on locks" do not provide synchronization, as locks are defined in Clause 30.   | Change "operations on locks" to "locking operations".  | ACCEPTED<br>See paper <a href="#">N3196</a>        |
| CA 20           | 1.10p1                                      | 1.10p1                                       | Te                           | <p>Reading from the last element in a vsse?</p> <p>Paul wrote:</p> <p>&gt; If every element in a vsse happens-before a given value</p> <p>&gt; computation, then that value computation must return</p>  | Please clarify.  | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5  | (6)                       | (7)  |
|-----------------|---|--|------------------------------|--|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>&gt; the value stored by the last element in the vsse.</p> <p>We wrote:</p> <p>We're not sure about that. Consider the following, with two relaxed writes to x on one thread that are not sequenced-before related to each other (eg in different arguments to the same function), but are followed by a release/acquire on a different variable y to another thread that then reads x. We think the final read (e) could read from either (a) or (b), regardless of how (a) and (b) are related in modification order.</p> <pre> a:Wrlx x=1      b:Wrlx x=2   \            /   sb          /sb    \          /    c:Wrel y-----                         \sw                         d:Racq y                            sb                           e:Rrlx x=?                     </pre> <p> </p> <p>Paul&gt; In that case IIRC, the standard does not specify</p> <p>Paul&gt; the order, but the code will be generated in some</p> <p>Paul&gt; order, and that arbitrary choice on the part of the</p> <p>Paul&gt; compiler will determine the modification order.</p> <p>We agree that in a normal implementation (eg where the</p> |                           |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)                       | (7)  |
|-----------------|---|--|------------------------------|--|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>argument evaluations are not spawned off to new threads - is that intended to be forbidden?), the two writes will indeed be ordered according to the generated-code order (and before the release fence), and hardware coherence will ensure that (e) reads the later one.</p> <p>But in the draft standard as written, that execution is allowed - the draft doesn't currently impose that aspect of coherence. To make the example more concrete, if there were another thread with</p> <pre>c --sw--&gt; f:Racq y --sb--&gt; g:Rrlx x</pre> <p>then e and g could read different values.</p> <p>Paul notes:</p> <ul style="list-style-type: none"> <li>&gt; But 1.10p1 says:</li> <li>&gt;</li> <li>&gt; A thread of execution (also known as a thread) is a             <ul style="list-style-type: none"> <li>&gt; single flow of control within a program, including</li> <li>&gt; the initial invocation of a specific top-level</li> <li>&gt; function, and recursively including every function</li> <li>&gt; invocation subsequently executed by the thread.</li> <li>&gt;</li> <li>&gt; This excludes the possibility of the implementation                 <ul style="list-style-type: none"> <li>&gt; spawning off a new thread -unless- the implementation</li> <li>&gt; can make things appear as if there was only one thread.</li> </ul> </li> </ul> </li> </ul> |                           |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)                       | (7)  |
|-----------------|---|--|------------------------------|--|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>&gt; From this viewpoint, your example shows just how<br/>                     &gt; careful an implementation must be if it is to fully<br/>                     &gt; comply with this as-if rule.</p> <p>We replied</p> <p>&gt;ok, thanks</p> <p>to this, but in fact the situation is still unclear.</p> <p>1.10p1 does indeed rule out the hypothetical implementation that we mentioned, but even if (a) and (b) would be ordered by any reasonable implementation, in terms of the concepts of the standard, that doesn't introduce a sequenced-before edge between (a) and (b).</p> <p>It seems that Paul is assuming the individual memory accesses in function arguments are indeterminately sequenced rather than unsequenced?</p> |                           |  |
| CA 19           | 1.10p5<br>1.10p13                           | 1.10p5<br>1.10p13                            | Te                           | <p>Alternative definition of the value read by an atomic operation</p> <p>Here's an interesting example involving a release/consume pair. We believe that in a direct implementation on hardware, this would be forbidden by coherence, but that the current text allows it. We don't know whether it should be allowed or not.</p> <pre>                     hb                     do                     rf Wx_release  ----&gt; Rx_consume                     ^                     \                      sb, hb                     </pre>  | Please clarify.           | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)                       | (7)  |
|-----------------|---|--|------------------------------|---|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p style="text-align: center;">mo \                    v<br/>                              --- Wx_release</p> <p>Paul claims this is forbidden by 1.10p5, but we don't see how that can be the case. We don't see much room in 1.10p5 for any other interpretation - it says:</p> <p>- "All modifications to a particular atomic object M occur in some particular total order, called the modification order of M"</p> <p>- "If A and B are modifications of an atomic object M and A happens before (as defined below) B, then A shall precede B in the modification order of M, which is defined below."</p> <p>Both of which seem very clear. The only wiggle room is in the Note</p> <p>- "[Note: This states that the modification orders must respect the "happens before" relationship]"</p> <p>We took that "must respect" to be a gloss rather than to add any additional constraint.</p> <p>Earlier we suggested a change, to the constraint on the value read by an atomic operation, that would forbid this example:</p> <p>The standard introduces visible side effects, which are used first to define the values read by non-atomic operations. They are then re-used to constrain the value read by atomic operations: 1.10p13 says that an atomic operation must read from somewhere in "the" visible sequence of side effects, which must start from *a* visible side effect, i.e. a side effect that (a) happens before the read, and (b) is not happens-before-hidden. We suspect that this re-use of the notion of visible side effect may be a</p> |                           |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |
|                 |   |  |                              | <p>drafting artifact, in which case one might remove the requirement that there is a vse for atomics, and replacing the first two sentences of 1.10p13 by</p> <p>"An atomic operation must read from somewhere in the modification order that is not happens-before-hidden and does not follow (in modification order) any side effect that happens-after the read."</p> <p>Now we're not sure how this would fit in with initialisation and reading of indeterminate values; we need to think about it more.</p> |   |   |
| CA 22           | 1.10p8                                      | 1.10p8                                       | Te                           | <p>Control dependencies for atomics</p> <p>Given the examples of compilers interchanging data and control dependencies, and that control dependencies are architecturally respected on Power/ARM for load-&gt;store (and on Power for load-&gt;load with a relatively cheap isync), we're not sure why carries-a-dependency-to does not include control dependencies between atomics.</p>   | Please clarify.   | <p>REJECTED</p> <p>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.</p> <p>See paper <a href="#">N3196</a></p> |
| CA 15           | 1.10p9                                      | 1.10p9                                       | Ed                           | <p>Intra-thread dependency-ordered-before</p> <p>The current draft has release/acquire synchronize-with edges only between a release on one thread and an acquire on a *different* thread, whereas the definition of dependency-ordered-before permits the release and consume to be on the same thread; it seems odd to permit</p>   | <p>We suggest changing the definition of dependency-ordered-before in 1.10p9 to the following:</p> <p>An evaluation A is dependency-ordered before an evaluation B if</p> | <p>ACCEPTED</p> <p>See paper <a href="#">N3196</a></p>  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>the latter. (At the moment function arguments can't race or sync with each other, but they can be dependency ordered before each other.)</p> <p>We don't currently have an example in which this makes a real difference, but for symmetry could suggest changing the definition of dependency-ordered-before in 1.10p9 to the following:</p> <p>An evaluation A is dependency-ordered before an evaluation B if</p> <ul style="list-style-type: none"> <li>- 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,</li> <li>or</li> <li>- for some evaluation X, A is dependency-ordered before X and X carries a dependency to B.</li> </ul> | <ul style="list-style-type: none"> <li>- 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</li> <li>- for some evaluation X, A is dependency-ordered before X and X carries a dependency to B.</li> </ul> |  |
| CA 11           | 1.10p12                                     | 1.10p12                                      | te                           | <p>"Subsequent" in vsse definition</p> <p>Remove the word "subsequent" from the definition of visible sequence of side effects in 1.10p12.</p> <p>(as suggested by Hans)</p> <p>Rationale: if every element in a vsse happens-before a read, the read should not take the value of the visible side effect.</p>  | Remove the word "subsequent" from the definition of visible sequence of side effects in 1.10p12.  | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
| CA 17           | 1.10p12                                     | 1.10p12                                      | Ed                           | 1.10p12 phrasing<br>1.10p12 last note:<br><br>"...as defined here..." should be<br><br>"...as defined below..."   | 1.10p12 last note:<br>"...as defined here..." should be<br><br>"...as defined below..."  | REJECTED<br><br>The reference really should be to "data races as defined in this International Standard", because the note compares this definition with the generally-understood meaning of data races in sequentially-consistent executions. That's far too stilted, and "as defined here" seems like a reasonable way to phrase it in the less formal context of a note. |
| CA 13           | 1.10p13                                     | 1.10p13                                      | ed                           | Wording of the read-read coherence condition<br>In 1.10p13 a coherence condition is stated on the values of atomic reads:<br><br>"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."<br><br>We suggest that this be replaced with the following: | In 1.10p13 a coherence condition is stated on the values of atomic reads:<br>"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, | ACCEPTED<br><br>See paper <a href="#">N3196</a>   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted     |
|                 |   |  |                              | <p>"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."</p> <p>Rationale: The words "corresponds to" are not used elsewhere in the standard, as far as we can see, and it is unclear whether they have a special meaning here. In addition taking the value of the read B from the value read by A seems unnecessarily indirect. B could take its value from X instead.</p> | <p>where Y follows X in the modification order of M."</p> <p>We suggest that this be replaced with the following:</p> <p>"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."</p> |  |
| CA 18           | 1.10p13                                     | 1.10p13                                      | Te                           | <p>Non-unique visible sequences of side effects and happens-before ordering</p> <p>In 1.10p13, replace</p> <p>"The visible sequence of side effects on..." by</p> <p>"A visible sequence of side effects on..."</p> <p>and</p>  | <p>In 1.10p13, replace</p> <p>"The visible sequence of side effects on..."</p> <p>by</p> <p>"A visible sequence of side effects</p>   | <p>ACCEPTED</p> <p>See paper <a href="#">N3196</a></p> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>"in the visible sequence of M with respect to B" by</p> <p>"in a visible sequence of M with respect to B"</p> <p>Rationale: the current standard allows multiple visible sequences of side effects (vsse's) for a given read (despite the use of "The" at the start of 1.10p13). We demonstrate this by constructing an execution with two vsse's. The following execution has five memory operations, four of which are read modify writes (RMW's). There are two threads, one with four operations each ordered by sequenced before (sb), the other with a single RMW release.</p> <pre> RMW1          +---RMW3_release    sb           do/ v R_consume&lt;----+    sb v RMW2    sb v RMW4                     </pre> <p>The modification order in this example is as follows:</p> <pre>           mod          mod          mod RMW1-----&gt;RMW2-----&gt;RMW3_release-----&gt;RMW4                     </pre> <p>With the modification order we give above, the happens-before relation has exactly these edges, according to 1.10p10:</p> <p>From sequenced-before:</p> | <p>on..."</p> <p>and</p> <p>"in the visible sequence of M with respect to B"</p> <p>by</p> <p>"in a visible sequence of M with respect to B"</p> |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | <p>RMW1 -&gt; R_consume, RMW2, RMW4<br/>                     R_consume -&gt; RMW2, RMW4<br/>                     RMW2 -&gt; RMW4</p> <p>From ithb:<br/>                     From dependency-ordered-before:<br/>                     RMW3_release -&gt; R_consume</p> <p>In particular, there are no edges<br/>                     RMW3_release -&gt; RMW2 or RMW4.</p> <p>As we understand it, this is the intended absence of transitivity from dependency-ordered-before to sequenced-before.</p> <p>1.10p5 says that if A happens-before B then A precedes B in the modification order, which is true for all the happens-before edges and the modification order above.</p> <p>RMW1 and RMW3_release are visible side effects</p> <p>RMW2 and RMW4 follow R_consume in happens-before, so cannot be in a visible sequence of side effects.</p> <p>Hence the two visible sequences of side effects are [RMW1] and [RMW3].</p> <p>The R_consume here must read from the later vsse in modification order for the dependency_ordered edge to exist. The existence of two vsse's relies on the lack of transitivity of happens before (which only occurs in the presence of consume operations).</p> |  |   |
| US 10           | 1.10  | Paragraph 14                                 | te                           | The definition of a data race does not take into account two overlapping atomic operations   | <p>Augment the first sentence:</p> <p>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 <u>or both are atomic and operate on overlapping, but</u></p> | REJECTED<br><br>The premise is incorrect; atomic objects may not overlap. The type argument to the atomic template must |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  | <a href="#">not-identical, memory locations</a> ), and neither happens before the other.   | 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.<br><br>See paper <a href="#">N3196</a> |
| US 11           | 1.10  | para7  | te                           | There is some confusion between locks and mutexes.   | Change "lock" when used as a noun to "mutex".  | ACCEPTED<br><br>See paper <a href="#">N3196</a>  |
| US 12           | 1.10  | P4,p14, p6,p12,p13                           | te                           | Adapt N3074:<br><a href="http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2010/n3074.html">http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2010/n3074.html</a>  | Proposed change in N3074:<br><a href="http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2010/n3074.html">http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2010/n3074.html</a>   | ACCEPTED<br><br>See paper <a href="#">N3196</a>  |
| CA 2            | various                                     | various                                      | various                      | Canada agrees with US 12, 14, 142, 145, 159  | Resolve as suggested in these comments   | ACCEPTED   |
| GB 8            | 1.10  | 4, 7   | Te                           | 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 | Change 1.10p4 as follows:<br><br>"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 | ACCEPTED<br><br>See paper <a href="#">N3196</a>  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | <p>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 ]"</p> <p>Change 1.10p7 as follows:</p> <p>"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</p> |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              |   | 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 ]” |  |
| GB 9            | 1.10  | 6  | Te                           | See (B) in attachment Appendix 1 - Additional Details   | Request the concurrency working group to determine if changes are needed  | ACCEPTED<br>See paper <a href="#">N3196</a>        |
| GB 10           | 1.10  | 10   | Te                           | See (C) in attachment Appendix 1 - Additional Details<br>The GB would like WG21 to confirm there is no issue related to this.<br>GB adds:<br>We agree that if the read from x reads the value written by the write to x the write to x inter-thread-happens-before the write to y. However, the read from y is sequenced before the write to x, so if the read from x reads the value written by the write to x, then the read from y also inter-thread-happens-before the write to y. Consequently, the read from y cannot see the value written by the write to y.<br>The reciprocal ordering also applies, but they cannot both apply in the same execution since if the write to x happens-before the read from x then the read from y happens-before the write to y, and vice-versa.<br>There is thus no contradiction.<br>[see comment below for proper formatting] | Request the concurrency working group to determine if changes are needed  | ACCEPTED<br>See paper <a href="#">N3196</a>        |
| CA 8            | 1.10p10                                     | 1.10p10                                      | te                           | Rationale: Without this the standard permits executions with a cyclic happens-before relation that it seems   | 1.10p10, before the Note, add:  | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>clear should be forbidden, e.g.</p> <pre> Rx_consume&lt;--+ +--&gt;Ry_consume                  rf\ /rf    sb                X     sb v                  / \   v Wy_release----+ +----Wx_release                     </pre> <p>One could instead impose acyclicity on happens-before; that would be equivalent.</p>  | "The inter-thread happens-before relation of an execution must be acyclic" |  |
| GB 11           | 1.10  | 12   | Te                           | <p>See (E) in attachment Appendix 1 - Additional Details<br/> The GB would like WG21 to confirm there is no issue related to this.<br/> GB adds:<br/> [see comment below for proper formatting]<br/> The variable in question has a single modification order, which is any of<br/> (a) RMW3, RMW1, RMW2, RMW4.<br/> (b) RMW1, RMW3, RMW2, RMW4.<br/> (c) RMW1, RMW2, RMW3, RMW4.<br/> (d) RMW1, RMW2, RMW4, RMW3.<br/> since RMW1, RMW2 and RMW4 occur in a single thread in that sequence, and<br/> RMW3 occurs in a separate thread with no other ordering constraints.<br/> Since the R_consume lies between RMW1 and RMW2 in that thread, it must<br/> either read the value written by RMW1 (which could happen if it<br/> immediately follows RMW1 in any of the sequences), or<br/> RMW3 (which could<br/> happen with sequence (b)).<br/> The visible sequence of side effects for R_consume is</p> | Request the concurrency working group to determine if changes are needed   | ACCEPTED<br>See paper <a href="#">N3196</a>        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                                     |
|                 |   |  |                              | thus either RMW3, RMW1 (from (a)), RMW1 (from (b), (c) or (d)), or RMW1, RMW3 (from (b)). Which sequence applies in practice may vary from execution to execution. There is however only a single sequence on any given execution.   |   |  |
| GB 12           | 1.10  | 13   | Te                           | See (F) in attachment Appendix 1 - Additional Details<br>The GB would like WG21 to confirm there is no issue related to this.<br>GB adds:<br>The cycle given is clearly forbidden by the current text. The read is sequenced-before the write in the same thread. If the read sees the value written by the other thread then that write is dependency-ordered-before the read, and thus happens-before the read, and happens-before the write from the reading thread. The write from the left-hand thread thus must occur before the write from the right-hand thread in the modification order of the object by 1.10p5. | Request the concurrency working group to determine if changes are needed  | ACCEPTED<br><br>See paper <a href="#">N3196</a>  |
| GB 13           | 1.10  | 13   | Te                           | See (G) in attachment Appendix 1 - Additional Details<br>GB suggests alternative wording to that in the attached paper:<br>"Furthermore, if a value computation A of an atomic object M happens before a value computation B of M, and A uses the value of M 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."  | Request the concurrency working group to determine if changes are needed  | ACCEPTED WITH MODIFICATIONS<br><br>See paper <a href="#">N3196</a>                     |
| GB 14           | 1.10  | 8  | Te                           | See (I) in attachment Appendix 1 - Additional Details<br>GB adds:<br>If an implementation can't guarantee the ordering it should refrain from performing the optimisation  | Request the concurrency working group to determine if changes are needed. | REJECTED<br><br>We agree that the implementation needs to address this, and it can. No |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  |   | changes are needed.  |
| GB 15           | 1.10  |  | Te                           | See (J) in attachment Appendix 1 - Additional Details  | Request the concurrency working group to determine if changes are needed. | REJECTED<br><br>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.<br><br>See paper <a href="#">N3196</a> |
| GB 16           | 1.10  | 12   | Ed                           | See (L) in attachment Appendix 1 - Additional Details  | "...as defined here..." should be "...as defined below...".               | REJECTED<br><br>"As defined here" refers to this standard, which is the intention. If it were normative text the correct phrase would be "in this International Standard", but for a note, the text as written is sufficient.  |
| US 13           | 2.2   | 1  | te                           | "Raw" strings are still only Pittsburgh-rare strings: the reversion in phase 3 only applies to an r-char-sequence.   | Make the reversion apply to the entire raw-string.                        | ACCEPTED   |
| US 14           | 2.2<br>2.3<br>2.5                           | P1   | te                           | Precedence of reversal and tokenization<br><br>The current paper implies that determination of the characters forming an <i>r-char-sequence</i> occurs while the | <b>In 2.14.5 [lex.string] paragraph 2:</b> Remove footnote 24:            | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 | 2.14.5                                      |  |                              | <p>transformations done in phase 1 and phase 2 are still in effect.</p> <p>Consider these cases:</p> <ul style="list-style-type: none"> <li>Line splicing occurred in translation phase 2; the backslash is not there on entry to phase 3 when we try to tokenize:<br/>const char str[] = R"a()\a"a";</li> <li>Trigraph replacement occurred in phase 1. The right parenthesis is not there on entry to phase 3:<br/>const char str[] = R"(??)";</li> <li>Trigraph replacement (again). In [lex.string] paragraph 2, there is a footnote 24 in N3092. Note that this provides fuel for anti-trigraph sentiment:<br/>const char str[] = R"#(???)#";</li> </ul> <p><b>Change in [lex.string] from N3077:</b></p> <p>Escape sequences and universal-character-names in non-raw string literals have the same meaning as in character literals ....</p> <p>should be reflected in [lex.phases] paragraph 1, phase 5 (CD2 wording):</p> <p>Each source character set member and universal-character-name in a character literal or a string literal, as well as each escape sequence in a character literal or a non-raw string literal, is converted to the corresponding member of the execution character set (2.14.3, 2.14.5); if there is no corresponding member, it is converted to an implementation-defined member other than the null (wide</p> | <p><b>In 2.2 [lex.phases] paragraph 1, phase 1; insert exception:</b></p> <p>Physical source file characters are mapped, in an implementation-defined manner, to the basic source character set (introducing new-line characters for end-of-line indicators) if necessary. The set of physical source file characters accepted is implementation-defined. Trigraph sequences (2.4) are replaced by corresponding single-character internal representations. Any source file character not in the basic source character set (2.3) is replaced by the universal-character-name that designates that character. (An implementation may use any internal encoding, so long as an actual extended character encountered in the source file, and the same extended character expressed in the source file as a universal-character-name (i.e., using the \uXXXX notation), are handled equivalently except where this replacement is reverted.)</p> <p><b>In 2.2 [lex.phases] paragraph 1, phase 3:</b></p> <p>The source file is decomposed into preprocessing tokens (2.5) and sequences of white-space characters (including comments). A source file shall not end in a partial preprocessing token or in a partial comment. Each comment is replaced by one space character. New-line characters are retained. Whether each nonempty sequence of</p> |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>character</p> <p>and [lex.charset] paragraph 2 (CD2 wording):</p> <p>Additionally, if the hexadecimal value for a universal-character-name outside the <i>c-char-sequence</i>, <i>s-char-sequence</i>, or <i>r-char-sequence</i> of a character or string literal corresponds to a control character (in either of the ranges 0x00–0x1F or 0x7F–0x9F, both inclusive) or to a character in the basic source character set, the program is ill-formed.</p> <p>UCNs simply do not occur in the grammar for <i>r-char-sequence</i> anyway.</p> | <p>white-space characters other than new-line is retained or replaced by one space character is unspecified. The process of dividing a source file's characters into preprocessing tokens is context-dependent. [ <i>Example</i>: see the handling of &lt; within a #include preprocessing directive. —<i>end example</i> ]</p> <p><b>In 2.2 [lex.phases] paragraph 1, phase 5:</b></p> <p>Each source character set member in a character literal or a string literal, as well as each escape sequence and universal-character-name in a character literal or a non-raw string literal, is converted to the corresponding member of the execution character set (2.14.3, 2.14.5); if there is no corresponding member, it is converted to an implementation-defined member other than the null (wide) character.</p> <p><b>In 2.3 [lex.charset] paragraph 2:</b></p> <p>.... Additionally, if the hexadecimal value for a universal-character-name outside the <i>c-char-sequence</i> or <i>s-char-sequence</i> of a character or string literal corresponds to a control character (in either of the ranges 0x00–0x1F or 0x7F–0x9F, both inclusive) or to a character in the basic source character set, the program is ill-formed. [ Footnote: A sequence of characters resembling a universal-character-name in an <i>r-char-sequence</i> (2.14.5 [lex.string]) does not form a universal-character-name. ]</p> |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | <p><b>In 2.5 [lex.pptoken] paragraph 3:</b></p> <p>If the input stream has been parsed into preprocessing tokens up to a given character:</p> <ul style="list-style-type: none"> <li>if the next character begins a sequence of characters that could be the prefix and initial double quote of a raw string literal, such as R", the next preprocessing token shall be a raw string literal and any transformations performed in phases 1 and 2 on this input stream (trigraphs, universal-character-names, and line splicing) are reverted for the remainder of the stream until said raw string literal (2.14.5) is matched; [ Footnote: A raw string literal formed through token concatenation (16.3.3) is not parsed from an input stream and is not subject to this reverting. Destringization (16.9) involves an alternate input stream, thus there are no phase 1 or phase 2 transformations to revert. ]</li> <li>otherwise, the next preprocessing token is the longest sequence of characters that could constitute a preprocessing token, even if that would cause further lexical</li> </ul> |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB                                     | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | analysis to fail.   |  |
| US 15           | 2.6   | para 2                                       | te                           | The <: digraph causes problem with users unfamiliar with digraphs when passing global objects as template arguments.   | Add a special hack for <:: much like the special hack for >>. | ACCEPTED   |
| CA 24           | 2.11  | Various                                      | Te                           | <p>A list of issues related TR 10176:2003</p> <p>1)<br/>"Combining characters should not appear as the first character of an identifier."<br/>Reference: ISO/IEC TR 10176:2003 (Annex A)<br/>This is not reflected in FCD.</p> <p>2)<br/>Restrictions on the first character of an identifier are not observed as recommended in TR 10176:2003. The inclusion of digits (outside of those in the basic character set) under <i>identifer-nondigit</i> is implied by FCD.</p> <p>3)<br/>It is implied that only the "main listing" from Annex A is included for C++. That is, the list ends with the Special Characters section. This is not made explicit in FCD. Existing practice in C++03 as well as WG 14 (C, as of N1425) and WG 4 (COBOL, as of N4315) is to include a list in a normative Annex.</p> <p>4)<br/>Specify width sensitivity as implied by C++03: \uFF21 is not the same as A<br/>Case sensitivity is already stated in [lex.name].</p> | Please clarify.   | ACCEPTED<br><br>See paper <a href="#">N3146</a>    |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| GB 17           | 2.14.2                                      | Table 5                                      | Ed                           | [lex.icon] 2.14.2/2 Table 5 - 'Types of integer constants' In the penultimate row for this table (for suffix 'll or LL') it gives the 'Octal or hexadecimal constant' in the third column as one of: long long int unsigned long int Unless I am misunderstanding something fundamental, this second should be: unsigned long long int | Replace the entry for "ll or LL" and "Octal or hexadecimal constant" in table 5 with "long long int unsigned long long int" | ACCEPTED   |
| JP 16           | 2.14.3                                      | 2 Note                                       | E                            | Typo, "wide-charater" should be "wide-character".  | Correct typo.<br>[ Note: the type wchar_t is able to represent all members of the execution wide-character set (see 3.9.1). | ACCEPTED   |
| RU 2            | 2.14.3                                      | p.23, par.3, line 1                          | ed                           | Reference missed   | Insert reference "(3.9.1)" after "extended integer type"  | ACCEPTED   |
| DE 2            | 2.14.4                                      |  | te                           | C++ does not support hexadecimal floating-point literals, although they are useful to specify exact floating-point constants.  | Consider supporting the C99 syntax for hexadecimal floating-point literals.   | REJECTED<br><br>There was no consensus to adopt this feature at this point in the standardization process. |
| US 16           | 2.14.5 [lex.string]                         |  | ge                           | Raw string literals have no implementation experience.   | Either demonstrate a complete implementation of this feature or remove N2146 from the working paper prior the FDIS.         | ACCEPTED<br><br>The feature has been implemented. No change to the Standard.                               |
| DE 3            | 2.14.7                                      |  | te                           | It is not sufficiently clear that std::nullptr_t is a distinct type and neither a pointer type nor a pointer-to-member type.   | Add a note in 2.14.7 stating that, preferably with cross-references to the normative statements in 3.9.                     | ACCEPTED   |
| RU 5            | 2.14.7                                      | p. 28  | ed                           | Page layout bug  | Move footnote 24 from page 28 to page 27  | ACCEPTED   |
| US 17           | 2.14.8                                      | 6  | te                           | In general, the parameter type of a literal operator must be the same as the argument passed to it. That is not the case for a <i>user-defined-character-literal</i> , where the   | Add the following phrase to the description in paragraph 6:<br><br>S shall contain a literal operator whose parameter       | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              | argument could inadvertently match a literal operator intended for use with <i>user-defined-integer-literals</i> :<br><br>typedef unsigned long long ULL;<br>int operator "" X(ULL);<br>int i = 'c'X; // operator"" X(ULL('c')) | type is the same as the type of <i>ch</i> .  |  |
| JP 17           | 2.14.8                                      | 3  | E                            | Typo, missing ",".<br>If S contains a raw literal operator the literal L is treated as  | Correct typo.<br>If S contains a raw literal operator, the literal L is treated as   | ACCEPTED   |
| JP 18           | 2.14.8                                      | 4  | E                            | Typo, missing ",".<br>If S contains a raw literal operator the literal L is treated as  | Correct typo.<br>If S contains a raw literal operator, the literal L is treated as   | ACCEPTED   |
| US 18           | 2.24.8 [lex.ext]                            |  | ge                           | User-defined literals have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N2750 from the working paper prior the FDIS.  | REJECTED<br>There was no consensus to adopt this change.<br><br>(Reference should be to 2.14.8, not 2.24.8.) |
| US 19           | 3   | 4  | te                           | It is not always clear when the term "use" is intended as a reference to the definition in 3.2 and when it has its normal English meaning. For example, 3 paragraph 4 reads, "A <i>name</i> is a use of an identifier..."       | Replace all occurrences of the word "use" that are not intended as references to 3.2 with some other term, such as "occurrence" or "appearance" or "reference to". | ACCEPTED<br><br>See paper <a href="#">N3214</a>  |
| US 20           | 3.1   | para 1 bullet 4                              | ed                           | Grammatical number mismatch in "an assignment expressions".   |  | ACCEPTED   |
| US 21           | 3.1   | 2  | ed                           | using N::d; does not declare N::d.  | using N::d; // declares d  | ACCEPTED   |
| US 22           | 3.2   | 4  | te                           | The type of the <i>expression</i> of a <i>decltype-specifier</i> is apparently required to be complete.   | Make an exception so that a template specialization type is not instantiated merely because it's the type of the <i>expression</i> in                              | ACCEPTED with MODIFICATIONS<br><br>The premise is incorrect: it is   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              |   | <code>decltype( expression )</code>   | the function call, not <code>decltype</code> , that requires a complete type.<br>See paper <a href="#">N3276</a> |
| JP 19           | 3.2   | 4  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(10)" to "(Clause 10)".   | ACCEPTED   |
| JP 20           | 3.3.2                                       | 7  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(9)" to "(Clause 9)".   | ACCEPTED   |
| US 23           | 3.4.5                                       | para 1                                       | te                           | Global class templates should not hide member templates.  | Strike the end of para 1 starting with "If the lookup in the class of the object expression finds a template,". See Appendix 1 - Additional Details | ACCEPTED   |
| US 24           | 3.5   | 3  | te                           | One of the criteria for giving a name internal linkage is "a variable that is explicitly declared const and neither explicitly declared extern nor previously declared to have external linkage." This should presumably apply to variables declared <code>constexpr</code> as well.  | Add parallel wording for the <code>constexpr</code> specifier.  | ACCEPTED   |
| DE 4            | 3.5   |  | te                           | It is odd that "N" has no linkage and "g" has external linkage in this example:   |   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | <pre>namespace {   namespace N // has no linkage   {     void g(); // has external linkage   } }</pre>   |  |   |
| DE 5            | 3.7.3                                       |  | te                           | The term "local" was changed globally to "block-scope", but this section still contains the term "local" (see also core issue 642).  | Change "local" to "block-scope" in the first paragraph.  | ACCEPTED  |
| RU 3            | 3.7.4.3                                     | p.65, line 7                                 | ed                           | Reference missed   | Insert reference "(5.7)" after "well-define pointer arithmetic"  | ACCEPTED  |
| RU 4            | 3.7.4.3                                     | p. 65, line 8                                | ed                           | Reference missed   | Insert references "(4.10, 5.4)" after "well-define pointer conversion"   | ACCEPTED  |
| GB 18           | 3.8   | 9  | Te                           | It isn't clear that the comment in the example actually reflects the result of the placement new. If the intended placement operator new is supposed to be the one given by the standard library ,by including , the example is ill-formed as the placement-new expression &b is const B* which doesn't implicitly convert to void*. | Replace:<br>new (&b) const B;<br>With:<br>new (const_cast<B*>(&b)) const B;  | ACCEPTED  |
| US 25           | 3.11  |  | Te                           | C/C++ compatibility problems defined in WG21/N3093.  | Make the changes proposed in WG21/N3093  | ACCEPTED<br>See paper <a href="#">N3190</a>   |
| US 26           | 3.7.4, 5.3.5, 12.5, 17.6.3.6, 18.6          |  | te                           | Programmers may define a static member function operator delete that takes a size parameter indicating the size of the object to be deleted. The equivalent global operator delete is not available. This omission has unfortunate performance consequences.   | Permit implementations and programmers to define sized versions of the global operator delete for use in preference to the unsized version. See Appendix 1 - Additional Details  | REJECTED<br>There was no consensus for making the suggested change at this point in the standardization process.                            |
| US 27           | 3.8   | 4  | te                           | Related to core issue 1027, consider:<br><br><pre>int f() {   union U { double d; } u1, u2;   (int&amp;)u1.d = 1;   u2 = u1;</pre>   | Clarify that this testcase is undefined, but that adding an array of unsigned char to union U would make it well-defined--if a storage location is allocated with a particular type, it should be undefined to create an object in that storage if it would be undefined to access the stored value of | REJECTED. There was no consensus to adopt this proposed change for this revision. However, an issue on this item has been opened for future |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <pre>return (int&amp;)u2.d; }</pre> <p>Does this involve undefined behavior? 3.8/4 seems to say that it's OK to clobber u1 with an int object. Then union assignment copies the object representation, possibly creating an int object in u2 and making the return statement well-defined. If this is well-defined, compilers are significantly limited in the assumptions they can make about type aliasing. On the other hand, the variant where U has an array of unsigned char member must be well-defined in order to support std::aligned_storage.</p> | the object through the allocated type.  | consideration.   |
| US 28           | 4.4   | para 3                                       | te                           | A const member function pointer could safely be applied to a non-const object without violating const correctness.   | Add an implicit conversion. See Appendix 1 - Additional Details   | REJECTED<br><br>There was no consensus for adding this feature at this point in the standardization process. |
| FI 7            | 4.11 [conv.mem], 5.2.9 [expr.static.cast]   |  | te                           | The CD1 comment CH1 should be reconsidered. The request for being able to cast a pointer to member to a pointer to a base class (or any other implicitly convertible type) of the member is a bugfix rather than an extension. It's a safe conversion, thus it should be allowed. There are valid use cases for such conversions that are currently forbidden.   | The standard should allow implicit conversions from "pointer to member of T of type cv D" to "pointer to member of T of type cv B", where D is of class type and B is a public base of D. It should allow explicit conversion the other way around.   | REJECTED<br><br>There was no consensus for adding this feature at this point in the standardization process. |
| CH 3            | 4.11 and 5.2.9                              |  | te                           | With respect to the target type, pointer to members should behave like normal pointers. The current situation creates an inconsistency in the C++ type system and is therefore a defect in the Standard.   | The standard should allow implicit conversions from "pointer to member of T of type cv D" to "pointer to member of T of type cv B", where D is of class type and B is a public base of D. It should allow explicit conversion in the other direction. | REJECTED<br><br>There was no consensus for adding this feature at this point in the standardization process. |
| JP 21           | 4.13  | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not  | Change "(5)" to "(Clause 5)".   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  |   |  |
| JP 1            | 5   | Paragraph 6                                  | TL                           | The first half of the Note(before "In general") indicates that the expression "E1.E2" produces xvalue if E1 is xvalue regardless of E2's type. It will be true even if E2 is of reference type. On the other hand, according to 5.2.5 paragraph 4, if E2 is of reference type, the result of "E1.E2" is lvalue regardless of E1's type. These two descriptions contradict each other. As 5.2.5 paragraph 4 seems correct, 5 paragraph 6 should be corrected. | Modify 5 paragraph 6 so that the result of E1.E2 is lvalue instead of xvalue when E2 is of reference type.  | ACCEPTED   |
| FI 8            | 5.1.2 [expr.prim.lambda]                    |  | te                           | As requested in JP 9 on CD, capturing by moving should be allowed for lambdas. Roshan Naik presents a very compelling use case in the Core Reflector message c++std-core-16341.  | Allow specifying capture by move.   | REJECTED<br><br>There was no consensus for making the suggested change at this point in the standardization process.   |
| CH 4            | 5.1.2                                       | p1   | ed                           | Document N3067 changed the position of attribute specifiers in various places. However, it left out lambda expressions as an oversight, so that the position of attribute-specifier opt in a lambda-declarator is inconsistent with a function declarator  | change the rule for lambda-declarator to<br><i>lambda-declarator:</i><br><i>( parameter-declaration-clause ) mutable<sub>opt</sub> exception-specification<sub>opt</sub> attribute-specifier<sub>opt</sub> trailing-return-type<sub>opt</sub></i> | ACCEPTED   |
| CH 5            | 5.1.2                                       | p4 first bullet                              | ed                           | typo   | Change second 'if' to 'is'.   | ACCEPTED   |
| US 29           | 5.1.2                                       | 5  | te                           | default arguments should be allowed in lambdas (core issue 974)  | See Appendix 1 - Additional Details   | REJECTED<br><br>There was no consensus for making the suggested change at this point in the standardization process. However, core language issue 974 remains open for consideration in a future |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  |  | revision.  |
| US 30           | 5.1.2                                       | 4  | te                           | lambda return type deduction should allow arbitrary function structure (core issue 975)  | See Appendix 1 - Additional Details  | REJECTED<br><br>There was no consensus for making the suggested change at this point in the standardization process. However, core language issue 975 remains open for consideration in a future revision. |
| JP 22           | 5.1.2                                       | 7  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  | Change "(5)" to "(Clause 5)".  | ACCEPTED   |
| CH 6            | 5.1.2                                       | p8 and p10                                   | te                           | The current capturing rules seem too restrictive.  | Consider to make those rules less restrictive.   | REJECTED<br><br>The comment made no specific suggestions for change.   |
| GB 19           | 5.1.2                                       | 16   | Ed                           | [expr.prim.lambda] 5.1.2/16 has text which begins "If a lambda-expression m1 captures an entity and that entity is captured by an immediately enclosing lambda expression m2..." - that is, it describes a situation with m2 enclosing m1, and then describes the capture transformation in these terms. The example given to support this, however, turns this all around and shows m1 enclosing m2. This doesn't make either the text or the example incorrect in any sense, but I would suggest that it adds a level of confusion that is | All references to m1 from the beginning of 5.1.2/16 up to the last occurrence before '[Example ' to be replaced by m2, and vice versa. Rationale for suggested wording: all other examples that use the 'mN' notation for lambda expressions and which involve nesting apply increasing N (by 1, from 1) to indicate increasing nesting depth. | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                  |
|                 |   |  |                              | easily avoided.   |  |   |
| GB 20           | 5.1.2                                       | 12   | Ed                           | [expr.prim.lambda] 5.1.2/12. In the example code given the local struct s1 has a member function with signature int s1::work(int n) whose definition does not include an appropriate return statement; neither does it include the conventional "// ..." to indicate that the example is intended to be incomplete.   | Suggested change: change the signature of this member function to void s1::work(int n), as the return of int does not contribute to the example.   | ACCEPTED  |
| GB 21           | 5.1.2                                       |  | Te                           | A lambda-capture can be &, which indicates it captures everything by reference (unless otherwise specified), or &/identifier/, which indicates it captures the /identifier/ by reference. It can also be =, to capture everything by value, or /identifier/, for a single thing. Why is = /identifier/ not allowed, for consistency?  | Add "= identifier" to the grammar in 5.1.2p1. (The wording already covers the semantics of this, since it refers to captures that "are preceded by &" or "do not contain &")   | REJECTED<br>There was no consensus for making the suggested change. |
| FI 19           | 5.1.2 [expr.prim.lambda]                    | 21   | te                           | "When the lambda-expression is evaluated, the entities that are captured by copy are used to direct-initialize each corresponding non-static data member of the resulting closure object. " This apparently means that if the capture-default is to copy, entities captured by default, implicitly, are copied even in cases where the copy constructors of such entities are explicit.   | Don't implicitly copy entities that have explicit copy constructors declared. Require that such entities be captured explicitly, by enumerating them in the capture list. This seems related to Core Issue 1020, so I'd like that issue to be resolved as well.<br>See Appendix 1 - Additional Details | REJECTED<br>There was no consensus for making the suggested change. |
| GB 23           | 5.2.2                                       | 4  | Ed                           | Order of initialization of arguments in a function is fundamental to the memory model of C++, and the obvious place to look for the definition is in the clause defining function call operators - which currently says nothing. The rules covering this are buried in paragraph 15 of [1.9]. A cross-reference to these words would be most helpful. In particular, it should be made clear that such initialization is indeterminately sequenced (and not unsequenced.) | Add a non-normative note with cross-reference after the first sentence of paragraph 4: "[Note - such initializations are indeterminately sequenced with respect to each other [1.9] - end note]"   | ACCEPTED  |
| US 31           | 5.2.9;7.2                                   | 10   | te                           | it is unclear from the text in 7.2 and 5.2.9 that the "values of the enumeration" term does not exclude a pvalue of an enumeration type from having other values representable  | clarify this. "The value is unchanged if it is in the range of enumeration values of the enumeration type; otherwise the resulting enumeration value is  | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | in its underlying type (c++std-core-15652).   | unspecified ( <b>and might not be in that range</b> ).<br>Also add a note after paragraph 7 "[Footnote: this set of values is used to define promotion and conversion semantics for the enumeration type; it does not exclude an expression of enumeration type from having a value that falls outside this range.]"   |  |
| US 32           | 5.2.5                                       | 5  | te                           | The description of ambiguity ("...if the class of which E2 is directly a member is an ambiguous base (10.2) of the naming class (11.2) of E2") does not cover the following case:<br><br><pre>struct A { int i; }; struct B: A {}; struct C: A, B {}; void f(C* p) {     p-&gt;i; // Should be ambiguous }</pre>                        | Change the wording to apply also to the case when the naming class is an ambiguous base of the class of the object expression.   | ACCEPTED   |
| JP 64           | 5.2.8                                       | 5  | E                            | In some code examples, ellipsis(...) is used in ill-formed. In these cases, "..." represents omission of some codes like this:<br><pre>class A { /* ... */ };</pre> But in some cases, it is used without commented-out as below:<br><pre>class A { ... };</pre> It is an inconsistent usage. They all should be enclosed in a comment. | Change to:<br><pre>class D { /* ... */};</pre>   | ACCEPTED   |
| GB 22           | 5.2.10                                      |  | Te                           | It is not legal to use reinterpret_cast<> with pointers to void.  | Here's an additional paragraph to add to §5.2.10 that would fix this:<br>* A pointer to an object type can be explicitly converted to a pointer to void, and vice versa.[1]<br>The result of such a pointer conversion will have the same result as the standard pointer conversion described in §4.10. A value of type "pointer to object" converted to "pointer to void" | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|--|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)      | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |  |  |                              |   | and back, possibly with different cv-qualification, shall have its original value.<br>[1] The types may have different cv-qualifiers, subject to the overall restriction that a reinterpret_cast cannot cast away constness. |   |
| US 33           | 5.3.1  | 3  | te                           | The resolution of issue 983 added an error for finding the named member in an ambiguous base. This is an unnecessary special case, since the value of the expression is pointer to member of base. If this value is then converted to pointer to member of derived, an error will be given at that point.   | Revert the change for issue 983 (and in the issues list, add a link to issue 203).   | ACCEPTED  |
| GB 24           | 5.3.3  | 6  | Te                           | The return type of the sizeof operator is defined as being of type std::size_t, defined in library clause 18.2. This, in turn, says that size_t is defined in the C standard, which in turn says that size_t is defined as the type of the result of the sizeof operator!<br>The C definition of sizeof returns an implementation-defined unsigned integer type, recommended not to have "an integer conversion rank greater than signed long int, unless the implementation supports objects large enough to make this necessary." | The result type of the sizeof operator should explicitly be implementation defined in clause 5.3.3.  | ACCEPTED  |
| US 34           | 5.3.4, 5.3.5                                     |  | te                           | Allocations functions are missing happens-before requirements and guarantees.   | Add requirements. See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br><br>See LWG 1524   |
| US 35           | 5.3.7 [expr.unary.no except], 15.4 [except.spec] |  | ge                           | noexcept has no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N3050 from the working paper prior the FDIS.  | REJECTED<br>The feature has been implemented.   |
| FI 17           | 5.3.7 [expr.unary.no except]                     |  | te                           | Destructors should by default be noexcept. Such a rule should, I think, be obeyed even for cases where a destructor is defaulted. Then a throwing destructor would need to be declared noexcept(false), and I think the resulting code breakage is acceptable.  | Clarify the implicit generation rules and defaulting rules so that destructors are noexcept unless explicitly declared noexcept(false).  | ACCEPTED with MODIFICATIONS<br><br>The rule that was adopted makes destructors noexcept if all the base and member destructors are. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              |   |  | See paper <a href="#">N3204</a>   |
| JP 2            | 5.5   | 6  | TL                           | Should be corrected because it contradicts with rules in 5.2.5 paragraph 4.   | Add the condition that a type of e2 is not a reference type  | REJECTED<br><br>There are no pointers to members of reference type (8.3.3p3).   |
| US 36           | 5.19 [expr.const]                           |  | ge                           | Generalized constant expressions have no implementation experience.   | Either demonstrate a complete implementation of this feature or remove N2235 from the working paper prior the FDIS.  | REJECTED<br><br>An implementation is in progress and is expected to be complete before the publication of the FDIS. Also, this feature is needed in the Standard library, thus should not be removed, and there was no consensus to do so.  |
| DE 6            | 5.19  |  | te                           | Paragraph 1 is interpreted by some readers to restrict the following definition of a "constant expression" to apply only where a constant expression is required. | Replace the first two normative sentences of paragraph 1 by "This sub-section defines <i>constant expressions</i> ." | ACCEPT WITH MODIFICATIONS<br>Changed to read:<br><br>Certain contexts require expressions that satisfy additional requirements as detailed in this sub-clause; other contexts have different semantics depending on whether or not an expression satisfies these requirements. Expressions that satisfy these requirements are called constant expressions.<br>[ Note: Constant expressions can be evaluated during translation.<br>-- end note ] |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |   |
|-----------------|---|--|------------------------------|---|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |   |
| DE 7            | 5.19  | 4, note                                      | te                           | The note in paragraph 4 alludes to the possibility that compile-time and run-time evaluations of floating-point expressions might yield different results. There is no clear normative statement (other than the absence of a restriction) that gives such permission.  | Move the second sentence of the note into normative text.                     | REJECTED<br><br>The possibility of differing results of calculations is implicit in the absence of normative statements constraining their accuracy, so no normative change is needed; the existing note is sufficient to point out this implication. |   |
| DE 8            | 5.19  | 6  | te                           | In the definition of "potential constant expression" in paragraph 6, it is unclear how "arbitrary" the substitution of the function parameters is. Does it mean "there exists a value for which the result is a constant expression" or does it mean "for all possible values, the result needs to be a constant expression"? Example:<br><br>constexpr int f(int x){return x + 1; }<br>is a constant expression under the first interpretation, but not under the second (because overflow occurs for x == INT_MAX, cf. 5.19p2 bullet 5). The answer also affects expressions such as:<br>constexpr int f2(bool v) { return v ? throw 0 : 0; }<br>constexpr int f3(bool v) { return v && (throw 0, 0); } |   |   | ACCEPTED<br><br>See paper <a href="#">N3218</a> |
| GB 25           | 5.19  |  | Te                           | In trying to pin down the behaviour of constexpr functions, it became apparent that there is disagreement over whether or not the following example is well-formed.<br><br>constexpr int f() { return 42 + 84; }<br>const int sz = f();<br>int a[sz];<br><br>This should have the same effect as<br><br>const int sz = 42 + 84;   | Update the wording in 5.19 to make it clear that both the examples are valid. | ACCEPTED  |   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | int a[sz];<br>otherwise constexpr functions are broken.   |   |  |
| GB 26           | 5.19  |  | Te                           | <p>It is not clear how overload resolution applies within a constexpr function. In particular, if an expression in the function body yields an integral value of 0 for some parameter values, and not for others, is it usable as a null pointer constant when evaluated as a constant expression.</p> <pre> typedef char (&amp;One)[1]; typedef char (&amp;Two)[2]; One f(void*); // #1 Two f(...); // #2  constexpr int test(int n) { return sizeof f(n); } constexpr int test2(int n) { return sizeof f(n*0); }  int q = 0;  #include  int main() { char a[test(0)]; std::cout &lt;&lt; sizeof(a) &lt;&lt; std::endl; // #3 std::cout &lt;&lt; test(q) &lt;&lt; std::endl; // #4  char b[test2(0)]; std::cout &lt;&lt; sizeof(b) &lt;&lt; std::endl; // #5 std::cout &lt;&lt; test2(q) &lt;&lt; std::endl; // #6 } </pre> <p>#3 and #4 should print 2, since n is not an integral constant expression with value 0 in the body of test() --- though it is a constant expression when test() is evaluated as a constant expression, it's value is dependent on the invocation. Permitting different results of overload</p> | Updated 5.19 to make it clear that overload resolution in a constexpr function is not dependent on the context of use, or the value of the arguments. | ACCEPTED<br>See paper <a href="#">N3218</a>        |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <p>resolution within the same function body in different calling contexts would violate ODR.</p> <p>On the other hand, in test2(), the answer is no longer dependent on the value of n, since "n*0" always evaluates to 0. However, it is not clear from the FCD whether "n*0" is thus a constant expression (and therefore a valid null pointer constant) inside the body of test2. Either way both #5 and #6 should print the same value; it would violate ODR for #5 to print "1" (indicating that "n*0" was a valid null pointer constant when test2() is evaluated in a constant expression context) whilst #6 prints "2" (since n*0 is not a constant expression if n is not constant). #5 and #6 should thus both print "1", or both print "2".</p> |  |  |
| US 37           | 6.5   |  | te                           | <p>"for (auto e : range)" creates copies of elements. This seems like a gotcha for new users. Not only are copies inefficient for reading, but writing to copies won't modify the original elements.</p> <p>Permitting "for ( identifier : expression )" and giving it the same meaning as "for ( auto&amp; identifier : expression )" would make the range-based for statement easier to teach and to use, and should be trivial to specify and to implement.</p>   | Permit "for ( identifier : expression )" or similar, with the same meaning as "for ( auto& identifier : expression )". | <p>REJECTED</p> <p>There was no consensus for making the suggested change at this point in the standardization process.</p>  |
| CA 3            | Various                                     | various                                      | various                      | Canada agrees with US 37, 44, 47, 85, 77, 92, 97, 102, 105, 109  | Resolve as suggested in these comments   | <p>ACCEPTED with MODIFICATIONS</p> <p>US 37, US 92, and US 105 were rejected. US 44, US 77, US 97, US 102, and US 109 were accepted. US 47 and US 85 were accepted with modifications.</p> |
| US 38           | 6.5   | 5  | te                           | 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  | As a strawman, replace the paragraph with "The implementation may assume that any                                      | <p>ACCEPTED</p> <p>See paper <a href="#">N3196</a></p>   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | would take advantage of this can often no longer identify the origin of such a loop.  | <p>program will eventually do one of the following:</p> <ul style="list-style-type: none"> <li>- terminate,</li> <li>- make a call to a library I/O function,</li> <li>- access or modify a volatile object, or</li> <li>- perform a synchronization operation (1.10) or atomic operation (Clause 29)."</li> </ul> <p>Possibly move this and the attached note to section 1.9, after p8.</p> |   |
| JP 69           | 6.35.1                                      | 2  | E                            | Constant width font should be used for 3 "while"s in the paragraph as described in Syntax notation (1.6).   | Change the font for "while" to constant width type. When the condition of a <b>while</b> statement is a declaration, the scope of the variable that is declared extends from its point of declaration (3.3.2) to the end of the <b>while</b> statement. A <b>while</b> statement of the form   | ACCEPTED  |
| GB 27           | 6.5.4p1                                     |  | Te                           | <p>6.5.4/1 requires that range-based for loops behave as if they had "{ auto&amp;&amp; __range = (expression); for (auto __begin = begin_expr, __end = end_expr; ...)" which implies that __begin and __end must have the same type. However, this prevents stateful iterators with an end sentinel of a different type. Since range-based for loops' equivalent code already introduces a block (for the __range variable), could __begin and __end be placed there, as "auto __begin = begin_expr; auto __end = end_expr;"?</p> <p>Example of what this change would allow, only the relevant details shown with ctors, op*, op++, etc. omitted: (apologies if the formatting is lost)</p> <pre>struct End {}; struct Counter { Counter&amp; begin() { return *this; } // used by begin_expr End end() { return End(); } //</pre> | <p>Change the "as if" for a range-based for-loop in 6.5.4p1 to move the initialization of __begin and __end outside the loop into the enclosing block:</p> <pre>"{ auto&amp;&amp; __range = (expression); auto __begin = begin_expr; auto __end = end_expr; for (; ..."</pre>  | REJECTED  |
|                 |   |  |                              |   |  | There was no consensus for making the suggested change. Use of iterators of different types is incompatible with the Standard library containers and with the earlier concepts-based specification. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              | used by end_expr bool operator!=(End) const { return _current != _end; }<br><br>Counter(int begin, int end) : _current(begin), _end(end) {}<br>int _current, _end; };<br><br>void use_example() { for (auto n : Counter(0, 10)) { // n takes values from 0..9 } }   |  |  |
| CH 7            | 6.5.4                                       | p1   | te                           | The phrasing "is equivalent to" is too restrictive and might constrain future improvements.   | Make clear that the specification is not necessarily the implementation, i.e. that the expressions in the specification are not necessarily called at all and that the order in which the statement is executed for different values of for-range-declaration is not necessarily the same as if the for loop would have been written the way in the specification. | REJECTED<br><br>There was no consensus for making the suggested change. The Standard needs to specify the meaning of the statement; removing the "equivalent to" phrasing would leave the feature underspecified. Implementations have a good deal of latitude under the "as-if" rule. |
| GB 28           | 7   | 1  | Ed                           | "Attributes" is misspelled  | Replace "Attributes" with "Attributes"   | ACCEPTED   |
| US 39           | 7.1   | 1  | te                           | The current wording is, "The optional <i>attribute-specifier</i> in a <i>decl-specifier-seq</i> appertains to the type determined by the <i>decl-specifier-seq</i> ." However, the rule for <i>decl-specifier-seq</i> in the grammar is recursive, and the intent is for the <i>attribute-specifier</i> to appertain to the top <i>decl-specifier-seq</i> , not the one in which the <i>attribute-specifier</i> directly appears. | Change the wording to indicate that the complete or outermost <i>decl-specifier-seq</i> is intended.   | ACCEPTED   |
| GB 29           | 7.1.5                                       |  | Te                           | A constexpr function is not permitted to return via an exception. This should be recognised, and a function declared 'constexpr' without an explicit exception specification should be treated as if declared 'noexcept(true)' rather than the usual 'noexcept(false)'. For a function template declared constexpr without an explicit  | Give constexpr functions an implicit non-throwing exception specification.   | ACCEPTED with MODIFICATIONS<br><br>The premise is not correct: an exception is forbidden only when a constexpr   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2  | (3)   | 4                                       | 5   | (6)  | (7)  |
|-----------------|--|---|---|---|--|--|
| MB <sup>1</sup> | Clause No./<br>Subclause No./<br>Annex<br>(e.g. 3.1) | Paragraph/<br>Figure/Table/<br>Note<br>(e.g. Table 1) | Type<br>of<br>com-<br>ment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations<br>on each comment submitted  |
|                 |  |   |   | exception specification, it should be considered 'noexcept(true)' if and only if the constexpr keyword is respected on a given instantiation.   |  | function is invoked in a context that requires a constant expression. Used as an ordinary function, it can throw. Instead of changing the default exception specification, the result of the noexcept operator was changed, based on whether the invocation of the function is a constant expression or not. |
| US<br>40        | 7.1.6.2  | 4   | te                                      | The description of decltype does not specify whether the type of a parameter is the declared type or the type as adjusted in 8.3.5¶5:<br><br>auto f(int a[])->decltype(a);<br>// ill-formed or int*?<br>auto g(const int i)->decltype(i);<br>// int or const int?                       | Clarify the wording to indicate that the type of a parameter is after the array- and function-to-pointer decay but before the removal of cv-qualification. | ACCEPTED   |
| DE<br>9         | 7.1.6.2  | p4  | te                                      | decltype applied to a function call expression requires a complete type (5.2.2 paragraph 3 and 3.2 paragraph 4), even though decltype's result might be used in a way that does not actually require a complete type. This might cause undesired and excessive template instantiations. | When immediately applying decltype, do not require a complete type, for example for the return type of a function call.                                    | ACCEPTED with<br>MODIFICATIONS<br><br>The change was applied only to the return type of a function call operand of decltype.<br><br>See paper <a href="#">N3276</a>  |
| US<br>41        | 7.1.6.3  | 2   | te                                      | The current wording disallows use of <i>typedef-names</i> in <i>elaborated-type-specifiers</i> . This prohibition should also apply to template aliases:<br><br>struct A { };<br>template<typename T> using X = A;<br>struct X<int>* p2; // ill-formed                                  | Add the necessary wording to prohibit a specialization of a template alias in an <i>elaborated-type-specifier</i> .  | ACCEPTED   |
| US              | 7.1.6.4  |   | te                                      | The overloaded meaning of the auto specifier is confusing   | Choose another keyword to indicate a late-   | REJECTED   |

<sup>1</sup> MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

<sup>2</sup> Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                   |
| 42              |   |  |                              | and prevents possible future language enhancements.  | specified return type. The identifiers lateret and postret have no Google code search hits. The identifiers late, latetype, and posttype have 30-40 hits. | There was no consensus to make the suggested change.                 |
| US 43           | 7.6   | 4  | ed                           | The contexts in which an <i>attribute-specifier</i> can appear include statements, described in clause 6, but the cross-references to clauses describing those contexts do not include clause 6.   | Add clause 6 to the list of cross-references in the first sentence.   | ACCEPTED   |
| GB 30           | 7.6.1                                       | 6  | Te                           | Making the use of <code>[]</code> illegal except where introducing an attribute specifier is just reintroducing the problem we had with <code>&gt;&gt;</code> for closing nested templates, albeit in a minor and less common form.<br>As Jason Merrill commented in c++std-core-16046, there is no ambiguity in practice because code couldn't actually be well-formed under interpretation as both an attribute specifier and a lambda introducer. A small amount of lookahead would be required to differentiate the cases, but this should not be a problem.<br>This restriction also means that lambdas in macros must be enclosed in parentheses to avoid accidental interpretation as an illegal attribute specifier if used as an array index. | Delete 7.6.1 paragraph 6.   | REJECTED<br><br>There was no consensus to make the suggested change. |
| JP 23           | 7.6.1                                       | 4  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.   | Change "(clause 7, clause 8)" to "(Clause 7, Clause 8)".  | ACCEPTED   |
| GB 31           | 7.6.2                                       |  | Te                           | After reviewing the case for attributes, wg14 has opted not to adopt this feature, and is instead using keywords for the few important cases identified in the attributes proposal.  | Revert the changes in the initial alignment proposal that changed the 'alignas' keyword into an attribute.  | ACCEPTED<br>See paper <a href="#">N3190</a>                          |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                   |
|                 |   |  |                              | For compatibility with C, the change of the 'alignas' keyword to the '[[align]]' attribute should be reversed.   |   |  |
| GB 32           | 7.6.3                                       |  | Ge                           | C has rejected the notion of attributes, and introduced the noreturn facility as a keyword. To continue writing clean, portable code we should replace the [[noreturn]] attribute with a 'noreturn' keyword, following the usual convention that while C obfuscates new keywords with _Capital and adds a macro to map to the comfortable spelling, C++ simply adopts the all-lowercase spelling.  | Replace the [[noreturn]] attribute with a new keyword, 'noreturn', with identical semantics. Note that this implies the keyword will not be something that a function can be overloaded upon. | REJECTED<br><br>There was no consensus to make the suggested change. |
| US 44           | 7.6.5                                       |  | te                           | <p>Even if attributes continue to be standardized over continued objections from both of the two vendors who are cited as the principal prior art, we can live with them with the exception of the virtual override controls. This result is just awful, as already shown in the example in 7.6.5 (excerpted):</p> <pre>class D [[base_check]] : public B {     void sone_func [[override]] ();     virtual void h [[ hiding]] (char*); };</pre> <p>Here we have six keywords (not counting void and char) — three normal keywords, and three [[decorated]] keywords. There has already been public ridicule of C++0x about this ugliness. This is just a poor language design, even in the face of backward compatibility concerns (e.g., that some existing code may already use those words as identifiers) because those concerns have already been resolved in other ways in existing practice (see below).</p> <p>More importantly, this is exactly the abuse of attributes as disguised keywords that was objected to and was explicitly promised not to happen in order to get this proposal passed. The use of attributes for the virtual control keywords is the most egregious abuse of the attribute syntax, and at least that use of attributes must be</p> | Change the syntax for virtual override control to not use attributes.   | ACCEPTED<br><br>See paper <a href="#">N3272</a>                      |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>fixed by replacing them with non-attribute syntax.</p> <p>These virtual override controls are language features, not annotations.</p> <p>It is possible to have nice names and no conflicts with existing code by using contextual keywords, such as recognizing the word as having the special meaning when it appears in a grammar position where no user identifier can appear, as demonstrated in C++/CLI which has five years of actual field experience with a large number of customers (and exactly no name conflict or programmer confusion problems reported in the field during the five years this has been available):</p> <pre>class D : public B {     void some_func() override; // same meaning as     [[override]] – explicit override     virtual void h (char*) new; // same meaning as     [[hiding]] – a new function, not an override }; int override = 42; // ok, override is not a reserved keyword</pre> <p>The above forms are implementable, have been implemented, have years of practical field experience, and work. Developers love them. Whether the answer is to follow this existing practice or something else, there needs to be a more natural replacement for the currently [[attributed]] keywords for virtual override control which is an ugly novelty that has no field experience and that developers have already ridiculed.</p> |  |  |
| US 45           | 7.6.5                                       | 6  | ed                           | <p>The example includes a line reading</p> <pre>class D [[base_check]] : public B {</pre> <p>However, the current syntax in 9.1 places the <i>attribute-</i></p>   | <p>Change the example to read</p> <pre>class [[base_check]] D : public B {</pre> | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <i>specifier</i> before the class name.   |   |  |
| CH 8            | 8.1   | p1 (syntax)                                  | ed                           | 'nopt <sup>r</sup> -abstract-declarator:' rule misses 'opt' subscript from the constant expression within the array brackets. This seems to be an editorial oversight   | change: " <i>nopt<sup>r</sup>-abstract-declarator<sub>opt</sub> [ constant-expression ] attribute-specifier<sub>opt</sub></i> " to " <i>nopt<sup>r</sup>-abstract-declarator<sub>opt</sub> [ constant-expression<sub>opt</sub> ] attribute-specifier<sub>opt</sub></i> "  | ACCEPTED<br>See paper <a href="#">N3262</a>  |
| US 46           | 8.3.2<br>20.7.6.2                           | all<br>Table 49                              | te                           | There is no way to create a prvalue of array type, so there ought to be no way create a (nonsensical) rvalue reference to array type.   | In [dcl.ref]/2, disallow declarations of T (&&A)].<br><br>In [dec.ref]/6 add a sentence: <u>If a typedef, a type template-parameter, or a decltype-specifier denotes a type A that is an array type (of known or unknown size), an attempt to create the type "rvalue reference to cv A" creates the type A&amp;.</u><br><br>In [meta.trans.ref]/Table 49 change the third row as follows:<br><br><u>If T names an array type, then the member typedef type shall name T&amp;, otherwise if T names an object or function type...</u> | REJECTED<br><br>It is possible to create an array prvalue: an array member of a class prvalue is an array prvalue.   |
| GB 33           | 8.3.5                                       | 5  | Ed                           | The register keyword is deprecated, so does not make for a good example. Suggest substituting the new storage class specifier, 'thread_local', instead.   | Use 'thread_local' in place of 'register' in the following sentence: "[ Example: register char* becomes char* —end example ]"   | ACCEPTED with MODIFICATIONS<br><br>thread_local cannot be used for a parameter. However, storage class specifiers do not affect the type, so mentioning storage class specifiers was incorrect and has been removed. |
| US 47           | 8.4.2                                       |  | te                           | 8.4.2 [dcl.fct.def.default]/4 says: "A special member function is user-provided if it is user-declared and not explicitly defaulted on its first declaration. A user-provided explicitly-defaulted function is..." The second sentence here should say "A user-declared explicitly-defaulted function is...". | Change<br><br>"A user-provided explicitly-defaulted function is..."<br><br>to<br><br>"A user-declared explicitly-defaulted function is..."  | ACCEPTED with MODIFICATIONS<br>See CWG 1134  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
| GB 34           | 8.4.2                                       | p2   | Ed                           | It is confusing when a normative paragraph starts with a note. The note starting this paragraph, with its reference to 'this' relating to the previous paragraph and not the content that follows, should be moved into the first paragraph, or the rest of this paragraph after the note should start a new numbered paragraph.                            | The note starting this paragraph should be moved into the first paragraph, or the rest of this paragraph after the note should start a new numbered paragraph. | ACCEPTED   |
| FI 1            | 8.4.2 [dcl.fct.def.default]                 | Paragraph 2                                  | te                           | It should be allowed to explicitly default a non-public special member function on its first declaration. It is very likely that users will want to default protected/private constructors and copy constructors without having to write such defaulting outside the class.   | Strike the "it shall be public" bullet.  | ACCEPTED   |
| FI 2            | 8.4.2 [dcl.fct.def.default]                 | Paragraph 2                                  | te                           | It should be allowed to explicitly default an explicit special member function on its first declaration. It is very likely that users will want to default explicit copy constructors without having to write such defaulting outside of the class.   | Strike the "it shall not be explicit" bullet.<br>See Appendix 1 - Additional Details   | ACCEPTED   |
| FI 3            | 8.4.2 [dcl.fct.def.default]                 | Paragraph 2                                  | te                           | It should be allowed to explicitly default a virtual special member function on its first declaration. It is very likely that users will want to default virtual copy assignment operators and destructors without having to write such defaulting outside of the class.  | Strike the "it shall not be virtual" bullet.<br>See Appendix 1 - Additional Details  | ACCEPTED   |
| GB 35           | 8.5.1                                       | 7  | Te                           | With the removal of the deprecated string-literal-to-non-const-char* conversion, member 'b' of struct S should be declared as a 'const' char*.  | Fix struct S as: "struct S { int a; const char* b; int c; };"  | ACCEPTED   |
| JP 71           | 8.5.1                                       | 7  | E                            | "char*" should be "const char*"<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed. | Change to:<br>struct S { int a; const char* b; int c; };<br>S ss = { 1, "asdf" };  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
| JP 72           | 8.5.1                                       | 15   | E                            | "char*" should be "const char *".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed.   | Change to:<br>union u { int a; const char* b; };<br>u a = { 1 };<br>u b = a;<br>u c = 1; // error<br>u d = { 0, "asdf" }; // error<br>u e = { "asdf" }; // error | ACCEPTED   |
| GB 36           | 8.5.1                                       | 17   | Ed                           | The 'b' member of union u should be declared const char * to better illustrate the expected cause of failures.  | Update union u as: "union u { int a; const char* b; };"  | ACCEPTED   |
| US 48           | 8.5.3                                       | 5  | te                           | The rule "...or the reference shall be an rvalue reference and the initializer expression shall be an rvalue or have a function type" is stated in terms of the rvalue-ness of the expression rather than the eventual target of the reference; this leads to some undesirable results, such as<br><br>struct A { };<br>struct B {<br>operator A&();<br>};<br><br>A&& aref = B(); // binds to lvalue<br><br>(c++std-core-16305) | Correct the formulation to deal with the rvalue-ness of the initializer after conversion to the appropriate type.  | ACCEPTED   |
| US 49           | 8.5.3                                       | 5  | te                           | The FCD does not specify direct binding for this example:<br><br>int i;<br>int main()<br>{<br>int&& ir = static_cast<int&&>(i);<br>ir = 42;<br>return (i != 42);<br>}<br><br>(c++std-core-16181)  | See Appendix 1 - Additional Details  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
| GB 37           | 8.5.3                                       |  | Te                           | <p>It seems that lvalues of any sort don't bind to non-const rvalue ref args, even if an intermediate temporary would be created.</p> <p>See the discussion at <a href="http://stackoverflow.com/questions/2748866/c0x-rvalue-references-and-temporaries">http://stackoverflow.com/questions/2748866/c0x-rvalue-references-and-temporaries</a>.</p> <p>I'll summarise that here: Assume that std::vector has push_back overloads declared as follows, with SFINAE omitted for clarity:</p> <pre>void push_back(const T &amp;); //Copies the const T &amp; into the vector's storage  void push_back(T &amp;&amp;); //Moves the T &amp;&amp; into the vector's storage</pre> <p>Then this code appears to behave as commented, as of N3090:</p> <pre>const char * cchar = "Hello, world"; std::vector&lt;std::string&gt; v; v.push_back(cchar); //makes a temporary string, copies the string into vector storage using push_back(const T&amp;)  v.push_back(std::string(cchar)); //makes a temporary string, moves the string into vector storage using push_back(T&amp;&amp;)  v.push_back(std::move(cchar)); //makes a temporary string, moves the string into the vector using push_back(T&amp;&amp;)</pre> <p>Johannes Schaub (litb) convincingly argued that the reason for this is clause 8.5.3/5 describing reference</p> | <p>Possible wording: amend the second list item in 8.5.3/5:</p> <p>Otherwise, the reference shall be an lvalue reference to a non-volatile const type (i.e., cv1 shall be const), or the reference shall be an rvalue reference [deleted the rest of the sentence].</p> <p>[The last example would also need to be deleted.]</p> | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)                       | (7)  |
|-----------------|---|--|------------------------------|--|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>binding - it allows direct binding of lvalues, bindings that require conversion sequences to const lvalue refs, and rvalues to rvalue refs. But it doesn't allow for lvalues to ever bind to rvalue refs, even if an intermediate temporary would otherwise need to be created.</p> <p>This isn't what I (as a user of <code>std::vector</code>) expect to happen. I expect all of these <code>push_back</code> calls to do the same thing, namely, select the <code>push_back(T&amp;&amp;)</code> overload, create a temporary string object from 'cchar', bind the temporary string to the argument, and hence move (not copy) the temporary string into vector's storage.</p> <p>It seems particularly strange that <code>v.push_back(cchar)</code> "requires" an extra copy, but <code>v.push_back(std::move(cchar))</code> does not. It almost seems like indicating that 'cchar' is potentially-movable (by casting it to an rvalue ref using <code>std::move</code>) allows moving from a completely different object - the temporary string.</p> <p>I suggest extending the rules for initializing const lvalue refs via implicit conversions (8.5.3/5), to also apply to rvalue refs.</p> <p>This also raises an additional question of whether lvalues of <code>_copyable_</code> types should be copied into a temporary object so that they may bind to an rvalue ref. Allowing this would not affect const T&amp;/T&amp;&amp; overload pairs. But it could be potentially useful when writing functions that wish to accept a number of rvalue refs to copyable-but-not-movable types (such as all C++03 classes with user-defined copy constructors), or when writing functions that "take apart" a number their arguments in a way that is different from a straightforward move (perhaps some tree operations would want to do this).</p> <p>Conversely, it might seem odd that declarations such as:</p> |                           |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)   | 4                            | 5  | (6)  | (7)   |
|-----------------|---|---|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1)          | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |   |                              | <p>string &amp;&amp; s = string_lvalue;<br/>                     string &amp;&amp; s = string_rvalue_ref_variable;<br/>                     //mistake, std::move(string_rvalue_ref_variable) was intended</p> <p>...would both</p> <p>silently copy their arguments and bind to the copy, instead of being invalid as they are now.</p> <p>We believe this is core issue 953</p>   |  |   |
| DE 10           | 8.5.3                                       |   | te                           | <p>Reference binding rules for rvalue references should consider temporaries generated from lvalues by implicit conversions. Consider to postpone the lvalue/rvalue analysis of an expression to after the implicit conversion chain has been deduced. Example:</p> <pre>void f(std::string&amp;&amp;); void g() {     f(std::string("hello")); // #1, ok     f("hello");             // #2, error, // but should be the same as #1 }</pre>  |  | ACCEPTED  |
| US 50           | 9   | 9   | te                           | <p>the class "struct A { const int i;}" was a POD in C++98, but is not a POD under the FCD rules because it does not have a trivial default constructor; I believe that C++0x POD was intended to be a superset of C++98 POD.</p>  | change POD to be standard layout and trivially copyable?   | ACCEPTED  |
| FI 16           | 9<br>9<br>9<br>3.9                          | 5 trivial<br>6 std-layout<br>9 POD<br>10 literal type | ge                           | <p>There are definitions for these types in the text, yet it is left unclear what use these classifications have. The types are very close to each other, which makes them confusing. If the reader must rely on external references, then these references should be specified (which is undesirable, or even disallowed by ISO(?)). As it stands, there is an example for using standard-layout classes (with other programming languages). There are also uses specified for literal types. One can imagine many uses for these four/five types, so it is important to have a clear</p> | <p>It is necessary to have detailed information on the expected uses of standard-layout, trivial, trivially copyable, literal and POD types.</p> | <p>REJECTED</p> <p>The current wording is clear enough. The Standard is correct as written.</p> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                        |
|                 |   |  |                              | specification of the intent as to where each of these types is expected to be used.   |  |   |
| JP 81           | 9   | 9  | E                            | Missing description of acronym "POD", which existed in C++03: The acronym POD stands for "plain old data."  | Add "The acronym POD stands for "plain old data." as an annotation.  | ACCEPTED  |
| US 51           | 9.2 [class.mem]                             |  | ge                           | Non-static data member initializers have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N2756 from the working paper prior the FDIS.  | REJECTED<br>C++/CLI has a very similar feature that has been implemented. |
| US 52           | 9.3   | 7  | te                           | The current wording allows friend declarations to name member functions "after their class has been defined." This appears to prohibit a friend declaration in a nested class defined inside its containing class that names a member function of the containing class, because the containing class is still considered to be incomplete at that point.      | Change the wording to allow a friend declaration of a "previously-declared member function."   | ACCEPTED  |
| US 53           | 9.3.1 [class.mfct.no n-static]              |  | ge                           | Move semantics for *this have no implementation experience.   | Either demonstrate a complete implementation of this feature or remove N1821 from the working paper prior the FDIS.  | REJECTED. The Committee found no consensus for adopting this change.      |
| JP 73           | 9.3.1                                       | 3  | E                            | "char*" should be "const char *".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed. | Change to:<br>struct tnode {<br>char tword[20];<br>int count;<br>tnode *left;<br>tnode *right;<br>void set(const char*, tnode* l, tnode* r);<br>};<br>void tnode::set(const char* w, tnode* l, tnode* r) {<br>count = strlen(w)+1;<br>if (sizeof(tword)<=count)<br>perror("tnode string too long");<br>strcpy(tword,w);<br>left = l;<br>right = r; | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted  |
|                 |   |  |                              |  | <pre> } void f(tnode n1, tnode n2) { n1.set("abc",&amp;n2,0); n2.set("def",0,0); } </pre>   |   |
| US 54           | 9.5 [class.union]                           |  | ge                           | Unrestricted unions have no implementation experience.   | Either demonstrate a complete implementation of this feature or remove N2544 from the working paper prior the FDIS.   | ACCEPTED<br><br>The feature has been implemented,no change is needed.   |
| JP 74           | 9.5   | 6  | E                            | "char*" should be "const char **".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed.   | Change to:<br>void f() {<br>union { int a; const char* p; };<br>a = 1;<br>p = "Jennifer";<br>}  | ACCEPTED  |
| GB 38           | 9.6   |  | Te                           | The signedness of bit-fields is the only time when 'signed int' is any different to just 'int'.<br><br>In C it is possible to remember whether a typedef uses 'signed' but in C++ it doesn't make sense and will result in ODR violations if A<long> and A<signed long> are not exactly equivalent.<br><br>This also causes portability problems because it is not specified whether typedefs such as int32_t are defined with 'signed' so using the <stdint> types in bitfields is problematic.<br><br>It is common to want to guarantee a bit-field has a minimum number of bits, for which the <stdint> types are | 'signed int' should always be equivalent to 'int' in all contexts.<br><br>A possible alternative would be to specify that signed types in <stdint> are declared with 'signed' so that using them for bit-fields has predictable results, but this doesn't address the ODR issue with A<long> and A<signed long> | REJECTED<br><br>Resolving this question was not deemed essential for this revision of the Standard, but core language issue 675 remains open for possible consideration in a future revision. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              | useful, except that the signedness of a bit-field using int32_t might depend on both unspecified and implementation-defined behaviour.   |  |  |
| US 55           | 10.3  | Paragraph 5                                  | te                           | <p>The following code not only does not compile on the compilers I've tested, but cannot be fixed through any combinations of forward references</p> <pre> class B { public:     virtual B *f() = 0; }; class D1 : public B { public:     virtual D2 *f(); }; class D2 : public B { public:     virtual D1 *f(); };                     </pre> | <p>In the core mailing list, Daniel Krugler points out that the current wording is ambiguous as to whether this is legal (although an editorial example suggests otherwise), and observes that it should be OK as long as D2 is complete at the point of definition of D1::f. The standard should resolve the ambiguity by saying that D2 only needs to be complete at the point of definition of D1::f.</p> <p>The core mailing list message text is below:<br/>                     I would be happy, if the standard would just allow this, but IMO the current wording seems to be readable in different ways (e.g. Comeau rejects the code). I think the reason is that [class.virtual]/5 just says:</p> <p>"The return type of an overriding function shall be either identical to the return type of the overridden function or covariant with the classes of the functions.[.]"</p> <p>This restriction is IMO only necessary for the <b>*definition*</b> of D::f.<br/>                     Secondly p. 6 says:</p> <p>"If the return type of D::f differs from the return type of B::f, the class type in the return type of D::f shall be complete at the point of</p> | <p>REJECTED</p> <p>It is necessary to be able to calculate the offsets applied to covariant return values when only the class definition is available.</p> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              |   | declaration of D::f or shall be the class type D[..]"<br><br>and shows the following example that explicitly forbids that (simplified):<br><br><pre>struct Base {     virtual B* vf5(); };  class A;  struct Derived : public Base {     A* vf5(); // error: returns pointer to incomplete class };</pre>   |  |
| US 56           | 11.3 [class.access.decl]                    |  | te                           | Access declarations were deprecated in the 1998 standard and have no benefits over using declarations.  | Remove access declarations from the working paper.  | ACCEPTED   |
| FI 4            | 12.1 [class.ctor]                           | Paragraph 5                                  | te                           | What effect does defaulting have on triviality? Related to FI 1, non-public special members defaulted on their first declaration should retain triviality, because they shouldn't be considered user-provided. Related to FI 3, defaulted member functions that are virtual should not be considered trivial, but there's no reason why non-virtuals could not be.<br><br>Furthermore, a class with a non-public explicitly-defaulted constructor isn't ever trivially constructible under the current rules. If such a class is used as a subobject, the constructor of the aggregating class should be trivial if it can access the non-public explicitly defaulted constructor of a subobject. | Change the triviality rules so that a class can have a trivial default constructor if the class has access to the default constructors of its subobjects and the default constructors of the subobjects are explicitly defaulted on first declaration, even if said defaulted constructors are non-public.<br><br>See Appendix 1 - Additional Details | ACCEPTED   |
| FI 15           | 12.3.1                                      | 2  | ge                           | 12.3.1. 2: "A default constructor may be an explicit constructor; such a constructor will be used to perform default-initialization or value-initialization (8.5)."   | The difference between a no argument default constructor and an explicit no argument default constructor should be explained in the standard.   | REJECTED<br><br>The current wording is clear       |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <p>12.3. 1 also says that an explicit ctor is different from a non-explicit ctor in that it is only invoked when direct-initialization (T a(1); T a{1}. presumably also T a;) or casting is used.</p> <p>What are the scenarios for the default ctor where explicit actually matters? Temporaries, arrays, ???</p> <p>When, if ever, is an explicit default ctor different from a non-explicit ctor?</p>  | If there is no difference, this should be explicitly specified.   | enough.  |
| US 57           | 12.3.2 [class.conv.fct ]                    |  | ge                           | Explicit conversion operators have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N2437 from the working paper prior the FDIS.   | ACCEPTED<br>The feature has been implemented.  |
| GB 39           | 12.4  | 4  | Te                           | Contradiction between the note and normative language describing when defaulted function definitions might have an exception specification. (See 8.4.2p2 for requirement to provide the exception specification)  | Either strike the second sentence of this note, or update it to reflect under which conditions a defaulted definition might have an exception specification.                                    | ACCEPTED   |
| GB 40           | 12.4  |  | Te                           | It is very difficult to write correct programs that do not call 'terminate' in the presence of destructors that can throw exceptions, and this practice has long been discouraged. Many implicitly declared destructors already carry noexcept declarations (mostly types with trivial destructors) and it is anticipated it will become common practice to want a user-declared destructor to be declared noexcept. This becomes important evaluating the noexcept operator, where any of the unevaluated sub-expressions may produce a temporary object. As this is expected to be the overwhelmingly common case, a user-declared destructor that does not supply an exception specification should be considered as if declared noexcept(true) rather than noexcept(false), the default for every other function. | a user-declared destructor that does not supply an exception specification should be considered as if declared noexcept(true) rather than noexcept(false), the default for every other function | ACCEPTED with MODIFICATIONS<br>The rule that was adopted makes destructors noexcept if all the base and member destructors are.<br>See paper <a href="#">N3204</a> |
| CH              | 12.4 and 15.4                               |  | te                           | Destructors should generally not throw exceptions. Consider giving an explicit rule for this.   | Add in 12.4 or 15.4 a paragraph to the effect that all destructors not having an exception  | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted  |
| 9               |   |  |                              |  | soecification are considered noexcept(true).  | The rule that was adopted makes destructors noexcept if all the base and member destructors are.<br><br>See paper <a href="#">N3204</a> |
| US 58           | 12.5  | foot 117                                     | ed                           | Missing comma in "is not virtual the size might".  | Add the comma.  | REJECTED<br><br>Because of the comma later in the sentence, a comma here would be wrong.  |
| US 59           | 12.6.2 [class.base.initt]                   |  | ge                           | Delegating constructors have no implementation experience.   | Either demonstrate a complete implementation of this feature or remove N1986 from the working paper prior the FDIS.   | ACCEPTED<br><br>The feature has been implemented, no changes are needed.  |
| US 60           | 12.8 [class.copy]                           |  | ge                           | Implicitly-defined move constructors and move assignment operators have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N3053 from the working paper prior the FDIS.   | ACCEPTED<br><br>The feature has been implemented, no changes are needed.  |
| DE 11           | 12.8  |  | te                           | It is unclear whether copy elision can or cannot apply to a case like C f(C c) { return c; }, i.e. where a parameter of class type is returned. Furthermore, if copy elision cannot apply there, it should still be possible to move (instead of copy) the return value.   | Amend paragraph 34 to explicitly exclude function parameters from copy elision. Amend paragraph 35 to include function parameters as eligible for move-construction.  | ACCEPTED  |
| FI 5            | 12.8 [class.copy]                           | Paragraph 13, paragraph 27                   | te                           | Same as FI 4, the parts involving copy constructors and copy assignment operators.<br><br>A class with a non-public explicitly-defaulted copy constructor isn't ever trivially copyable under the current rules. If such a class is used as a subobject, the copy constructor of the aggregating class should be trivial if it can access the non-public explicitly defaulted copy constructor of a subobject. | Change the triviality rules so that a class can have a trivial copy constructor if the class has access to the copy constructors of its subobjects and the copy constructors of the subobjects are explicitly defaulted on first declaration, even if said defaulted copy constructors are non-public.<br><br>See Appendix 1 - Additional Details | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted      |
| GB 41           | 12.8  | 15, 29                                       | Te                           | Contradiction between the notes claiming that defaulted definitions to not have exception specifications, and the normative language in 8.4.2 which declares that they might.   | Either strike the second sentence of each note, or update it to reflect under which conditions a defaulted definition might have an exception specification. | ACCEPTED  |
| US 61           | 12.8;20.2.5                                 | 17,31;table 42                               | ed                           | static_cast is broken across two lines  | do not hyphenate static_cast   | ACCEPTED  |
| US 62           | 12.8  | 16   | te                           | The new wording describing generated copy constructors does not describe the initialization of members of reference type. (Core issue 1051 in N3083.)   | Add the required description.  | ACCEPTED  |
| US 63           | 12.8  | 16-18  | te                           | The new wording specifies the behavior of an implicitly-defined copy constructor for a non-union class (§16), an implicitly-defined move constructor for a non-union class (§17), and an implicitly-defined copy constructor for a union (§18), but not an implicitly-defined move constructor for a union. (Core issue 1064 in N3083.)   | Add the required description.  | ACCEPTED  |
| US 64           | 12.8  | 28   | te                           | The current wording reads, "A copy/move assignment operator that is defaulted and not defined as deleted is <i>implicitly defined</i> when an object of its class type is assigned a value of its class type or a value of a class type derived from its class type or when it is explicitly defaulted after its first declaration." This sounds as if any assignment to a class object, regardless of whether it is a copy or a move assignment, defines both the copy and move operators. Presumably an assignment should only define the assignment operator chosen by overload resolution for the operation. (Compare the corresponding wording in §14 for the copy/move constructors: "... <i>implicitly defined</i> if it is used to initialize an object of its class type..." (Core issue 1066 in N3083.) | Clarify the wording so that only the operator needed for the operation is implicitly defined.  | ACCEPTED  |
| US 65           | 12.9 [class.inhctor]                        |  | ge                           | Inheriting constructors have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N2540 from the working paper prior the FDIS.  | REJECTED<br><br>There was no consensus for this change. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2  | (3)   | 4                               | 5  | (6)  | (7)   |
|-----------------|--|---|---------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./<br>Subclause No./<br>Annex<br>(e.g. 3.1) | Paragraph/<br>Figure/Table/<br>Note<br>(e.g. Table 1) | Type of<br>comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations<br>on each comment submitted   |
| JP<br>65        | 13.1   | 3   | E                               | In some code examples, ellipsis(...) is used in ill-formed. In these cases, "..." represents omission of some codes like this:<br>class A { /* ... */ };<br>But in some cases, it is used without commented-out as below:<br>class A { ... };<br>It is an inconsistent usage. They all should be enclosed in a comment.  | Change to:<br>int f (int) { /* ... */ } // definition of f(int)<br>int f (clnt) { /* ... */ } // error: redefinition of f(int) | ACCEPTED  |
| US<br>66        | 13.3.1.7   | 1   | te                              | overload resolution should first look for a viable list constructor, then look for a non-list constructor if no list constructor is viable   | See Appendix 1 - Additional Details  | ACCEPTED<br>See paper <a href="#">N3262</a>   |
| US<br>67        | 13.3.2   | 3   | ge                              | To determine whether there is an ICS, 13.3.2 uses 13.3.3.1 instead of just saying "there is an ICS if-and-only-if a copy init would be well-formed." Apparently this is desired, but to a casual reader or an implementor reading these rules for the first time for a new implementation, it's not clear <i>why</i> that's desired.                           | Insert a note or annex explaining why 13.3.2 does things as it does.   | ACCEPTED with<br>MODIFICATIONS<br><br>The resolution for core language issue 1152 fixes a related problem, but a rationale was not deemed necessary at this time. |
| JP<br>75        | 13.2   | 1   | E                               | "char*" should be "const char **".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed. | Change to:<br>struct B {<br>int f(int);<br>};<br>struct D : B {<br>int f(const char*);<br>};                                   | ACCEPTED  |
| JP<br>76        | 13.2   | 2   | E                               | "char*" should be "const char **".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)   | Change to:<br>void f(const char*);<br>void g() {<br>extern void f(int);<br>f("asdf"); // error: f(int) hides f(const char*)    | ACCEPTED  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed.   | // so there is no f(const char*) in this scope<br>}   |  |
| JP 77           | 13.3.1.2                                    | 1  | E                            | "char*" should be "const char *".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed.  | Change to:<br>struct String {<br>String (const String&);<br>String (const char*);<br>operator char* ();<br>};<br>String operator + (const String&, const String&);<br>void f(void) {<br>char* p= "one" + "two"; // ill-formed because neither<br>// operand has user-defined type<br>int l = 1 + 1; // Always evaluates to 2 even if<br>// user-defined types exist which<br>// would perform the operation.<br>} | ACCEPTED   |
| US 68           | 13.4  | 1  | te                           | Overload resolution within the operand of a unary & operator is done by selecting the function "whose type matches the target type required in the context." The criterion for determining whether the types match, however, is not defined. At least three possibilities suggest themselves:<br><br>1. The types are identical.<br><br>2. The source type can be implicitly converted to the target type.<br><br>3. The expression would be well-formed if the function under consideration were not overloaded.<br><br>This question arises for pointer-to-member types, where there is an implicit conversion from a pointer-to-base- | Specify the intended criterion for determining whether the types match.   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | member to a pointer-to-derived-member, as well as when the context is an explicit type conversion (which allows for static_cast a conversion from pointer-to-derived-member to a pointer-to-base-member and, in the reinterpret_cast interpretation of functional and old-style casts, essentially any conversion).  |  |  |
| JP 82           | 13.5.8                                      | 8  | E                            | Typo, "lteral" should be "literal".<br>// error: invalid lteral suffix identifier  | Correct typo.<br>// error: invalid literal suffix identifier   | ACCEPTED   |
| US 69           | 14.3.2                                      | para 1                                       | te                           | The standard permits the address of thread_local variable as a non-type template parameter. The addresses of these variables are not constant, however.  | Require static storage duration for non-type parameters.   | ACCEPTED   |
| DE 12           | 14.3.2                                      |  | te                           | Now that local classes can be used as template arguments, it seems odd that there are "external linkage" restrictions on non-type template parameters.   | Permit addresses of objects and functions with internal linkage as arguments for non-type template parameters.   | ACCEPTED   |
| JP 78           | 14.3.2                                      | 2  | E                            | "char*" should be "const char *".<br>If not corrected, type mismatch is another cause of error in the example below, which is not appropriate for an example here.<br>template<class T, char* p> class X {<br>X();<br>X(const char* q) { /... /}<br>};<br>X<int, "Studebaker"> x1; // error: string literal as template-argument<br>char p[] = "Vivisectionist";<br>X<int,p> x2; // OK | template<class T, const char* p> class X {   | ACCEPTED   |
| JP 83           | 14.3.2                                      | 2  | E                            | Constructors in template declaration are not essential for this example to explain string literal error use.<br><br>template<class T, char* p> class X {<br>X();<br>X(const char* q) { /* ... */}<br>};  | Delete two constructors in template declaration as follows.<br><br>template<class T, const char* p> class X {<br>/* ... */<br>};<br>X<int, "Studebaker"> x1; // error: string literal as | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | X<int, "Studebaker"> x1; // error: string literal as template-argument<br>char p[] = "Vivisectionist";<br>X<int,p> x2; // OK  | template-argument<br>char p[] = "Vivisectionist";<br>X<int,p> x2; // OK  |   |
| GB 42           | 14.5.3                                      | 1  | Ed                           | The name "eror" should probably be "error".   | Replace:<br>Tuple<0> eror;<br>With:<br>Tuple<0> error;   | ACCEPTED  |
| JP 24           | 14.5.3                                      | 4  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(10)" to "(Clause 10)".  | ACCEPTED  |
| US 70           | 14.5.6.2                                    | 2, 4   | te                           | 14.8.2.4¶3 specifies that the deduction used in partial ordering in a non-call context is based on the complete function type of the function templates. The wording in 14.5.6.2¶2 (and echoed in ¶4) reflects an earlier specification, however, saying that the deduction uses only "the function parameter types, or in the case of a conversion function the return type." This is a contradiction.                         | Update the wording in 14.5.6.2 to say only that deduction is performed as described in 14.8.2.4 and not to specify which types are used.   | ACCEPTED  |
| US 71           | unused                                      | unused                                       |                              | unused  |  | NA  |
| CA 7            | 14.5.6.2p3                                  | P3   | te                           | r country In FCD sub-clause 14.5.6.2 [temp.func.order] paragraph 3, we are told to synthesize, "for each type, non-type, or template template parameter... a unique type, value, or class template respectively."<br><br>These are then substituted for each occurrence of the respective parameter in the function type of the function  | This is a test of the interpretation of the resolution to Issue 214. In particular, we would like for the committee to spell out how properties of the synthetics produced for partial ordering are determined | REJECTED<br><br>Resolving this issue was not deemed essential for this revision of the Standard. Core language issue 1157 was opened, however, to |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)                       | (7)   |
|-----------------|---|--|------------------------------|---|---------------------------|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB | Secretariat observations on each comment submitted      |
|                 |   |  |                              | <p>template.</p> <p>It is not specified what the properties of said synthetics are, for example, members of a dependent type referred to in non-deduced contexts are not specified to exist, although the transformed function type would be invalid if they do not exist.</p> <p>We may assume, for example, that each synthetic will be given minimal properties such that the transformed function type is valid at the point of definition of the template.</p> <p><b>Example 1:</b></p> <pre>template &lt;typename T, typename U&gt; struct A; template &lt;typename T&gt; void foo(A&lt;T, typename T::u&gt; *) {} // #1 // synthetic T1 has member T1::u template &lt;typename T&gt; void foo(A&lt;T, typename T::u::v&gt; *) {} // #2 // synthetic T2 has member T2::u and member T2::u::v // T in #1 deduces to synthetic T2 in partial ordering; // deduced A for the parameter is A&lt;T2, T2::u&gt; *--this is not necessarily compatible // with A&lt;T2, T2::u::v&gt; * and it does not need to be. See Note 1. The effect is that // (in the call below) the compatibility of B::u and B::u::v is respected. // T in #2 cannot be successfully deduced in partial ordering from A&lt;T1, T1::u&gt; *; // invalid type T1::u::v will be formed when T1 is substituted into non-deduced contexts. struct B { struct u { typedef u v; }; }; int main() { foo((A&lt;B, B::u&gt; *)0); // calls #2 } *Note 1: Template argument deduction is an attempt to match a P and to a deduced A; however, template argument deduction is not specified to fail if the P and the deduced A are incompatible. This may occur in the presence of non-deduced contexts. Notwithstanding the parenthetical statement in [temp.deduct.partial] paragraph 9, template argument deduction may succeed in determining a template argument for every template parameter while producing a deduced A that is not</pre> |                           | allow for possible consideration for a future revision. |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                      |
|                 |   |  |                              | compatible with the corresponding P.<br><br><b>Example 2:</b><br><br><pre> template &lt;typename T, typename U, typename V&gt; struct A; template &lt;typename T&gt; void foo(A&lt;T, struct T::u, struct T::u&gt; *); #2.1 // synthetic T1 has member non-union class T1::u template &lt;typename T, typename U&gt; void foo(A&lt;T, U , U&gt; *); #2.2 // synthetic T2 and U2 has no required properties // T in #2.1 cannot be deduced in partial ordering from A&lt;T2, U2, U2&gt; *; // invalid types T2::u and T2::u::u will be formed when T2 is substituted in non- deduced contexts. // T and U in #2.2 deduces to, respectively, T1 and T1::u from A&lt;T1, T1::u, struct T1::u&gt; * unless // struct T1::u does not refer to the injected-class-name of the class T1::u (if that is possible). struct B { struct u { }; }; int main() { foo((A&lt;B, B::u, struct B::u&gt; * )0); // calls #2.1 } It is however unclear to what extent an implementation will have to go to determine these minimal properties.                     </pre> |   |   |
| US 72           | 14.5.7 [temp.alias]                         |  | ge                           | Template aliases have no implementation experience.  | Either demonstrate a complete implementation of this feature or remove N2258 from the working paper prior the FDIS. | REJECTED<br><br>The feature has been implemented, no changes ed needed. |
| US 73           | 14.5.7                                      | 1  | te                           | The current wording of 7.1.3¶2 requires that the <i>identifier</i> in an <i>alias-declaration</i> "...shall not appear in the <i>type-id</i> ." With template aliases, however, the name of the alias can be used indirectly:<br><br><pre> template&lt;typename T&gt; struct A; template&lt;typename T&gt; using B=typename A&lt;T&gt;::U; template&lt;typename T&gt; struct A {     typedef B&lt;T&gt; U; }; B&lt;short&gt; b;                     </pre>   | Add wording to indicate that such usages in template aliases are ill-formed.  | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | Here the instantiation of B<short> causes the instantiation of A<short>, and the typedef in A<short> ends up attempting to use B<short>, which is still in the process of being instantiated.   |  |   |
| US 74           | 14.5.7                                      | 1  | te                           | An <i>alias-declaration</i> allows a class or enumeration type to be defined in its <i>type-id</i> (7.1.6¶3). However, it's not clear that this is desirable when the <i>alias-declaration</i> is part of a template alias:<br><br>template<typename T> using A = struct { void f(T) {} };  | Either prohibit the definition of classes and enumerations in template aliases, or prohibit the use of template parameters in such definitions, or add an example illustrating this usage. | ACCEPTED<br><br>Definition of a class or enumeration is now prohibited in a template alias. |
| FI 10           | 14.5.7 [temp.alias]                         |  | te                           | Can template aliases be declared in class scope?  | Allow declaring template aliases in class scope, if not allowed by the current grammar.  | ACCEPTED  |
| FI 11           | 14.5.7 [temp.alias]                         |  | te                           | We have class templates and function templates, shouldn't we call template aliases alias templates for consistency?   | Change "template alias" -> "alias template" everywhere.  | ACCEPTED  |
| JP 25           | 14.5.7                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(clause 7)" to "(Clause 7)".   | ACCEPTED  |
| JP 26           | 14.6.2.1                                    | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only   | Change "(9)" to "(Clause 9)".  | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | a number Z in parentheses to confer Clause or Table number Z.   |   |  |
| GB 43           | 14.6.2.1p3                                  |  | Te                           | <p>If I'm not missing something, 14.6.2.1/3 says that in the definition of a class template or member of a class template, the injected-class-name refers to the current instantiation. The template name followed by the argument list enclosed in "&lt;.&gt;" also refers to the current instantiation, but only in the definition of a primary class template. That results in an odd situation:</p> <pre>template&lt;typename T&gt; struct A { typedef int type; void f(type); }; // here we are outside the definition of "A"</pre> <pre>template&lt;typename T&gt; void A::f(A::type) {} // OK: "A" is the injected-class-name</pre> <pre>template&lt;typename T&gt; void A::f(A&lt;T&gt;::type) {} // ill-formed: "A&lt;T&gt;" is not the injected-class-name. Needs "typename "!</pre> <p>If you would define the member-function within the primary class template, bullet 2 would apply:</p> <pre>template&lt;typename T&gt; struct A {   typedef int type; void f(A&lt;T&gt;::type) {} // OK: name of A followed by arguments enclosed in &lt;.&gt; };</pre> <p>I think that this state of affairs isn't any good.</p> <p>-&gt; Suggested solution: Change 14.6.2.1/1 bullet2 to apply also to members of the primary class template. The same for bullet4 talking about partial specializations. Since partial specializations are also class templates, i wonder whether one could also smelt together bullet2 and bullet4 and only talk about "class template".</p> | Updated 14.6.2.1 [temp.dep.type] p1 bullet 2: "— in the definition of a primary class template or a member of a class template, the name of the class template followed by the template argument list of the primary template (as described below) enclosed in <>," | ACCEPTED   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
| GB 44           | 14.6p5                                      |  | Te                           | <p>C++0x does not allow this code</p> <pre>template&lt;typename T&gt; struct id { typedef T type; }; template&lt;typename T&gt; void f() { int id&lt;T&gt;::type::*p = 0; } struct A { };  int main() { f&lt;A&gt;(); }</pre> <p>The reason is that it requires “typename” before “id&lt;T&gt;::type”, but “typename” is not allowed at that place by the syntax. Ultimately, current compilers accept this.</p>   | <p>Change 14.6/5 to</p> <p>A qualified name used as the name in a mem-initializer-id, a base-specifier, an elaborated-type-specifier or the nested-name-specifier of a pointer-to-member declarator is implicitly assumed to name a type, without the use of the typename keyword.</p> | ACCEPTED   |
| CA 6            | 14.6p5                                      | P5   | te                           | <p>Given the change in N1376=02-0034 to [temp.res], found in FCD in [temp.res] paragraph 5:</p> <p>A qualified name used as the name in a mem-initializer-id, a base-specifier, or an elaborated-type-specifier is implicitly assumed to name a type, without the use of the typename keyword</p> <p>the following appears to be well-formed, with templates foo() being distinct since any type T will produce an invalid type for the second parameter for at least one foo() when T is replaced within the non-deduced context:</p> <pre>template &lt;typename T&gt; bool *foo(T *, enum T::u_type * = 0) { return 0; } template &lt;typename T&gt; char *foo(T *, struct T::u_type * = 0) { return 0; } struct A { enum u_type { l }; }; int main(void) { foo((A*)0); }</pre> <p>In particular, while determining the signature (1.3.11 [defns.signature]) for the function templates foo(), an elaborated-type-specifier qualifies as part of the decl-</p> | <p>Please clarify as the following case appears to be expensive to implement with little functional value to the language.</p>   | <p>REJECTED</p> <p>The specification is as intended; compilers should handle cases like these.</p> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                      |
|                 |   |  |                              | <p>specifier-seq under 8.3.5 [dcl.fct] paragraph 5 in determining the type of a parameter in the parameter-type-list (absent additional wording).</p> <p>Also, the return type is included in the signature of a function template.</p> <p>A portion of the GCC 4.5.0 output:</p> <p>Internal compiler error: Error reporting routines re-entered. Please submit a full bug report, with preprocessed source if appropriate. See &lt; <a href="http://gcc.gnu.org/bugs.html">http://gcc.gnu.org/bugs.html</a> &gt; for instructions.</p> <p>Implementations do not appear to support this case and the ability to do so brings little value since type traits such as <code>is_enum</code> and <code>is_class</code> cannot be defined using this and equivalent functionality can be achieved using the aforementioned type traits.</p> <pre>template &lt;typename T&gt; struct MY_is_enum : std::false_type { };  template &lt;typename T&gt; struct MY_is_enum&lt;enum T&gt; : std::true_type { }; // ill-formed,</pre> <p>elaborated-type-specifier resolves to typedef-name</p> |   |   |
| US 75           | 14.7  |  | TE                           | As described in c++std-core-16425 and its followup messages, writing metaprograms is needlessly hard because specializing template members inside a class is (inadvertently?) not permitted. In addition, this surprising restriction makes C++ less simple and more arbitrary-seeming.  | Accept the code like that in c++std-core-16425, like Visual C++ does already                | REJECTED<br><br>There was no consensus for making the suggested change. |
| JP 79           | 14.7.1                                      | 10   | E                            | "char*" should be "const char *".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.   | Change to:<br>namespace N {<br>template<class T> class List {<br>public:<br>T* get();<br>}; | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                      |
|                 |   |  |                              | There are many code fragments depending on the removed rule. They are ill-formed.   | <pre> } template&lt;class K, class V&gt; class Map { public: N::List&lt;V&gt; lt; V get(K); }; void g(Map&lt;const char*,int&gt;&amp; m) { int i = m.get("Nicholas"); } a call of lt.get() from Map&lt;const char*,int&gt;::get() would place List&lt;int&gt;::get() in the namespace N rather than in the global namespace. —end example ] </pre>   |   |
| FI 9            | 14.7.3<br>[temp.expl.spec]                  | Paragraph 2                                  | te                           | Explicit specializations in class scope inside class templates should be allowed. It's weird, confusing and inconsistent that they can be declared/defined in some scopes but not in others.  | Allow explicit specialization of member templates inside class templates.  | REJECTED<br><br>There was no consensus for making the suggested change. |
| JP 80           | 14.8.1                                      | 5  | E                            | "char*" should be "const char *".<br>The special rule to convert character literal to pointer has been removed from "4.2 Array-to-pointer conversion [conv.array]".<br>char * p1 = "..."; // ill-formed.(removing const'ness)<br>char const *p2 = "..."; // well-formed.<br>There are many code fragments depending on the removed rule. They are ill-formed. | Change to:<br><pre> template&lt;class X, class Y, class Z&gt; X f(Y,Z); template&lt;class ... Args&gt; void f2(); void g() { f&lt;int,const char*,double&gt;("aa",3.0); f&lt;int,const char*&gt;("aa",3.0); // Z is deduced to be double f&lt;int&gt;("aa",3.0); // Y is deduced to be const char*, and // Z is deduced to be double f("aa",3.0); // error: X cannot be deduced f2&lt;char, short, int, long&gt;(); // OK } </pre> | ACCEPTED  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
| US 76           | 14.8.2                                      | para 9                                       | te                           | "extern template" prevents inlining functions not marked inline.  | Remove wording about "suppressing the implicit instantiation". See Appendix 1 - Additional Details | REJECTED<br><br>Section reference should be 14.7.2. The suggested change could result in a silent change of meaning between implementations if speculative instantiation for inlining is allowed. There was no consensus to adopt this change. |
| US 77           | 14.8.2.1                                    |  | te                           | Core Issue 1014 claims that calling f(const T&) and f(T&&) with a const int lvalue is ambiguous. It's unambiguous because the partial ordering rules consider f(const T&) to be more specialized than f(T&&), for the same reasons that they consider h(const T&) to be more specialized than h(T&).<br><br>However, calling z(T&) and z(T&&) with an int lvalue is ambiguous. Because z(T&) accepts a strict subset of the things that z(T&&) accepts, it seems that the partial ordering rules should be modified to consider z(T&) to be more specialized than z(T&&). There may be additional subtleties. | Modify the partial ordering rules to consider z(T&) to be more specialized than z(T&&).            | ACCEPTED   |
| US 78           | 15.2  | 2  | te                           | This paragraph says that "An object that is... partially destroyed will have destructors executed... for subobjects for which the principal constructor (12.6.2) has completed execution and the destructor has not yet begun execution." This would presumably apply to an example like<br><br>struct S { ~S(); } s[10];<br><br>If the destructor for s[5] throws an exception, elements 0-4 should still be destroyed. However, the wording specifically refers to "fully constructed base classes and  | Clarify the intent with respect to array elements and storage duration.                            | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | non-variant members," even though array elements are subobjects of the array (1.8¶2). This is presumably the effect of stack unwinding (¶1), which applies to "all automatic objects constructed since the try block was entered," but whether that should also be expected for arrays of static, thread, and dynamic storage duration is not clear.   |  |  |
| GB 45           | 15.3  | 16   | Te                           | The phrasing of this clause suggests all exception-declarations produce objects. There should be some additional wording to clarify that exception-declarations that declare references bind to the exception object by appropriate initialization, and *are* allowed to be references to abstract classes. Likewise, the elipsis form does not initialize any object or temporary.  | Distinguish between initializing objects, initializing references, and initializing nothing in the case of an elipsis. | ACCEPTED   |
| CA 5            | 15.3p8, 15.1p7                              | P8, p7                                       | te                           | <p>There is an issue with the definition of "currently handled exception" in 15.3 [except.handle] paragraph 8:</p> <p style="padding-left: 40px;">The exception with the most recently activated handler that is still active is called the <i>currently handled exception</i>.</p> <p>This wording implies that the currently handled exception may be changed by another thread. Thus, by 15.1 [except.throw] paragraph 7,</p> <p style="padding-left: 40px;">A throw-expression with no operand rethrows the currently handled exception (15.3).</p> <p>the following may throw an exception that is not of type int. try { throw 0; } catch (...) { throw; } Any solution should also specify what the currently handled exception will be for a thread that is spawned during the handling of an exception by its parent.</p> | Clarify and fix as suggested.  | ACCEPTED   |
| US 79           | 15.4  | ¶1   | te                           | Because C has no exception mechanism, functions having "C" language linkage should implicitly be declared  | <p>Insert a sentence such as:</p> <p>Any function declared to have "C" linkage shall be</p>                            | REJECTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                      |
|                 |   |  |                              | noexcept.   | treated as if declared noexcept even if the declaration lacks the noexcept keyword.  | There was no consensus for making the suggested change.                 |
| GB 46           | 15.4  |  | Te                           | It is not entirely clear that a function-try-block on a destructor will catch exceptions for a base or member destructor; whether such exceptions might be swallowed with a simple return statement rather than being rethrown; and whether such a clause might be entered multiple times if multiple bases/members throw, or if that is an automatic terminate call. | [Add words to 15.4 clarifying the problem cases.]  | ACCEPTED  |
| CH 10           | 15.4  | p9   | te                           | In case of incorrect program specification, the general rule is that the behaviour is undefined. This should be true for noexcept as well.  | Change the second bullet of p9 to "otherwise, the behaviour is undefined.  | REJECTED<br><br>There was no consensus for making the suggested change. |
| GB 47           | 15.5.1                                      | 1  | Te                           | 15.5.1:1 [except.terminate] lists the situations in which "exception handling must be abandoned for less subtle error handling techniques".<br><br>The list seems to omit some new situations added by other c++0x features.  | The list should be augmented with the following:<br>* when function std::nested_exception::rethrow_nested is called for an object that stores a null exception pointer.<br>* when execution of a function registered with std::at_quick_exit exits using an exception.<br><br>Also, after the list, add the following sentence:<br><br>Function std::terminate is also called by the implementation, when the destructor or a copy constructor of a class std::thread is called for the object that is joinable. | ACCEPTED  |
| GB 48           | 15.5.2                                      | 1-4  | Te                           | This subclause is dealing exclusively with dynamic exception specifications, and should clearly say so.   | Replace each italicised occurrence of 'exception-specification' with 'dynamic-exception-specification' in clause 15.5.2, [except.unexpected]   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
| GB 49           | 15.5.2                                      | all  | Ge                           | Dynamic exception specifications are deprecated, so clause 15.5.2 that describes how they work should move to Annex D.   | Move 15.5.2 [except.unexpected] to Annex D.  | REJECTED<br><br>The convention in the core clauses is to note the deprecation but leave the specification <i>in situ</i> . |
| CH 11           | 16.1  | p3   | ed                           | constant-expression as given by C++'s grammar allows far more syntactical constructs than those that are allowed in preprocessor context. The footnote 145 provides some hints on the limitations, but in our opinion it should be made more clear that the "constant-expression" allows less than a lookup of the corresponding grammar rule suggests | Add a note or extend footnote 145 with "Note that constant-expression is much more limited than the C++ grammar rule would suggest. See the following paragraphs how it is limited in the context of conditional inclusion."                       | REJECTED<br><br>The current wording is the same as in C++03, and hasn't caused confusion.                                  |
| CH 12           | 16.3.5                                      | p5   | ed                           | missing space between '~' and '5' in expansion   | line "f(2 * (2+(3,4)-0,1))   f(2 * (~5)) & f(2 * (0,1))^m(0,1);" should read "f(2 * (2+(3,4)-0,1))   f(2 * (~5)) & f(2 * (0,1))^m(0,1);"   | ACCEPTED   |
| CH 13           | 16.3.5                                      | p7   | ed                           | superfluous braces in source   | change "int j[] = { t(1,2,3), t(,4,5), t(6,{,}7), t(8,9,)," to "int j[] = { t(1,2,3), t(,4,5), t(6,,7), t(8,9,),"  | ACCEPTED   |
| CH 14           | 16.3.5                                      | p9   | ed                           | superfluous spaces after/before parentheses  | change<br>fprintf(stderr, "Flag" );<br>fprintf(stderr, "X = %d\n", x );<br>puts( "The first, second, and third items." );<br><br>to<br>fprintf(stderr, "Flag");<br>fprintf(stderr, "X = %d\n", x);<br>puts("The first, second, and third items."); | ACCEPTED   |
| DE 13           | 16.8  |  | te                           | Committee Draft comment DE 18 has only been partially addressed, and the record of response ignores the missing item, namely the absence of a macro  | Add the macro to the list of predefined macros in 16.8.  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <code>__STDCPP_STRICT_POINTER_SAFETY__</code> that indicates that a given implementation has strict pointer safety (see 3.7.4.3).   |   |  |
| US 80           | [library]<br>17 - 30                        |  | te                           | Consider applying <code>noexcept</code> to all of the <code>std::lib</code> .   |   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3155</a>   |
| GB 60           | 17-30                                       |  | Ge                           | Dyanamic exception specifications are deprecated; the library should recognise this by replacing all non-throwing exception specifications of the form 'throw()' with the 'noexcept' form.  | Replace all non-throwing exception specifications of the form 'throw()' with the 'noexcept' form.   | ACCEPTED with MODIFICATIONS<br><br>See documents N3148, N3150, N3195, N3155, N3156, N3199, and N3180.<br><br>See paper <a href="#">N3148</a> |
| GB 61           | 17-30                                       |  | Te                           | All library types should have non-throwing move constructors and move-assignment operators unless wrapping a type with a potentially throwing move-operation. When such a type is a class-template, these operations should have a conditional <code>noexcept</code> specification.<br>There are many other places where a <code>noexcept</code> specification may be considered, but the move operations are a special case that must be called out, to effectively support the <code>move_if_noexcept</code> function template. | Review every class and class template in the library. If <code>noexcept</code> move constructor/assignment operators can be implicitly declared, then they should be implicitly declared, or explicitly defaulted. Otherwise, a move constructor/move-assignment operator with a 'noexcept' exception specification should be provided. | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3279</a>   |
| GB 62           | 17-30                                       |  | Te                           | Issues with efficiency and unsatisfactory semantics mean many library functions document they do not throw exceptions with a <code>Throws: Nothing</code> clause, but do not advertise it with an exception specification. The semantic issues are largely resolved with the new 'noexcept' specifications, and the <code>noexcept</code> operator means we will want to detect these guarantees programatically in order to construct programs taking advantage of the guarantee.  | Add a 'noexcept' exception specification on each library API that offers an unconditional <code>Throws: Nothing</code> guarantee. Where the guarantee is conditional, add the appropriate <code>noexcept(constant-expression)</code> if an appropriate constant expression exists.  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3195</a>   |
| GB 63           | 17-30                                       |  | Ge                           | Since the newly introduced operator <code>noexcept</code> makes it easy (easier than previously) to detect whether or not a   | Review the whole library, and apply the <code>noexcept</code> specification where it is appropriate.  | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | function has been declared with the empty exception specification (including noexcept) library functions that cannot throw should be decorated with the empty exception specification. Failing to do so and leaving it as a matter of QoI would be detrimental to portability and efficiency.  |   | See paper <a href="#">N3155</a>                    |
| GB 64           | 17-30                                       |  | Ge                           | There are a number of unspecified types used throughout the library, such as the container iterators. Many of these unspecified types have restrictions or expectations on their behaviour in terms of exceptions. Are they permitted or required to use exception specifications, more specifically the new noexcept specification? For example, if <code>vector&lt;T&gt;::iterator</code> is implemented as a native pointer, all its operations will have an (effective) nothrow specification. If the implementation uses a class type to implement this iterator, is it permitted or required to support that same guarantee? | Clearly state the requirements for exception specifications on all unspecified library types. For example, all container iterator operations should be conditionally noexcept, with the condition matching the same operation applied to the allocator <code>pointer_type</code> , a certain subset of which are already required not to throw.                       | REJECTED<br>The Standard is correct as written.    |
| GB 65           | 17-30                                       |  | Te                           | Nothrowing swap operations are key to many C++ idioms, notably the common copy/swap idiom to provide the strong exception safety guarantee.  | Where possible, all library types should provide a swap operation with an exception specification guaranteeing no exception shall propagate. Where <code>noexcept(true)</code> cannot be guaranteed to not terminate the program, and the swap in questions is a template, an exception specification with the appropriate conditional expression could be specified. | ACCEPTED with MODIFICATIONS<br>See LWG 1349        |
| GB 66           | 17-30                                       |  | Ed                           | The syntax for attributes was altered specifically to make it legal do declare attributes on functions by making the attribute the first token in the declaration, rather than inserting it between the function name and the opening paren of the parameter list. This is much more readable, and should be the preferred style throughout the library. For example, prefer:<br>[[noreturn]] void exit(int status);<br>to<br>void exit [[noreturn]] (int status);   | Update all function declarations throughout the library to use the preferred placement of function-level attributes.  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
| CH 15           | Library clauses                             |  | te                           | Due to the new rules about implicit copy and move constructors some library facilities are now move-only.   | Make them copyable again.  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3266</a>   |
| CH 16           | Library clauses                             |  | te                           | Dynamic exception specifications are deprecated. Deprecated features shouldn't be used in the Standard.   | Replace dynamic exception specifications with noexcept.                                | ACCEPTED with MODIFICATIONS<br><br>See documents N3148, N3150, N3195, N3155, N3156, N3199, and N3180.<br><br>See paper <a href="#">N3148</a> |
| CH 17           | Library clauses                             |  | te                           | The introduction of noexcept makes "Throws: Nothing" clauses looking strange.   | Consider replacing "Throws: Nothing." clause by the respective noexcept specification. | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3195</a>   |
| CH 18           | 17  |  | te                           | The general approach on moving is that a library object after moving out is in a "valid but unspecified state". But this is stated at the single object specifications, which is error prone (especially if the move operations are implicit) and unnecessary duplication.  | Consider putting a general statement to the same effect into clause 17.                | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3264</a>   |
| JP 27           | 17.1  | 9  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(25)" to "(Clause 25)".  | ACCEPTED   |
| GB 52           | 17.3.7                                      |  | Te                           | The definition of deadlock in 17.3.7 excludes cases involving a single thread making it incorrect.  | The definition should be corrected.  | ACCEPTED with MODIFICATIONS<br>See LWG 1354  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
| GB 50           | 17.3.13                                     |  | Te                           | This definition of move-assignment operator redundant and confusing now that the term move-assignment operator is defined by the core language in subclause 12.8p21.  | Strike suclause 17.3.13 [defns.move.assign.op]. Add a cross-reference to (12.8) to 17.3.12.  | ACCEPTED with MODIFICATIONS<br><br>See paper N3142                 |
| GB 51           | 17.3.14                                     |  | Te                           | This definition of move-constructor redundant and confusing now that the term constructor is defined by the core language in subclause 12.8p3.  | Strike subclause 17.3.14, [defns.move.ctor]  | ACCEPTED with MODIFICATIONS<br><br>See paper N3142                 |
| JP 67           | 17.5.2.1.2                                  | 2  | E                            | In some code examples, ellipsis(...) is used in ill-formed. In these cases, "..." represents omission of some codes like this:<br>class A { /* ... */ };<br>But in some cases, it is used without commented-out as below:<br>class A { ... };<br>It is an inconsistent usage. They all should be enclosed in a comment.<br>In addition, the number of period should be 3 rather than 5. | Change to:<br>enum bitmask {<br>V0 = 1 << 0, V1 = 1 << 1, V2 = 1 << 2, V3 = 1 << 3, ...<br>};<br>static const bitmask C3 (V3);<br>/* ... */                            | REJECTED<br><br>This is pseudo-code; okay as written.              |
| GB 53           | 17.5.2.1.3                                  |  | Te                           | The bitmask types defined in 27.5.2 and 28.5 contradict the bitmask type requirements in 17.5.2.1.3, and have missing or incorrectly defined operators.   | See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3110</a> |
| GB 54           | 17.5.2.1.4.x                                |  | Ge                           | The defined terms NTC16S, NTC32S, NTWCS, char16-character sequence, null-terminated char16-character string, char32-character sequence, null-terminated char32-character string, wide-character sequence and null-terminated wide-character string do not occur at any point in the standard outside their definitional subclauses and associated footnotes.                            | Strike 17.5.2.1.4.3, 17.5.2.1.4.4, 17.5.2.1.4.5 and associated footnotes 170, 171, 172, 173, 174 and 175.  | ACCEPTED   |
| GB 55           | 17.6.1.3                                    | Table 15                                     | Te                           | The thread header uses duration types, found in the <chrono> header, and which rely on the ratio types declared in the <ratio> header.  | Add the <chrono> and <ratio> headers to the freestanding requirements.<br>It might be necessary to address scaled-down expectations of clock support in a freestanding | REJECTED<br><br>There was no consensus to make this change. We are |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted      |
|                 |   |  |                              |   | environment, much like <thread>.   | not adding new headers to freestanding at this point.   |
| GB 56           | 17.6.1.3                                    | Table 15                                     | Ge                           | The <utility> header provides support for several important C++ idioms with move, forward and swap. Likewise, declval will be frequently used like a type trait. In order to complete cycles introduced by std::pair, the <tuple> header should also be made available. This is a similarly primitive set of functionality, with no dependency of a hosted environment, but does go beyond the minimal set of functionality otherwise suggested by the freestanding libraries. Alternatively, split the move/forward/swap/declval functions out of <utility> and into a new primitive header, requiring only that of freestanding implementation. | Add <utility> and <tuple> to table 15, headers required for a free-standing implementation.  | REJECTED<br><br>No consensus for a change at this time. |
| GB 57           | 17.6.1.3                                    | Table 15                                     | Te                           | The atomic operations facility is closely tied to clause 1 and the memory model. It is not easily supplied as an after-market extension, and should be trivial to implement of a single-threaded serial machine. The consequence of not having this facility will be poor interoperability with future C++ libraries that memory model concerns seriously, and attempt to treat them in a portable way.   | Add <atomic> to table 15, headers required for a free-standing implementation.   | ACCEPTED  |
| GB 58           | 17.6.2                                      |  | Te                           | It is not clear whether a library header specified in terms of a typedef name makes that same typedef name available for use, or if it simply requires that the specified type is an alias of the same type, and so the typedef name cannot be used without including the specific header that defines it. For example, is the following code required to be accepted:<br>#include <vector><br>std::size_t x = 0;<br>Most often, this question concerns the typedefs defined in header <cstdlib>  | Add a paragraph under 17.6.2 clarifying whether or not headers specified in terms of std::size_t can be used to access the typedef size_t, or whether the header <cstdlib> must be included to reliably use this name. | REJECTED<br><br>The standard is correct as written.     |
| GB 59           | 17.6.3.6                                    | 2  | Ed                           | The replaceable functions in header <new> are all described in clause 18.6 [support.dynamic], where it can be seen that all the listed functions have an exception  | Narrow the reference to (Clause 18) to (Clause 18.6).<br>Add the missing exception specification on each   | ACCEPTED with MODIFICATIONS                             |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | specification which must be compatible with any replacement function.  | function signature:<br>void* operator new(std::size_t size)<br>throw(std::bad_alloc);<br>void* operator new(std::size_t size, const std::nothrow_t&) throw();<br>void operator delete(void* ptr) throw();<br>void operator delete(void* ptr, const std::nothrow_t&) throw();<br>void* operator new[](std::size_t size)<br>throw(std::bad_alloc);<br>void* operator new[](std::size_t size, const std::nothrow_t&) throw();<br>void operator delete[](void* ptr) throw();<br>void operator delete[](void* ptr, const std::nothrow_t&) throw();<br>(note that other comments might further want to replace 'throw()' with 'noexcept') | Cross-reference changed as requested. Additions to name mentions not done: these are not declarations. |
| US 81           | 17.6.3.9                                    | ¶1   | ed                           | International Standards do not make “statements”; they “specify” or “require” instead.   | s/statements/ specifications/   | ACCEPTED with MODIFICATIONS<br><br>Removed the offending word.   |
| US 82           | 17.6.3.9                                    | ¶1, bullet 3                                 | te                           | The second Note can benefit by adopting recent nomenclature.   | Rephrase the Note in terms of <i>xvalue</i> .   | ACCEPTED with MODIFICATIONS<br>See LWG 1362  |
| US 83           | 17.6.3.10                                   | ¶2, first sent.                              | ed                           | The word “constructor” is misspelled.  | s/constructor/constructor/  | ACCEPTED   |
| GB 68           | 18.2  | 4  | Te                           | There is no reason for the offsetof macro to invoke potentially throwing operations, so the result of noexcept(offsetof(type,member-designator)) should be true. | Add to the end of 18.2p4:<br>"No operation invoked by the offsetof macro shall throw an exception, and<br>noexcept(offsetof(type,member-designator)) shall be true."  | ACCEPTED   |
| JP 68           | 18.3.1.5                                    | 2  | E                            | In some code examples, ellipsis(...) is used in ill-formed. In these cases, "..." represents omission of some codes like this:                                   | Change to:<br>inline static constexpr float infinity() throw() {<br>return /* ... */; }   | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

|                    |                           |
|--------------------|---------------------------|
| Date: 7 April 2011 | Document: SC22 WG21 N3289 |
|--------------------|---------------------------|

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                                    |
|                 |   |  |                              | <p>class A { /* ... */ ;<br/>                     But in some cases, it is used without commented-out as below:<br/>                     class A { ... } ;<br/>                     It is an inconsistent usage. They all should be enclosed in a comment.<br/>                     More over, in this case, "implementation-defined" would be better than "...".</p>        | <pre>inline static constexpr float quiet_NaN() throw() { return /* ... */; } inline static constexpr float signaling_NaN() throw() { return /* ... */; }</pre>                        | Replaced "..." by "value".  |
| JP 84           | 18.5  | 5 note                                       | E                            | <p>Note in paragraph 5 says "the atexit() functions shall not introduce a data race" and Note in paragraph 10 says "the at_quick_exit() functions do not introduce ...". Such different expressions in similar functions are confusing.<br/>                     If these notes are written for unspecified behaviors just before the sentence, "do" would be preferred.</p> | Replace "shall" with "do".  | ACCEPTED  |
| GB 69           | 18.5  | 14   | Ed                           | <p>("The function quick_exit() never returns to its caller.") should be removed as redundant. The function is already attributed with [[noreturn]].</p>  | Remove paragraph 14   | ACCEPTED  |
| CH 19           | 18.8.5                                      |  | te                           | <p>It's not clear how exception_ptr is synchronized.</p>   | <p>Make clear that accessing in different threads multiple exception_ptr objects that all refer to the same exception introduce a race.</p>   | <p>ACCEPTED with MODIFICATIONS</p> <p>See paper <a href="#">N3278</a></p>             |
| GB 70           | 18.6  |  | Te                           | <p>std::nothrow_t is a literal type (being an empty POD) so the preferred form of declaration for std::nothrow is as a constexpr literal, rather than an extern symbol.</p>  | <p>Replace:<br/>                     extern const nothrow_t nothrow;<br/>                     with<br/>                     constexpr nothrow_t nothrow{};</p>                        | <p>REJECTED</p> <p>There was no consensus to adopt this change for this revision.</p> |
| GB 71           | 18.6.2.4 / 18.8.2.2 / 18.8.3.2              |  | Te                           | <p>The thread safety of std::set_new_handler(), std::set_unexpected(), std::set_terminate(), is unspecified making the the functions impossible to use in a thread safe manner.</p>  | <p>The thread safety guarantees for the functions must be specified and new interfaces should be provided to make it possible to query and install handlers in a thread safe way.</p> | <p>ACCEPTED with MODIFICATIONS</p> <p>See paper <a href="#">N3189</a></p>             |
| DE              | 18.6.1.4                                    |  | te                           | <p>It is unclear how a user replacement function can</p>   | <p>Offer a non-mutating interface to query the current</p>  | <p>ACCEPTED with MODIFICATIONS</p>  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
| 14              |   |  |                              | simultaneously satisfy the race-free conditions imposed in this clause and query the new-handler in case of a failed allocation with the only available, mutating interface <code>std::set_new_handler</code> .   | new-handler.   | See paper <a href="#">N3189</a>   |
| GB 72           | 18.8.2                                      |  | Ge                           | Dynamic exception specifications are deprecated, so clause 18.8.2 that describes library support for this facility should move to Annex D, with the exception of the <code>bad_exception</code> class which is retained to indicate other failures in the exception dispatch mechanism (e.g. calling <code>current_exception()</code> ).  | With the exception of 18.8.2.1 [ <code>bad.exception</code> ], move clause 18.8.2 directly to Annex D. [ <code>bad.exception</code> ] should simply become the new 18.8.2. | ACCEPTED  |
| GB 73           | 18.8.4                                      |  | Te                           | The thread safety <code>std::uncaught_exception()</code> and the result of the function when multiple threads throw exceptions at the same time are unspecified. To make the function safe to use in the presence of exceptions in multiple threads the specification needs to be updated.  | Update this clause to support safe calls from multiple threads without placing synchronization requirements on the user.   | ACCEPTED with MODIFICATIONS<br>Change 18.8.4 [ <code>uncaught</code> ] p. 1 as follows:<br><i>Returns</i> : true after <u>the current thread has initialized initializing an exception object (15.1) until a handler for the exception (including <code>unexpected()</code> or <code>terminate()</code>) is activated (15.3).</u> [ <i>Note</i> : This includes stack unwinding (15.2). — <i>end note</i> ] |
| GB 74           | 18.8.5                                      | 10   | Te                           | One idea for the <code>exception_ptr</code> type was that a reference-counted implementation could simply 'reactivate' the same exception object in the context of a call to <code>rethrow_exception</code> . Such an implementation would allow the same exception object to be active in multiple threads (such as when multiple threads join on a <code>shared_future</code> ) and introduce potential data races in any exception handler that catches exceptions by reference - notably existing library code written before this capability was | Throws: a copy of the exception object to which p refers.  | REJECTED<br>There was no consensus to adopt the proposed change.  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |
| US 84           | 18.8.6 [except.nested]                      | 6-7  | te                           | <p>added. 'rethrow_exception' should *always* make a copy of the target exception object.</p> <p>The throw_with_nested spec passes in its argument as T&amp;&amp; (perfect forwarding pattern), but then discusses requirements on T without taking into account that T may be an lvalue-reference type. It is also not clear in the spec that t is intended to be perfectly forwarded.</p>   | Patch 6-7 to match the intent with regards to requirements on T and the use of <code>std::forward&lt;T&gt;(t)</code> .  | <p>ACCEPTED with MODIFICATIONS</p> <p>See LWG 1370 <i>Change 18.8.7 nested_exception [except.nested] as indicated:</i><br/> [[noreturn]] template &lt;class T&gt; void<br/> throw_with_nested(T&amp;&amp; t);<br/> Let <u>U</u> be <u>remove_reference&lt;T&gt;::type</u><br/> 6 <i>Requires:</i> <math>\mp</math> <u>U</u> shall be CopyConstructible.<br/> 7 <i>Throws:</i> If <math>\mp</math> <u>U</u> is a non-union class type not derived from nested_exception, an exception of unspecified type that is publicly derived from both <math>\mp</math> <u>U</u> and nested_exception <u>and</u> constructed from <u>std::forward&lt;T&gt;(t)</u>, otherwise <u>throws std::forward&lt;T&gt;(t)</u>.</p> |
| GB 75           | 19  |  | Te                           | <p>None of the exception types defined in clause 19 are allowed to throw an exception on copy or move operations, but there is no clear specification that the operations have an exception specification to prove it. Note that the implicitly declared constructors, taking the exception specification from their base class (ultimately <code>std::exception</code>) will implicitly generate a <code>noexcept</code> exception specification if all of their data members similarly declare <code>noexcept</code> operations. As the</p> | Add a global guarantee that all exception types defined in clause 19 that rely on implicitly declared operations have a non-throwing exception specification on those operations. | <p>REJECTED</p> <p>The standard is correct as written. See 18.8.1 [exception], p2</p>   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                                  |
|                 |   |  |                              | representation is unspecified, we cannot assume non-throwing operations unless we explicitly state this as a constraint on the implementation.   |  |   |
| GB 76           | 19.5.1.5                                    |  | Te                           | The C++0x FCD recommends, in a note (see 19.5.1.1/1), that users create a single error category object for each user defined error category and specifies error_category equality comparisons based on equality of addresses (19.5.1.3). The Draft apparently ignores this when specifying standard error category objects in section 19.5.1.5, by allowing the generic_category() and system_category() functions to return distinct objects for each invocation. | Append a new sentence to 19.5.1.5 [syserr.errcat.objects]/1, which reads "All calls of this function return references to the same object.", and append the same sentence to 19.5.1.5/3.   | ACCEPTED  |
| GB 77           | 19.5.6.2                                    | 14   | Ed                           | The description for system_error::what (19.5.6.2/14) changed between the C++ Working Papers N2914 and N2960. The latter document indicates that the Returns clause shall read "Returns: An NTBS incorporating the arguments supplied in the constructor.". Instead, in the FCD it now reads "Returns: An NTBS incorporating and code().message() the arguments supplied in the constructor.".  | Remove the extra words from 19.5.6.2 [syserr.syserr.members]/14: "Returns: A NTBS incorporating the arguments supplied in the constructor."  | ACCEPTED  |
| GB 78           | 19.5.6.2                                    |  | Te                           | The FCD contains a requirement that all standard classes which are derived from std::exception have a copy constructor and copy assignment operator which essentially copy the stored what() message (See 18.8.1/2). In contrast, it is unspecified whether copies of system_error return the same error_code on calls to system_error::code().  | Insert a new paragraph after 19.5.6.1 [syserr.syserr.overview]/1 which reads "The copy constructor and copy assignment operator of class system_error both have a strengthened postcondition which supplements the basic postcondition for standard library exception classes copy construction and copy assignment (18.8.1): If two objects lhs and rhs | REJECTED<br><br>This is already addressed by 17.5.2.2 [functions.within.classes]/1. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              |   | both have type system_error and lhs is a copy of rhs, then lhs.code() == rhs.code() shall hold."  |  |
| GB 79           | 20, 22, 24, 28                              |  | Te                           | The library provides several traits mechanisms intended a customization points for users. Typically, they are declared in headers that are growing quite large. This is not a problem for standard library vendors, who can manage their internal file structure to avoid large dependencies, but can be a problem for end users who have no option but to include these large headers.   | Move the following traits classes into their own headers, and require the existing header to #include the traits header to support backwards compatibility:<br>iterator_traits (plus the iterator tag-types)<br>allocator_traits<br>pointer_traits<br>char_traits<br>regex_traits | REJECTED<br><br>This suggestion is not a defect, as the likely benefit is small, if any, compared to the cost of not just implementing the feature, but also explaining/teaching it. There was no consensus to make this change. |
| US 85           | 20.2.1                                      | Table 34                                     |                              | 20.2.1 Table 34 "MoveConstructible requirements" says "Note: rv remains a valid object. Its state is unspecified".<br><br>Some components give stronger guarantees. For example, moved-from shared_ptrs are guaranteed empty (20.9.11.2.1/25).<br><br>In general, what the standard really should say (preferably as a global blanket statement) is that moved-from objects can be destroyed and can be the destination of an assignment. Anything else is radioactive. For example, containers can be "emptier than empty". This needs to be explicit and required generally.<br><br>Note: The last time that one of us mentioned "emptier than empty" (i.e. containers missing sentinel nodes, etc.) the objection was that containers can store sentinel nodes inside themselves in order to avoid dynamically allocating them. This is unacceptable because (a) it forces existing implementations (i.e. Dinkumware's, Microsoft's, IBM's, etc.) to change for no good reason (i.e. permitting more operations on moved-from objects), and (b) it invalidates end iterators when swapping containers. (The Working Paper currently permits end iterator invalidation, which we consider to be wrong, but that's a separate argument. In | State as a general requirement that moved-from objects can be destroyed and can be the destination of an assignment. Any other use is undefined behavior.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3264</a>   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              | any event, *mandating* end iterator invalidation is very different from permitting it.)  |  |  |
| GB 80           | 20.2.3                                      | 2  | Te                           | See (A) in attachment std_issues.txt   | as stated in the attached paper  | REJECTED<br>There was no consensus to make this change.  |
| CA 10           | 20.2.3p2                                    | 20.2.3p2                                     | te                           | <p>Reads of indeterminate value result in undefined behaviour</p> <p>In 20.2.3p2, NullablePointer requirements [nullablepointer.requirements], the standard specifies the behaviour of programs that read indeterminate values:</p> <p>... A default-initialized object of type P may have an indeterminate value. [ Note: Operations involving indeterminate values may cause undefined behaviour. end note ]</p> <p>We suggest changing the note to:</p> <p>[ Note: Operations involving indeterminate values cause undefined behaviour. end note ]</p> <p>Rationale: The note uses the word "may", but we believe the intention is that such reads will cause undefined behaviour, but implementations are not required to produce an error.</p> <p>Clark adds:</p> <p>&gt; Unfortunately, this issue goes deeper than can be</p> | <p>In 20.2.3p2, the standard specifies the behaviour of programs that read indeterminate values:</p> <p>... A default-initialized object of type P may have an indeterminate value. [ Note: Operations involving indeterminate values cause undefined behaviour. end note ]</p> <p>We suggest changing the note to:</p> <p>[ Note: Operations involving indeterminate values cause undefined behaviour. end note ]</p> | <p>REJECTED</p> <p>It is an accurate statement because there is at least one operation, assignment with the indeterminate value on the left hand side, that does not involve undefined behavior.</p> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>&gt; addressed by deleting the word "may" from a note in</p> <p>&gt; clause 20. The term "indeterminate value" and its</p> <p>&gt; meaning were introduced in C99. While the term is</p> <p>&gt; generally understood to be applicable to C++ (and</p> <p>&gt; section 20.2.3 reflects that), the term isn't actually</p> <p>&gt; defined in the C++ WD, and worse yet, there's no</p> <p>&gt; statement that use of an indeterminate value can result</p> <p>&gt; in undefined behavior (so the existing note can't be</p> <p>&gt; deduced from the normative statements of the</p> <p>&gt; standard).</p> <p>&gt; This is tracked by core issue 616. The wording in</p> <p>&gt; 20.2.3 should be noted as relating to that issue, and</p> <p>&gt; should be handled as part thereof.</p> <p>Further on this, in the current draft standard, we can construct executions in which an atomic read happens before the initialization of the atomic object, so there is no place to take a read value from. We imagine that such reads should also be of indeterminate values and result in undefined behaviour?</p> |  |  |
| US 86           | 20.2.5                                      | Table 42                                     | ed                           | In the row for X::propagate_on_container_move_assignment, the note says "copied" when it should say "moved"   | Change the note as follows: true_type only if an allocator of type X should be <del>copied</del> moved | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted       |
| US 87           | 20.2.5                                      | Table 42                                     | te                           | reference_type should not have been removed from the allocator requirements. Even if it is always the same as value_type&, it is an important customization point for extensions and future features.  | Add a row (after value_type) with columns:<br>Expression: X::reference_type<br>Return type: T&<br>Assertion/note...: (empty)<br>Default: T&<br><br>[allocator.traits]: Add reference_type to allocator_traits template, defaulted to value_type&.   | REJECTED<br>There was no consensus to adopt this change. |
| US 88           | 20.2.5                                      |  | Te                           | Allocator interface is not backward compatible.  | See Appendix 1 - Additional Details   | REJECTED<br><br>Withdrawn by the submitter.              |
| US 89           | 20.3 [utility]                              |  | ed                           | make_pair is missing from the <utility> synopsis.  | Add template <class T1, class T2> pair<V1, V2> make_pair(T1&&, T2&&); to the synopsis in [utility].   | ACCEPTED   |
| GB 81           | 20.3  | 1  | Ed                           | make_pair should be declared in the synopsis of <utility>  | add to 20.3 [utility] paragraph 1:<br>template<br>see below make_pair(T1&&, T2&&);  | ACCEPTED   |
| GB 82           | 20.3  | 1  | Ed                           | The <utility> synopsis precedes the tuple_size and tuple_element declarations with a comment saying "// 20.3.5, tuple-like access to pair:" but the sub-clause should be 20.3.5.3 and the comment should probably be below those declarations (since they, like the tuple declaration above them, are defined in <tuple> and are not related to pair.)<br><br>Also, there should be a comment above piecewise_construct_t giving the sub-clause 20.3.5.5, and it should be at the end of the synopsis, | correct the sub-clause reference for tuple-like access to 20.3.5.3, move the comment after the tuple_element declaration, and add "// defined in <tuple>" to the tuple_size and tuple_element declarations. Move the piecewise_construct_t and piecewise_construct declarations to the end of the synopsis and precede them with "// 20.3.5.5 Piecewise construction" | ACCEPTED   |
| US 90           | 20.3  | 3  | te                           | In n3090, at variance with previous iterations of the idea discussed in papers and incorporated in WDs, std::forward is constrained via std::is_convertible, thus is   |   | ACCEPTED with MODIFICATIONS                              |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)                             | (7)   |
|-----------------|---|--|------------------------------|--|---------------------------------|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB       | Secretariat observations on each comment submitted          |
|                 |   |  |                              | <p>not robust wrt access control. This causes problems in normal uses as implementation detail of member functions. For example, the following snippet leads to a compile time failure, whereas that was not the case for an implementation along the lines of n2835 (using enable_ifs instead of concepts for the constraining, of course)</p> <pre>#include &lt;utility&gt; struct Base { Base(Base&amp;&amp;); }; struct Derived : private Base {     Derived(Derived&amp;&amp; d)     : Base(std::forward&lt;Base&gt;(d)) {} };</pre> <p>In other terms, LWG 1054 can be resolved in a better way, the present status is not acceptable.</p> |                                 | See paper <a href="#">N3143</a>                             |
| JP 28           | 20.3.1                                      | 2  | E                            | <p>Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.</p>   | Change "(31)" to "(Clause 31)". | ACCEPTED with MODIFICATIONS<br><br>Changed to "(Table 31)". |
| JP 29           | 20.3.1                                      | 4, 6, 8                                      | E                            | <p>Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and</p>  | Change "(32)" to "(Clause 32)". | ACCEPTED with MODIFICATIONS                                 |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |
|                 |   |  |                              | subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  |   | Changed to "(Table 32)".  |
| JP 30           | 20.3.2                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(34)" to "(Table 34)".<br>Change "(36)" to "(Table 36)".                | ACCEPTED  |
| US 91           | Merged with US 90                           |  |                              |   |   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3143</a>                                    |
| US 92           | 20.3.3                                      |  | te                           | std::identity was removed from 20.3.3 [forward], apparently because std::forward() no longer needs it. However, std::identity was useful by itself (to disable template argument deduction, and to provide no transformation when one is requested).  | Restore std::identity, possibly in a different section.                         | REJECTED<br><br>No consensus to restore at this time. common_type can be used for the use case given. |
| US 93           | 20.3.3                                      | 6  | ge                           | When teaching C++0x, students have consistently found the name std::move confusing because it doesn't actually move the object (it just makes it possible to move). It was also confusing for me.   | Choose a name that expresses the semantics more clearly. Suggestion: std::unpin | REJECTED<br><br>There was no consensus to adopt this change.  |
| US 94           | 20.3.3                                      | 9  | ed                           | Returns clause for move_if_noexcept d refers to a non-existent symbol, t, which should be x.  | Modify the Returns clause:<br><i>Returns:</i> std::move( <del>t</del> x).       | ACCEPTED  |
| DE              | 20.3.5.2,                                   |  | te                           | Several function templates of pair and tuple allow for too  | Consider to add wording to constrain these                                      | ACCEPTED with MODIFICATIONS   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted             |
| 15              | 20.4.2.1                                    |  |                              | many implicit conversions, for example:<br><pre>#include &lt;tuple&gt; std::tuple&lt;char*&gt; p(0); // Error? struct A { explicit A(int){} }; A a = 1; // Error std::tuple&lt;A&gt; ta = std::make_tuple(1); // OK?</pre>  | function templates.   | See paper <a href="#">N3140</a>                                |
| JP 31           | 20.5.1                                      | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(section 3.9)" to "(3.9)".  | ACCEPTED   |
| GB 84           | 20.3.5.2                                    |  | Ed                           | [pairs.pair] defines the class template pair as well as related non-member functions such as comparison operators and make_pair. The related non-member functions should be in a separate sub-clause, otherwise it's not clear that paragraphs below 17 do not refer to members of pair.  | Create a new "Pair specialized algorithms" section containing everything below paragraph 17 in [pairs.pair]   | ACCEPTED   |
| US 95           | 20.3.5.2                                    | 9  | te                           | Copy-assignment for pair is defaulted and does not work for pairs with reference members. This is inconsistent with conversion-assignment, which deliberately succeeds even if one or both elements are reference types, just as for tuple. The copy-assignment operator should be consistent with the conversion-assignment operator and with tuple's assignment operators.  | Add to pair synopsis:<br><u><a href="#">pair&amp; operator=(const pair&amp; p);</a></u><br>Add before paragraph 9:<br><u><a href="#">pair&amp; operator=(const pair&amp; p);</a></u><br><u><a href="#">_Requires: T1 and T2 shall satisfy the requirements of CopyAssignable.</a></u><br><u><a href="#">_Effects: Assigns p.first to first and p.second to second.</a></u><br><u><a href="#">_Returns: *this.</a></u> | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3140</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted  |
| DE 16           | 20.3.5.2, 20.4.2.1                          |  | te                           | Several pair and tuple functions in regard to move operations are incorrectly specified if the member types are references, because the result of a std::move cannot be assigned to lvalue-references. In this context the usage of the requirement sets MoveConstructible and CopyConstructible also doesn't make sense, because non-const lvalue-references cannot satisfy these requirements.   | Replace the usage of std::move by that of std::forward and replace MoveConstructible and CopyConstructible requirements by other requirements.   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3140</a>  |
| GB 85           | 20.3.5.4                                    |  | Te                           | While std::pair may happen to hold a pair of iterators forming a valid range, this is more likely a coincidence than a feature guaranteed by the semantics of the pair template. A distinct range-type should be supplied to enable the new for-loop syntax rather than overloading an existing type with a different semantic.  | Strike 20.3.5.4 and the matching declarations in 20.3 header synopsis.<br>If a replacement facility is required for C++0x, consider n2995.   | ACCEPTED with MODIFICATIONS<br>See LWG 1381<br><br>Strike 20.3.5.4 and the matching declarations in 20.3 header synopsis. |
| ES 1            | 20.3.5.4 [pair.range]                       |  | Te                           | Using pair to represent a range of iterators is too general and does not provide additional useful restrictions (see N2995 and preceding papers).  | Provide a separate template range<Iterator>.   | REJECTED<br><br>No consensus to make this change.   |
| US 96           | 20.3.5.2<br>20.4.2.1<br>20.4.2.2            | ¶ 6-14<br>¶ 6-20<br>¶ 6-18                   | te                           | pair and tuple constructors and assignment operators use std::move when they should use std::forward. This causes lvalue references to be erroneously converted to rvalue references. Related requirements clauses are also wrong.   | See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3140</a>  |
| US 97           | 20.3.5.2 and 20.4.2                         |  | te                           | pair's class definition in N3092 20.3.5.2 [pairs.pair] contains "pair(const pair&) = default;" and "pair& operator=(pair&& p);". The latter is described by 20.3.5.2/12-13.<br><br>"pair(const pair&) = default;" is a user-declared explicitly-defaulted copy constructor. According to 12.8 [class.copy]/10, this inhibits the implicitly-declared move constructor. pair should be move constructible. (12.8/7 explains that "pair(pair<U, V>&& p)" will never be | Either remove "pair(const pair&) = default;" and "pair& operator=(pair&& p);" from pair's class definition in 20.3.5.2 and from 20.3.5.2/12-13, or give pair explicitly-defaulted copy/move constructors and copy/move assignment operators.<br><br>Change tuple to match. | ACCEPTED with MODIFICATIONS<br><br>See paper N3140  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)                       | (7)  |
|-----------------|---|--|------------------------------|---|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>instantiated to move pair&lt;T1, T2&gt; to pair&lt;T1, T2&gt;.)</p> <p>"pair&amp; operator=(pair&amp;&amp; p);" is a user-provided move assignment operator (according to 8.4.2 [dcl.fct.def.default]/4: "A special member function is user-provided if it is user-declared and not explicitly defaulted on its first declaration."). According to 12.8/20, this inhibits the implicitly-declared copy assignment operator. pair should be copy assignable, and was in C++98/03. (Again, 12.8/7 explains that "operator=(const pair&lt;U, V&gt;&amp; p)" will never be instantiated to copy pair&lt;T1, T2&gt; to pair&lt;T1, T2&gt;.)</p> <p>Additionally, "pair&amp; operator=(pair&amp;&amp; p);" is unconditionally defined, whereas according to 12.8/25, defaulted copy/move assignment operators are defined as deleted in several situations, such as when non-static data members of reference type are present.</p> <p>If "pair(const pair&amp;) = default;" and "pair&amp; operator=(pair&amp;&amp; p);" were removed from pair's class definition in 20.3.5.2 and from 20.3.5.2/12-13, pair would receive implicitly-declared copy/move constructors and copy/move assignment operators, and 12.8/25 would apply. The implicitly-declared copy/move constructors would be trivial when T1 and T2 have trivial copy/move constructors, according to 12.8/13, and similarly for the assignment operators, according to 12.8/27. Notes could be added as a reminder that these functions would be implicitly-declared, but such notes would not be necessary (the Standard Library specification already assumes a high level of familiarity with the Core Language, and casual readers will simply assume that pair is copyable and movable).</p> <p>Alternatively, pair could be given explicitly-defaulted copy/move constructors and copy/move assignment operators. This is a matter of style.</p> |                           |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted      |
|                 |   |  |                              | <p>tuple is also affected. tuple's class definition in 20.4.2 [tuple(tuple)] contains:</p> <pre>tuple(const tuple&amp;) = default; tuple(tuple&amp;&amp;); tuple&amp; operator=(const tuple&amp;); tuple&amp; operator=(tuple&amp;&amp;);</pre> <p>They should all be removed or all be explicitly-defaulted, to be consistent with pair. Additionally, 20.4.2.1 [tuple.cnstr]/8-9 specifies the behavior of an explicitly-defaulted function, which is currently inconsistent with pair.</p> |   |   |
| GB 86           | 20.4  |  | Te                           | tuple and pair are essentially two implementations of the same type, with the same public interface, and their specification in becoming increasingly intertwined. The tuple library should be merged into the <utility> header to reduce library dependencies and simplify user expectations. The <tuple> header could optionally be retained as a deprecated alias to the <utility> header.   | Merge everything declared in 20.4 into the <utility> header. Either remove the <tuple> header entirely, or move it to Annex D as a deprecated alias of the <utility> header.  | REJECTED<br>There was no consensus to make this change. |
| US 98           | 20.4.2.4                                    | Paragraph 4                                  | te/ed                        | pack_arguments is poorly named. It does not reflect the fact that it is a tuple creation function and that it forwards arguments.   | Rename pack_arguments to forward_as_tuple throughout the standard.  | ACCEPTED  |
| GB 88           | 20.4.2.4                                    |  | Te                           | The tuple_cat template consists of four overloads and that can concatenate only two tuples. A single variadic signature that can concatenate an arbitrary number of tuples would be preferred.  | Adopt a simplified form of the proposal in n2795, restricted to tuples and neither requiring nor outlawing support for other tuple-like types.  | ACCEPTED with MODIFICATIONS<br>See LWG 1385             |
| US 99           | 20.4.2.4                                    | 4 - 6  | te                           | pack_arguments is overly complex.   | <p>This issue resulted from a lack of understanding of how references are forwarded. The definition of pack_arguments should be simply:</p> <pre>template &lt;class... Types&gt; tuple&lt;ATypes&amp;&amp;&gt; pack_arguments(Types&amp;&amp;...t);</pre> | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                 |
|                 |   |  |                              |  | <p><del>Types: Let Ti be each type in Types....</del></p> <p>Effects: ...</p> <p>Returns:<br/>tuple&lt;A Types&amp;&amp;...&gt;(std::forward&lt;Types&gt;(t...))</p> <p>The synopsis should also change to reflect this simpler signature.</p>  |  |
| GB 87           | 20.4.2.10                                   |  | Te                           | There is no compelling reason to assume a heterogeneous tuple of two elements holds a pair of iterators forming a valid range. Unlike std::pair, there are no functions in the standard library using this as a return type with a valid range, so there is even less reason to try to adapt this type for the new for-loop syntax.  | Strike 20.4.2.10 and the matching declarations in the header synopsis in 20.4.  | ACCEPTED   |
| US 100          | 20.6.1 [ratio.ratio]                        |  | te                           | LWG 1281 was discussed in Pittsburgh, and the decision there was to accept the typedef as proposed and move to Review. Unfortunately the issue was accidentally applied to the FCD, and incorrectly. The FCD version of the typedef refers to ratio<N, D>, but the typedef is intended to refer to ratio<num, den> which in general is not the same type.  | Accept the current proposed wording of LWG 1281 which adds:<br><br>typedef ratio<num, den> type;  | ACCEPTED   |
| GB 89           | 20.6.2                                      |  | Te                           | The alias representations of the ratio arithmetic templates do not allow implementations to avoid overflow, since they explicitly specify the form of the aliased template instantiation. For example ratio_multiply, ratio<2, LLONG_MAX>> is *required* to alias ratio<2*LLONG_MAX, LLONG_MAX*2>, which overflows, so is ill-formed. However, this is trivially equal to ratio<1, 1>. It also contradicts the opening statement of 20.6.2p1 "implementations may use other algorithms to compute these values". | Change the wording in 20.6.2p2-5 as follows:<br><br>template <class R1, class R2> using ratio_add = see below;<br><br>The type ratio_add<R1, R2> shall be a synonym for<br><del>ratio<T1, T2></del> <ins>ratio<U, V> such that ratio<U, V>::num and ratio<U, V>::den are the same as the corresponding members of ratio<T1, T2> would be in the absence of arithmetic overflow</ins> where T1 has the value R1::num * R2::den + | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3210</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | <p>R2::num * R1::den and T2 has the value R1::den * R2::den. &lt;ins&gt;If the required values of ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den cannot be represented in intmax_t then the program is illformed.&lt;/ins&gt;</p> <p>template &lt;class R1, class R2&gt; using ratio_subtract = see below;</p> <p>The type ratio_subtract&lt;R1, R2&gt; shall be a synonym for &lt;del&gt;ratio&lt;T1,T2&gt;&lt;/del&gt; &lt;ins&gt;ratio&lt;U, V&gt; such that ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den are the same as the corresponding members of ratio&lt;T1,T2&gt; would be in the absence of arithmetic overflow&lt;/ins&gt; where T1 has the value R1::num * R2::den - R2::num * R1::den and T2 has the value R1::den * R2::den. &lt;ins&gt;If the required values of ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den cannot be represented in intmax_t then the program is illformed.&lt;/ins&gt;</p> <p>template &lt;class R1, class R2&gt; using ratio_multiply = see below;</p> <p>The type ratio_multiply&lt;R1, R2&gt; shall be a synonym for &lt;del&gt;ratio&lt;T1,T2&gt;&lt;/del&gt; &lt;ins&gt;ratio&lt;U, V&gt; such that ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den are the same as the corresponding members of ratio&lt;T1,T2&gt; would be in the absence of arithmetic overflow&lt;/ins&gt;</p> |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  | <p>where T1 has the value R1::num * R2::num and T2 has the value R1::den * R2::den. &lt;ins&gt;If the required values of ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den cannot be represented in intmax_t then the program is illformed.&lt;/ins&gt;</p> <p>template &lt;class R1, class R2&gt; using ratio_divide = see below;</p> <p>The type ratio_divide&lt;R1, R2&gt; shall be a synonym for &lt;del&gt;ratio&lt;T1,T2&gt;&lt;/del&gt; &lt;ins&gt;ratio&lt;U, V&gt; such that ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den are the same as the corresponding members of ratio&lt;T1,T2&gt; would be in the absence of arithmetic overflow&lt;/ins&gt; where T1 has the value R1::num * R2::den and T2 has the value R1::den * R2::num. &lt;ins&gt;If the required values of ratio&lt;U,V&gt;::num and ratio&lt;U,V&gt;::den cannot be represented in intmax_t then the program is illformed.&lt;/ins&gt;</p> |  |
| US 101          | 20.7  |  | te                           | Paper n2965 was largely rejected after the last CD on the grounds there was no associated national body comment, so I am submitting a national body comment this time.   | Consider n2965 in the context of a national body comment.   | REJECTED<br><br>There is no consensus to adopt this proposal at this time.   |
| GB 90           | 20.7  |  | Ed                           | type_traits is a core support facility offered by the compiler, and exposed with a library interface that is required in a free-standing implementation. It has far more in common with numeric_limits than the utility components in clause 20, and should move to clause 18. | Move clause 20.7 into clause 18.  | REJECTED<br><br>Type traits support queries about all types, not just built-in types. They do not belong in Clause 18. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| DE 17           | 20.7  |  | te                           | Speculative compilation for <code>std::is_constructible</code> and <code>std::is_convertible</code> should be limited, similar to the core language (see 14.8.2 paragraph 8).   |   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3142</a>   |
| DE 18           | 20.7  |  | te                           | Several type traits require compiler support, e.g. <code>std::is_constructible</code> or <code>std::is_convertible</code> . Their current specification seems to imply, that the corresponding test expressions should be well-formed, even in absense of access:<br><br><code>class X { X(int){} };<br/>constexpr bool test = std::is_constructible&lt;X, int&gt;::value;</code><br><br>The specification does not clarify the context of this test and because it already goes beyond normal language rules, it's hard to argue by means of normal language rules what the context and outcome of the test should be. | Specify that <code>std::is_constructible</code> and <code>std::is_convertible</code> will return true only for public constructors/conversion functions.  | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3142</a>   |
| US 102          | 20.7.4                                      |  | te                           | Despite Library Issue 520's ("Result_of and pointers to data members") resolution of CD1, the FCD's <code>result_of</code> supports neither pointers to member functions nor pointers to data members. It should.   | Ensure <code>result_of</code> supports pointers to member functions and pointers to data members.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3123</a>   |
| GB 91           | 20.7.4.3                                    | Table 45                                     | Ed                           | It is mildly distasteful to dereference a null pointer as part of our specification, as we are playing on the edges of undefined behaviour. With the addition of the <code>declval</code> function template, already used in these same expressions, this is no longer necessary.   | Replace the sub-expression <code>*(U*)0'</code> with the sub-expression <code>'declval&lt;U&gt;()'</code> in the specification for <code>has_nothrow_copy_assign</code> and <code>has_nothrow_move_assign</code> type traits. | ACCEPTED with MODIFICATIONS<br><br>The wording in question was removed by document N3142.<br><br>See paper <a href="#">N3142</a> |
| GB 92           | 20.7.4.3                                    | Table 45                                     | Te                           | Trivial functions implicitly declare a <code>noexcept</code> exception specification, so the references to <code>has_trivial_*</code> traits in the <code>has_nothrow_*</code> traits are redundant, and should be struck for clarity.  | For each of the <code>has_nothrow_something</code> traits, remove all references to the matching <code>has_trivial_something</code> traits.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3142</a>   |
| FI 18           | 20.7.4.3, Table 45                          |  | te                           | Related to the change proposed in FI 17, there should be a trait for checking whether a destructor throws.  | Add the following type predicate:   | REJECTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB         | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 | (Type property predicates)                  |  |                              |  | <p>Template</p> <pre>template &lt;typename T&gt; struct has_nothrow_destructor;</pre> <p>Condition</p> <pre>has_trivial_destructor&lt;T&gt;::value is true or the expression noexcept((*U*)0).~U()) is well-formed and true, where U is remove_all_extents&lt;T&gt;::type.</pre> <p>Precondition</p> <p>T shall be a complete type, (possibly cv-qualified) void, or an array of unknown bound.</p> <p>Reasoning:</p> <p>With this metafunction the destructor of a class template can adjust its noexcept specification depending on whether destructors of its unbound members (or unbound base classes) might throw:</p> <pre>template &lt;typename T&gt; struct C {     T t;     ~C() noexcept(has_nothrow_destructor&lt;T&gt;::value) {} };</pre> | <p>No consensus to make this change at this time.</p> <p>See paper <a href="#">N3142</a></p> |
| DE              | 20.7.4.3                                    |  | te                           | The fundamental trait is_constructible reports false | Remove all false positives from the domain of  | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted             |
| 19              |   |  |                              | positives, e.g.<br>is_constructible<char*, void*>::value<br>evaluates to true, even though a corresponding variable initialization would be ill-formed.   | is_constructible.   | See paper <a href="#">N3047</a>                                |
| JP 32           | 20.7.5                                      | 2, table 47                                  |                              | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(10)" to "(Clause 10)".   | ACCEPTED   |
| GB 93           | 20.7.6.2                                    | Table 49                                     | Te                           | The comments for add_rvalue_reference say "this rule reflects the semantics of reference collapsing", but reference collapsing is not defined anywhere.   | Add a new sentence at the end of 8.3.2p6 "This is called reference collapsing".<br>Add a reference to 8.3.2 to the use of "reference collapsing" in 20.7.6.2/table 49                     | REJECTED<br>There is no consensus to make this change.         |
| US 103          | 20.7.6.6<br>[meta.trans.other]              |  | te                           | The current definition of result_of works for function pointers but the condition statement outlaws them. There is even an example in the WP that shows result_of working for function pointers.  | Add "pointer to function" to the list of things that Fn shall be.   | REJECTED<br>Standard is correct as written.                    |
| US 104          | 28.8  |  | te                           | std::basic_regex should have an allocator for all the reasons that a std::string does. For example, I can use boost::interprocess to put a string or vector in shared memory, but not a regex.  | Add allocators to regexes   | REJECTED<br>No consensus for a change at this time             |
| GB 94           | 20.8  |  | Ed                           | This subclause has grown large, with many components, and so should stand alone.  | Promote 20.8 and all its contents up to a new numbered clause.  | REJECTED<br>Clause 20 has been extensively reorganized.        |
| GB 95           | 20.8  |  | Ge                           | The adaptable function protocol supported by unary_function/binary_function has been superceded by lambda expressions and std::bind. Despite the name, the protocol is not very adaptable as it requires intrusive  | Move clauses 20.8.3, 20.8.9, 20.8.11 and 20.8.12 to Annex D. Remove the requirements to conditionally derive from unary/binary_function from function, reference_wrapper, and the results | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3198</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                        |
|                 |   |  |                              | support in the adaptable types, rather than offering an external traits-like adaption mechanism. This protocol and related support functions should be deprecated, and we should not make onerous requirements for the specification to support this protocol for callable types introduced in this standard revision, including those adopted from TR1. It is expected that high-quality implementations will provide such support, but we should not have to write robust standard specifications mixing this restricted support with more general components such as function, bind and reference_wrapper.  | of calling mem_fn and bind.  |   |
| GB 96           | 20.8  |  | Te                           | The function templates 'hash', 'less' and 'equal_to' are important customization points for user-defined types to be supported by several standard containers. These are accessed through the <functional> header which has grown significantly larger in C++0x, exposing many more facilities than a user is likely to need through there own header, simply to declare the necessary specialization. There should be a smaller header available for users to make the necessary customization.   | Provide a tiny forwarding header for important functor types in the <functional> header that a user may want to specialize. This should contain the template declaration for 'equal_to', 'hash' and 'less'.    | REJECTED<br>No consensus to make a change.                                |
| GB 97           | 20.8.10                                     |  | Te                           | The bind template is intended as a single, simple to use replacement for the '98 adaptable function APIs and machinery. It works well in almost all respects, but lacks the ability to easily negate a predicate, or equivalently, act as a replacement for not1 and not2. Two easy ways to solve this omission would be to add a 'bind_not' function that produces a binder that negates its result. However, preference is given to requiring the unspecified bind result type to overload operator! to produce the same effect. This is preferred due to (i) its simpler usage, being the naively expected syntax, but more importantly (ii) some (limited) field experience. | Require the unspecified result of a bind expression to support unary operator! to yield another bind result that, when evaluated, yields 'res', where 'res' is the result of evaluating the original function. | REJECTED<br>There is no consensus to adopt this change for this revision. |
| JP 3            | 20.8.14.2                                   |  | TL                           | explicit default constructor is defined in std::function. Although it is allowed according to 12.3.1, it seems unnecessary to qualify the constructor as explicit.   | Remove explicit.<br><br>namespace std {<br>template<class> class function;   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | If it is explicit, there will be a limitation in initializer_list.  | <pre>// undefined  template&lt;class R, class... ArgTypes&gt; class function&lt;R(ArgTypes...)&gt; : public unary_function&lt;T1, R&gt; // iff sizeof...(ArgTypes) == 1 and ArgTypes contains T1  : public binary_function&lt;T1, T2, R&gt; // iff sizeof...(ArgTypes) == 2 and ArgTypes contains T1 and T2 { public:typedef R result_type;  // 20.8.14.2.1, construct/copy/destroy:function();</pre> |  |
| JP 33           | 20.8.15                                     | 1  | E                            | <p>Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.</p> <p>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.</p> | Change "(37)" to "(Table 37)".  | ACCEPTED   |
| JP 34           | 20.8.15                                     | 1  | E                            | <p>Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.</p> <p>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.</p> | Change "(33)" to "(Table 33)".  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted          |
| JP 4            | 20.8.14.2.1                                 | 1  | TL                           | Really does the function require that default constructor is explicit?   | Remove explicit.<br>function();<br>template <class A><br>function(allocator_arg_t, const A& a);                                | ACCEPTED  |
| US 105          | 20.8.15.2                                   |  | te                           | unique_ptr and shared_ptr are inconsistent in their handling of arrays. We can write:<br><br>unique_ptr<int[]> p( new int[10] );<br><br>// handles deletion correctly<br><br>But we cannot write:<br><br>shared_ptr<int[]> p( new int[10] );<br><br>// incorrect<br><br>This is an inconsistency. It is true that we have the following workaround:<br><br>std::shared_ptr<int> s(new int[5],<br>std::default_delete<int[]>());<br><br>But this is still inconsistent, not to mention awkward and error-prone because the programmer will occasionally forget the deleter and the code will silently compile and may appear to work on some platforms. | Support:<br>shared_ptr<int[]> p( new int[10] );<br><br>to handle deletion correctly by calling delete[] on the stored pointer. | REJECTED<br><br>There is no consensus to adopt this change. |
| GB 98           | 20.9  |  | Ed                           | This subclause has grown large, with many components, and so should stand alone.   | Promote 20.9 and all of its contents to a new, top-level, numbered clause.   | REJECTED<br><br>Clause 20 has been extensively reorganized  |
| GB 99           | 20.9  | 1  | Te                           | One reason that the unique_ptr constructor taking a nullptr_t argument is not explicit is to allow conversion of nullptr to unique_ptr in contexts like equality comparison. Unfortunately operator== for unique_ptr is a little more  | Add the following signatures to 20.9p1, <memory> header synopsis:<br><br>template<typename T, typename D> bool                 | ACCEPTED with MODIFICATIONS<br>See LWG 1401                 |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | clever than that, deducing template parameters for both arguments. This means that nullptr does not get deduced as unique_ptr type, and there are no other comparison functions to match.                                | <pre>operator==(const unique_ptr&lt;T, D&gt; &amp; lhs, nullptr_t);  template&lt;typename T, typename D&gt; bool operator==(nullptr_t, const unique_ptr&lt;T, D&gt; &amp; rhs);  template&lt;typename T, typename D&gt; bool operator!=(const unique_ptr&lt;T, D&gt; &amp; lhs, nullptr_t);  template&lt;typename T, typename D&gt; bool operator!=(nullptr_t, const unique_ptr&lt;T, D&gt; &amp; rhs);</pre> |  |
| GB 100          | 20.9  |  | Te                           | The unique_ptr and shared_ptr constructors taking nullptr_t delegate to a constexpr constructor, and could be constexpr themselves.  | In the 20.9.10.2 [unique.ptr.single] synopsis add "constexpr" to unique_ptr(nullptr_t).<br>In the 20.9.10.3 [unique.ptr.runtime] synopsis add "constexpr" to unique_ptr(nullptr_t).<br>In the 20.9.11.2 [util.smartptr.shared] synopsis add "constexpr" to shared_ptr(nullptr_t).   | ACCEPTED   |
| JP 85           | 20.9.1                                      |  | E                            | There are inconsistent definitions for allocator_arg. In 20.9 [memory] paragraph 1, constexpr allocator_arg_t allocator_arg = allocator_arg_t(); and in 20.9.1, const allocator_arg_t allocator_arg = allocator_arg_t(); | Change "const" to "constexpr" in 20.9.1 as follows.<br>constexpr allocator_arg_t allocator_arg = allocator_arg_t();   | ACCEPTED   |
| US 106          | 20.9.3                                      | all  | te                           | pointer_traits should have a size_type member for completeness.  | Add "typedef see below size_type;" to the generic pointer_traits template and "typedef size_t size_type;" to pointer_traits<T*>. Use pointer_traits::size_type and pointer_traits::difference_type as the defaults for allocator_traits::size_type and allocator_traits::difference_type.<br><br>See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br>See LWG 1404        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |
| GB 104          | 20.9.5.1                                    | 13   | Te                           | The ~ is missing from the invocation of the destructor of U.  | Add the missing ~ : Effects: p->~U()  | ACCEPTED  |
| GB 105          | 20.9.6                                      | 1  | Te                           | There is a missing '_' in the piecewise-construct call for pair in the class definition.  | Fix the declaration:<br>template <class T1, class T2, class... Args1, class... Args2> void construct(pair<T1, T2>* p, piecewise_construct_t, tuple<Args1...> x, tuple<Args2...> y); | REJECTED<br><br>The declaration is identical to the proposed text.  |
| US 107          | 20.9.6                                      |  | Te                           | scoped_allocator_adaptor should have its own header.  | See Appendix 1 - Additional Details   | ACCEPTED  |
| GB 101          | 20.9.10                                     | 5  | Ed                           | The first sentence of the paragraph says "Each object of a type U instantiated form the unique_ptr template..." "form" should be "from"   | Replace "form" with "from" in the opening sentence:<br>"Each object of a type U instantiated from the unique_ptr template..."   | ACCEPTED  |
| FI 13           | 20.9.10.                                    | 5  | ed                           | typo  | "form the unique_ptr" should be "from the unique_ptr."  | ACCEPTED  |
| GB 102          | 20.9.10                                     |  | Ed                           | unique_ptr is a smart pointer so [unique_ptr] should be a sub-clause of [util.smartptr]   | move [unique_ptr] to a sub-clause of [util.smartptr]  | REJECTED<br><br>[util.smartptr] deals with shared_ptr and weak_ptr, which work together. Adding unique_ptr would create a muddle. |
| JP 35           | 20.9.10.2                                   | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(38)" to "(Table 38)".  | ACCEPTED  |
| JP 36           | 20.9.10.2.1                                 | 1, 6   |                              | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are  | Change "(33)" to "(Table 33)".  | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                  |
|                 |   |  |                              | in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  |   |   |
| JP 37           | 20.9.10.2.1                                 | 18   |                              | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(34)" to "(Table 34)".  | ACCEPTED  |
| JP 38           | 20.9.10.2.3                                 | 1  |                              | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(36)" to "(Table 36)".  | ACCEPTED  |
| JP 5            | 20.9.11.2                                   | 1  | TL                           | Hash support based on ownership sharing should be supplied for shared_ptr and weak_ptr.<br><br>For two shared_ptr objects p and q, two distinct equivalence relations can be defined. One is based on equivalence of pointer values, which is derived from the expression p.get() == q.get() (hereafter called address-based equivalence relation), the other is based on equivalence of ownership sharing, which is derived from  | Add the following non-static member functions to shared_ptr and weak_ptr class template;<br>// 20.9.11.2 paragraph 1<br>namespace std{<br>template<class T> class shared_ptr {<br>public:<br>...<br>size_t owner_hash() const;<br>... | REJECTED<br><br>No consensus to make this change for this revision. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>the expression !p.owner_before(q) &amp;&amp; !q.owner_before(p) (hereafter called ownership-based equivalence relation). These two equivalence relations are independent in general. For example, a shared_ptr object created by the constructor of the signature shared_ptr(shared_ptr&lt;U&gt; const &amp;, T *) could reveal a difference between these two relations. Therefore, hash support based on each equivalence relation should be supplied for shared_ptr. However, while the standard library provides the hash support for address-based one (20.9.11.6 paragraph 2), it lacks the hash support for ownership-based one. In addition, associative containers work well in combination with the shared_ptr's ownership-based comparison but unordered associative containers don't. This is inconsistent.</p> <p>For the case of weak_ptr, hash support for the ownership-based equivalence relation can be safely defined on weak_ptrs, and even on expired ones. The absence of hash support for the ownership-based equivalence relation is fatal, especially for expired weak_ptrs. And the absence of such hash support precludes some quite effective use-cases, e.g. erasing the unordered_map entry of an expired weak_ptr key from a customized deleter supplied to shared_ptrs.</p> <p>Hash support for the ownership-based equivalence relation cannot be provided by any user-defined manner because information about ownership sharing is not available to users at all. Therefore, the only way to provide ownership-based hash support is to offer it intrusively by the standard library.</p> <p>As far as we know, such hash support is implementable. Typical implementation of such hash function could return the hash value of the pointer of the counter object that is internally managed by shared_ptr and weak_ptr.</p> | <pre>}; } // 20.9.11.3 paragraph 1 namespace std{ template&lt;class T&gt; class weak_ptr { public: ... size_t owner_hash() const; ... }; } </pre> <p>These functions satisfy the following requirements. Let p and q be objects of either shared_ptr or weak_ptr, H be a hypothetical function object type that satisfies the hash requirements (20.2.4) and h be an object of the type H. The expression p.owner_hash() behaves as if it were equivalent to the expression h(p). In addition, h(p) == h(q) must become true if p and q share ownership.</p> |  |
| CH              | 20.9.11.2                                   | p4   | te                           | Requiring shared_ptr and weak_ptr to always synchronize the use count makes it potentially slow and is inconsistent  | Strike 'not' from 'Changes in use_count() do not reflect modifications that can introduce data   | REJECTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted             |
| 20              |   |  |                              | with the general approach to leave the synchronization to the user of a facility.  | paces.' Possibly add additional synchronized constructors and assignments.   | No consensus to make this change for this revision.            |
| US 108          | 20.9.11.2.1 [util.smartptr.shared.const]    |  | te                           | shared_ptr should have the same policy for constructing from auto_ptr as unique_ptr. Currently it does not.  | Add "template <class Y> explicit shared_ptr(auto_ptr<Y>&); to [util.smartptr.shared.const] (and to the synopsis).  | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3109</a> |
| US 109          | 20.9.11.2.6                                 |  | te                           | 20.9.11.2.6 [util.smartptr.shared.create]/2 says: "the placement new expression ::new (pv) T() or ::new (pv) T(std::forward<Args>(args)...)". It should simply say "the placement new expression ::new (pv) T(std::forward<Args>(args)...)", because empty parameter packs expand to nothing. This would be consistent with the requirements in paragraph 1. | Change<br>"the placement new expression ::new (pv) T() or ::new (pv) T(std::forward<Args>(args)...)"<br>to<br>"the placement new expression ::new (pv) T(std::forward<Args>(args)...)"   | ACCEPTED   |
| GB 103          | 20.9.12                                     | 12   | Te                           | The precondition to calling declare_no_pointers is that no bytes in the range "have been previously registered" with this call. As written, this precondition includes bytes in ranges, even after they have been explicitly unregistered with a later call to 'undeclear_no_pointers'.  | Replace "have been previously registered" with "are currently registered"  | ACCEPTED   |
| GB 106          | 20.10.3                                     |  | Te                           | duration is an arithmetic type, unlike time_point, and so should provide a specialization of numeric_limits.   | Add a declaration of a partial specialization of numeric_limits for duration to the header synopsis in 20.10.<br>Add 20.3.8 [time.duration.limits]<br>"duration is an arithmetic type, and so provides an appropriate specialization of numeric_limits." | REJECTED<br>No consensus to adopt this change.                 |
| US 110          | 20.10.5                                     |  | te                           | Significant parts of the clock section are "unspecified", rather than "implementation-defined".  | Make those parts "implementation-defined".   | REJECTED<br>There is no consensus to adopt this change.        |
| US 111          | 20.10.5.2                                   | para 1                                       | te                           | What it means for monotonic_clock to be a synonym is undefined. If it may or may not be a typedef, then certain classes of programs become unportable.   | Require that it be a distinct class type.  | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3191</a> |
| GB 107          | 20.10.5.2                                   | 2  | Te                           | 1.4p9 states that which conditionally supported constructs are available should be provided in the documentation for   | Provide feature test macro for determining the presence of std::chrono::monotonic_clock. Add   | ACCEPTED with MODIFICATIONS                                    |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | the implementation. This doesn't help programmers trying to write portable code, as they must then rely on implementation-specific means to determine the availability of such constructs. In particular, the presence or absence of <code>std::chrono::monotonic_clock</code> may require different code paths to be selected. This is the only conditionally-supported library facility, and differs from the conditionally-supported language facilities in that it has standard-defined semantics rather than implementation-defined semantics. | <code>_STDCPP_HAS_MONOTONIC_CLOCK</code> to the <code>&lt;chrono&gt;</code> header, which is defined if <code>monotonic_clock</code> is present, and not defined if it is not present. | See paper <a href="#">N3191</a>   |
| DE 20           | 20.10.5.2                                   |  | te                           | The library component <code>monotonic_clock</code> is conditionally supported, but no compile-time flag exists that allows user-code to query its existence. Further-on there exist no portable means to simulate such a query. (To do so, user code would be required to add types to namespace <code>std::chrono</code> .)  | Provide a compile-time flag (preferably a macro) that can be used to query the existence of <code>monotonic_clock</code> .   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3191</a>  |
| CH 21           | 20.10.5.2                                   | p2   | te                           | Monotonic clocks are generally easy to provide on all systems and are implicitly required by some of the library facilities anyway.   | Make monotonic clocks mandatory, i.e. remove p2. Also change 30.2.4p2 accordingly.   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3191</a>  |
| US 112          | 20.10.5.3                                   | para 1                                       | te                           | What it means for <code>high_resolution_clock</code> to be a synonym is undefined. If it may or may not be a typedef, then certain classes of programs become unportable.   | Require that it be a distinct class type.  | REJECTED<br><br>The standard is correct as written. This is not a defect. There are a number of places in the standard where we allow implementations to choose their preferred technique, the most obvious example being the <code>iterator/const_iterator</code> types of <code>set</code> . Typically, this means it is not portable to declare function overloads that differ only in their use of these types. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted          |
| JP 39           | 21.2.2                                      | 4  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(35)" to "(Table 35)".  | ACCEPTED  |
| GB 108          | 21.2.3.1                                    | 2,3  | Te                           | The definition of streamoff/streampos defers to the definition of off_type and pos_type in 21.2.2, which defers back to 27.2.2 for the definition of streamoff/streampos. The actual definition appears to be supplied in 27.3, the synopsis of <iosfwd>.   | Update the reference in 21.2.3.1 to refer forward to 27.2.2, rather than back to 21.2.2. Add a cross-reference to from 27.2.2 to 27.3.  | ACCEPTED  |
| GB 109          | 21.2.3.2/3/4                                |  | Te                           | It is not clear what the specification means for u16streampos, u32streampos or wstreampos when they refer to the requirements for POS_T in 21.2.2, as there are no longer any such requirements. Similarly the annex D.7 refers to the requirements of type POS_T in 27.3 that no longer exist either.  | Clarify the meaning of all cross-reference to the removed type POS_T.   | ACCEPTED with MODIFICATIONS<br><br>See LWG 1414             |
| JP 40           | 21.4  | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(96)" to "(Table 96)".  | ACCEPTED  |
| GB 110          | 21.4.7.1                                    |  | Te                           | data() is the function of choice for calling into legacy 'C'-like APIs. Both vector and array designed this function to be callable in a const-correct way while allowing for functions that want to use the result to designate a return buffer.   | Add the following overload to basic_string data():<br>charT * data();<br>Relax the requirement that programs do not alter values in the array through the pointe retrieved through this new overload. | REJECTED<br><br>There is no consensus to adopt this change. |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted             |
| GB 111          | 21.5  |  | Te                           | Section 17.6.4.8, Data Race Avoidance, requires the C++ Standard Library to avoid data races that might otherwise result from two threads making calls to C++ Standard Library functions on distinct objects. The C standard library is part of the C++ Standard Library and some C++ Standard library functions (parts of the Localization library, as well as Numeric Conversions in 21.5), are specified to make use of the C standard library. Therefore, the C++ standard indirectly imposes a requirement on the thread safety of the C standard library. However, since the C standard does not address the concept of thread safety conforming C implementations exist that do not provide such guarantees. This conflict needs to be reconciled. | remove the requirement to make use of strtol() and sprintf() since these functions depend on the global C locale and thus cannot be made thread safe. | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3278</a> |
| JP 86           | 21.7  | 1  | E                            | Table numbers are listed incorrectly. "74,75. and" should be "74, 75, and".   | Correct typo as follows.<br>Tables 71, 72, 73, 74, 75, and 76 describe  | ACCEPTED   |
| JP 87           | 22.3.1                                      |  | E                            | While usage of "traits" and "Traits" are explained in 21.2 as template parameters and arguments respectively, specifying "Traits" as template parameter seems misuse.<br>template <class charT, class Traits, class Allocator><br>bool operator()(const<br>basic_string<charT,Traits,Allocator>& s1,<br>const basic_string<charT,Traits,Allocator>& s2) const;  | Change "Traits" to "traits" in three places.  | ACCEPTED   |
| JP 88           | 22.6  | 3, Table 91                                  | E                            | Typo, duplicated "ispunct" and missing "iswpunct".  | Correct as follows.<br>isprint<br>ispunct<br>isspace<br>:<br>iswprint<br>iswpunct<br>iswspace   | ACCEPTED   |
| JP 89           | 23.1  | 2, Table 92                                  | E                            | Typo, "<forwardlist>" should be "<forward_list>".   | Correct typo.<br><forward_list>   | ACCEPTED   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)   | 4                            | 5  | (6)  | (7)  |
|-----------------|---|---|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1)                          | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                                       |
| US 113          | 23.2  | 5   | te                           | Resolve LWG 579 one way or the other, but preferably in the direction of changing the two erase overloads to return void.  |  | REJECTED<br><br>There is no consensus to adopt this change.                              |
| US 114          | 23.2.1                                      | Paragraph 9   | te                           | Requirements on iterators swapping allegiance would disallow the small-string optimization.  | Add an exclusion for <code>basic_string</code> to the sentence beginning "Every iterator referring to an element...". Add a sentence to 21.4.6.8/2 saying that iterators and references to string elements remain valid, but it is not specified whether they refer to the same string or the other string.                | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3108</a>                       |
| DE 21           | 23.2.1, 23.3.3.4                            |   | te                           | 23.2.1/11 provides a general no-throw guarantee for <code>erase()</code> container functions, exceptions from this are explicitly mentioned for individual containers. Because of its different name, <code>forward_list</code> 's <code>erase_after()</code> function is not ruled by this but should so.   | Add a "Throws: Nothing" clause to both <code>erase_after</code> overloads in 23.3.3.4.   | ACCEPTED   |
| US 115          | 23.2.1                                      | Paragraph 15  | te                           | The terms <code>CopyConstructible</code> , <code>MoveConstructible</code> , and <code>constructible from</code> are redefined, then used inconsistently and often incorrectly within the section. New terms should have been introduced and used correctly.  | Better terms would be <i>X can copy-insert T</i> , <i>X can move-insert T</i> , and <i>X can construct-insert T with args</i> . See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3173</a>                       |
| US 116          | 23.2.1                                      | Table 96  | ed                           | The requirement for <code>X(rv)</code> that move construction of the allocator not throw can be misread as requiring that move construction of the whole container not throw.  | Add non-normative note:<br><br><i>Requires: move construction of A shall not exit via an exception. [Note: This requirement on allocators exists so that implementations can (optionally) provide a nothrow guarantee on move construction of some containers. – end note]</i>   | REJECTED<br><br>The text as is seems quite clear; the proposed note just muddles things. |
| FI 12           | 23.2.3 [sequence.requirements]              | Table 97 — Sequence container requirements (in addition to container) | te                           | The requirement for <code>insert</code> says: "Requires:T shall be <code>CopyConstructible</code> . For vector and deque, T shall also be <code>CopyAssignable</code> ". Why must T be <code>CopyAssignable</code> ? Is it for cases where the object may be already constructed and <code>insert</code> will first copy the existing data out of the way and assign to an existing object? Can't such an implementation do the same with <code>placement-new</code> ? | <code>CopyAssignable</code> seems like an overly strict requirement for <code>insert</code> , exposing implementation details in the specification. If such implementation details are achievable without assignment, eg. with <code>placement-new</code> , the <code>CopyAssignable</code> requirement should be removed. | REJECTED<br><br>No consensus to make the change.   |
| JP              | 23.2.3                                      | 17, Table 98  | E                            | In Operational semantics for "a.emplace_front(args)",  | Change <code>&lt;Arg&gt;</code> to <code>&lt;Args&gt;</code> .   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)   | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| 90              |   |  |                              | <Arg> should be <Args>. Prepends an object of type T constructed with std::forward<Arg>(args)...."  |   |  |
| JP 91           | 23.2.5  | 10, Table 100                                | E                            | Typo, unnecessary new-line. a_uniq. emplace(args)<br>Typo, unnecessary space. a_eq. emplace(args)   | Remove space characters between "." and "emplace" same as other word wrapped columns in the table.  | ACCEPTED with MODIFICATIONS<br><br>The newline is needed, but the space is unnecessary. Removed the space. |
| JP 92           | 23.2.5  | 11   | E                            | Typo, "unodified" should be "unmodified".   | Correct typo. unmodified  | ACCEPTED   |
| ES 2            | 23.2.5 [unord.req], Table 100 (Unordered associative container requirements (in addition to container)) | Row for expression a.erase(q)                | Te                           | The expression is required to return the iterator immediately following q prior to the erasure. As explained in N2023, this requirement makes it impossible to achieve average O(1) complexity for unordered associative containers implemented with singly linked lists. This has a theoretical as well as a practical impact, as reported by users of early implementations of these containers. Discussions among committee members have not found any way of remedying this deficiency (other than acknowledging it) by some smart modification of usual singly linked lists implementations. | Change the return type of the expression from iterator to void. Eliminate the sentence "Return value is the iterator immediately following q prior to the erasure". Change accordingly the appearances of "iterator erase(const_iterator position)" in 23.5.1, 23.5.2, 23.5.3 and 23.5.4. | REJECTED<br>There was no consensus to adopt this change.   |
| GB 112          | 23.3.1.7  | p3   | Te                           | Should the effect of calling front/back on a zero sized array really be implementation defined i.e. require the implementor to define behaviour?  | Change "implementation defined" to "undefined"  | ACCEPTED   |
| GB 113          | 23.3.2.2  | p1   | Te                           | There is no mention of what happens if sz==size(). While it obviously does nothing I feel a standard needs to say this explicitly.  | Append "If sz == size(), does nothing" to the effects.  | ACCEPTED with MODIFICATIONS<br><br>See LWG 1418  |
| US 117          | 23.3.3  |  | Te                           | forward_list::erase_after should return an iterator.  | See Appendix 1 - Additional Details   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| JP 41           | 23.3.3                                      | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(table 93)" to "(Table 93)".  | ACCEPTED   |
| GB 114          | 23.3.4.1                                    | p11 & p12                                    | Ed                           | It looks like the erase/insert effects given in p11 are intended for p12.   | Move the erase/insert effects down to p12   | ACCEPTED with MODIFICATIONS<br><br>The code in question is a remnant from a previous version. It has now been removed. |
| GB 115          | 23.3.4.2                                    | p1   | Te                           | There is no mention of what happens if sz==size(). While it obviously does nothing I feel a standard needs to say this explicitly.  | Express the semantics as pseudo-code similarly to the way it is done for the copying overload that follows (in p3). Include an else clause that does nothing and covers the sz==size() case.                  | ACCEPTED with MODIFICATIONS<br><br>See LWG 1420  |
| GB 116          | 23.3.5                                      |  | Ed                           | The sequence container adaptors consume sequence containers, but are neither containers nor sequences themselves. While they clearly belong in clause 23, they should not interrupt the presentation of the sequence container themselves.  | Move clause 23.3.5 out of clause 23.3. Recommending inserting as a 'new' 23.4 immediately following sequence containers, and before the current 23.4 (associative containers) which would be renumbered 23.5. | ACCEPTED   |
| DE 22           | 23.3.5.1, 23.3.5.2, 23.3.5.3                |  | te                           | With the final acceptance of move operations as special members and introduction of corresponding suppression rules of implicitly generated copy operations the some library types that were copyable in C++03 are no longer copyable (only movable) in C++03, among them queue, priority_queue, and stack.   |   | ACCEPTED with MODIFICATIONS<br><br>See LWG 1421  |
| JP 93           | 23.3.5.2                                    | 1  | E                            | Typo, missing ";".<br>template <class... Args> void emplace(Args&&... args)   | Correct typo.<br>template <class... Args> void emplace(Args&&... args);   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                        |
| GB 117          | 23.3.6.2                                    | p9   | Te                           | (Same as for 23.3.2.2p1 i.e. deque::resize). There is no mention of what happens if sz==size(). While it obviously does nothing I feel a standard needs to say this explicitly.   | Append "If sz == size(), does nothing" to the effects.   | ACCEPTED with MODIFICATIONS<br>See LWG 1525                               |
| GB 118          | 23.3.7                                      |  | Te                           | vector<bool> iterators are not random access iterators because their reference type is a special class, and not 'bool &'. All standard library operations taking iterators should treat this iterator as if it was a random access iterator, rather than a simple input iterator.         | Either revise the iterator requirements to support proxy iterators (restoring functionality that was lost when the Concept facility was removed) or add an extra paragraph to the vector<bool> specification requiring the library to treat vector<bool> iterators as-if they were random access iterators, despite having the wrong reference type. | REJECTED<br>There is no consensus to adopt this change for this revision. |
| JP 6            | 23.4.1                                      | 2  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.<br>namespace std {<br>template <class Key, class T, class Compare = less<Key>,<br>class Allocator = allocator<pair<const Key, T> > ><br>class map {<br>public:<br>...<br>map(const Allocator&); | Add explicit.<br>namespace std {<br>template <class Key, class T, class Compare = less<Key>,<br>class Allocator = allocator<pair<const Key, T> > ><br>class map {<br>public:<br>...<br>explicit map(const Allocator&);   | ACCEPTED  |
| JP 7            | 23.4.2                                      | 2  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.   | Add explicit.<br>namespace std {<br>template <class Key, class T, class Compare = less<Key>,<br>class Allocator = allocator<pair<const Key, T> > ><br>class multimap {<br>public:<br>...<br>explicit multimap(const Allocator&);   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted      |
| JP 8            | 23.4.3                                      | 2  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.  | Add explicit.<br><pre>namespace std { template &lt;class Key, class Compare = less&lt;Key&gt;, class Allocator = allocator&lt;Key&gt; &gt; class set { public: ... explicit set(const Allocator&amp;);</pre>   | ACCEPTED  |
| JP 9            | 23.4.4                                      | 2  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.  | Add explicit.<br><pre>namespace std { template &lt;class Key, class Compare = less&lt;Key&gt;, class Allocator = allocator&lt;Key&gt; &gt; class multiset { public: ... explicit multiset(const Allocator&amp;);</pre>   | ACCEPTED  |
| US 118          | 23.5  |  | te                           | Some unordered associative container operations have undesirable complexities when the container is implemented using singly linked lists. | See Appendix 1 - Additional Details  | REJECTED<br>There is no consensus to adopt this change. |
| JP 10           | 23.5.1                                      | 3  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.  | Add explicit.<br><pre>namespace std { template &lt;class Key, template &lt;class Key, class T, class Hash = hash&lt;Key&gt;, class Pred = std::equal_to&lt;Key&gt;, class Alloc = std::allocator&lt;std::pair&lt;const Key, T&gt; &gt; &gt; class unordered_map { public: ... </pre> | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | explicit unordered_map(const Allocator&);   |  |
| JP 11           | 23.5.2                                      | 3  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.                                | Add explicit.<br>namespace std {<br>template <class Key,<br>class T,<br>class Hash = hash<Key>,<br>class Pred = std::equal_to<Key>,<br>class Alloc = std::allocator<std::pair<const Key,<br>T> > ><br>class unordered_multimap<br>{<br>public:<br>...<br>explicit unordered_multimap(const Allocator&); | ACCEPTED   |
| JP 94           | 23.5.2                                      | 1  | E                            | "see below" should be in italic and need one space between words.<br>explicit unordered_multimap(size_type n = seebelow, | Change to:<br>explicit unordered_multimap(size_type n = see below,  | ACCEPTED   |
| JP 12           | 23.5.3                                      | 3  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.                                | Add explicit.<br>namespace std {<br>template <class Key,<br>class Hash = hash<Key>,<br>class Pred = std::equal_to<Key>,<br>class Alloc = std::allocator<Key> ><br>class unordered_set<br>{<br>public:<br>...<br>explicit unordered_set(const Allocator&);   | ACCEPTED   |
| JP 13           | 23.5.4                                      | 3  | TL                           | Constructor accepting an allocator as a single parameter should be qualified as explicit.                                | Add explicit.<br>namespace std {<br>template <class Key,  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              |   | <pre> class Hash = hash&lt;Key&gt;, class Pred = std::equal_to&lt;Key&gt;, class Alloc = std::allocator&lt;Key&gt; &gt; class unordered_multiset { public: ... explicit unordered_multiset(const Allocator&amp;); </pre> |  |
| US 119          | [input.iterators ]<br>24.2.3                | Table 104                                    | te                           | Although the section talks about operator==, there is no requirement that it exist.   | Add a == b to Table 104  | REJECTED<br>Standard is correct as is.             |
| JP 42           | 25.1  | 8, 9   | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(4)" to "(Clause 4)".  | ACCEPTED   |
| US 120          | 25.2.12                                     | para 1                                       | te                           | is_permutation is underspecified for anything but the simple case where both ranges have the same value type and the comparison function is an equivalence relation.  | Restrict is_permutation to the case where it is well specified. See Appendix 1 - Additional Details  | ACCEPTED   |
| ES 3            | 25.2.12                                     |  | Te                           | is_permutation does not require ForwardIterator1 and ForwardIterator2 to have the same value type. This opens the door to nonsense heterogeneous usage where both ranges have different value types   | Require both iterator types to have the same value type  | ACCEPTED   |
| JP 43           | 25.3.9                                      | 5  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the   | Change "(4)" to "(Clause 4)".  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              | form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.   |  |  |
| US 121          | 25.3.12 [alg.random.shuffle]                | 1  | te                           | random_shuffle and shuffle should be consistent in how they accept their source of randomness: either both by rvalue reference or both by lvalue reference.   | Change random_shuffle to accept its RandomNumberGenerator by lvalue reference. | ACCEPTED with MODIFICATIONS<br>See LWG 1432        |
| GB 119          | 25.3.12                                     |  | Te                           | The functions random_shuffle and shuffle both take arguments providing a source of randomness, but one take its argument by rvalue reference, and the other requires an lvalue reference. The technical merits of which form of argument passing should be settled for this specific case, and a single preferred form used consistently.   | [DEPENDS ON WHETHER RVALUE OR LVALUE REFERENCE IS THE PREFERRED FORM]          | ACCEPTED with MODIFICATIONS<br>See LWG 1433        |
| JP 44           | 25.4  | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(4)" to "(Clause 4)".  | ACCEPTED   |
| JP 45           | 25.4.7                                      | 1, 10, 19                                    | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(32)" to "(Table 32)".   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
| US 122          | 25.4.7 [alg.min.max]                        |  | te                           | It was the LWG's intent in Pittsburgh that N2772 be applied to the WP  | Apply N2772 to the WP.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3106</a> |
| US 123          | 25.5  | ¶5b, and all uses of lshift                  | ed                           | N3056, as adopted, calls for each use of lshift to be followed by a subscripted value.   | Adjust all occurrences of the lshift notation so as to match the notation of N3056.            | ACCEPTED   |
| FI 14           | 26.3.1.                                     | 3  | ed                           | typo   | "floating-point environmnet" should be "floating-point environment."                           | ACCEPTED   |
| GB 120          | 26.4.7                                      |  | Ge                           | The complex number functions added for compatibility with the C99 standard library are defined purely as a cross-reference, with no hint of what they should return. This is distinct from the style of documentation for the functions in the earlier standard. In the case of the inverse-trigonometric and hyperbolic functions, a reasonable guess of the functionality may be made from the name, this is not true of the cproj function, which apparently returns the projection on the Reimann Sphere. A single line description of each function, associated with the cross-reference, will greatly improve clarity. | [ONE LINE DESCRIPTIONS, AND ASSOCIATED PARAGRAPH NUMBERS, TO FOLLOW IF THE INTENT IS APPROVED] | ACCEPTED with MODIFICATIONS<br><br>See LWG 1435                    |
| JP 46           | 26.5.1.6                                    | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  | Change "Clause 21 and 27" to "Clauses 21 and 27".  | ACCEPTED   |
| GB 121          | 26.5.3                                      |  | Te                           | All the random number engine types in this clause have a constructor taking an unsigned integer type, and a constructor template for seed sequences. This means that an attempt to create a random number engine seeded by an integer literal must remember to add the appropriate unsigned suffix to the literal, as a signed integer will  | [WORDING TO FOLLOW ONCE A PREFERRED DIRECTION IS INDICATED]                                    | ACCEPTED with MODIFICATIONS<br><br>See LWG 1436                    |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|--|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)  | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |  |  |                              | attempt to use the seed sequence template, yielding undefined behaviour, as per 26.5.1.1p1a. It would be helpful if at least these anticipated cases produced a defined behaviour, either an erroneous program with diagnostic, or a conversion to unsigned int forwarding to the appropriate constructor. |   |  |
| US 124          | 26.5.3.2   | ¶4   | te                           | The Mersenne twister algorithm is meaningless for word sizes less than two, as there are then insufficient bits available to be "twisted".   | Insert the following among the relations that are required to hold: $2u < w$ .  | ACCEPTED   |
| US 125          | 26.5.4.1 [rand.adapt.di sc]  | 3  | ed                           | The synopsis for min() and max() is lacking "(" in the return statements.  | return Engine::min();<br>return Engine::max();  | ACCEPTED   |
| US 126          | 26.5.4.1 [rand.adapt.di sc],<br>26.5.4.2 [rand.adapt.ibi ts],<br>26.5.4.3 [rand.adapt.sh uf] | 3  | te                           | Each adaptor has a member function called base() which has no definition.  | Give it the obvious definition.   | ACCEPTED with MODIFICATIONS<br><br>See LWG 1438    |
| US 127          | 26.5.4.1   | synopsis after ¶3                            | ed/te                        | Engine::min is a function and ought be invoked in the context where mentioned, as should Engine::max.  | Append parentheses so as to become return Engine::min() and return Engine::max().   | ACCEPTED   |
| US 128          | 26.5.4.3   | synopsis after ¶3                            | ed/te                        | Engine::min is a function and ought be invoked in the context where mentioned, as should Engine::max.  | Append parentheses so as to become return Engine::min() and return Engine::max().   | ACCEPTED   |
| US 129          | 26.5.7.1   | ¶8b  | ed/te                        | The expression begin[x+q] is incorrect since x is unspecified .  | Replace x by k so as to obtain begin[k+q].  | ACCEPTED   |
| US 130          | 26.5.7.1   | ¶8c  | ed/te                        | The phrase "three more times" is misleading.   | s/three more times,/again,/   | ACCEPTED   |
| US 131          | 26.5.7.1   | ¶8c  | ed/te                        | Values $r_3$ and $r_4$ are correctly specified, but in their subsequent use they are interchanged with respect to the original algorithm by Mutso Saito.   | Exchange subscripts so as to read as follows: "... update begin[k + p] by xoring it with $r_3$ , update begin[k + q] by xoring it with $r_4$ , and ..." | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|--|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)                                      | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted |
| US 132          | 26.5.7.1<br>[rand.util.seed seq]   | 8b   | ed                           | The last sentence includes "begin[x + q]" but "x" is non-sensical here.  | Change "begin[x + q]" to "begin[k + q]"  | ACCEPTED   |
| US 133          | 26.5.7.1<br>[rand.util.seed seq]   | 8c   | te                           | The use of r3 and r4 is reversed in the final sentence of 8c according to the defect report on comp.std.c++ that this specification is based on:<br><a href="http://groups.google.com/group/comp.std.c++/browse_thread/thread/e34cbee1932efdb8/aad523dccec12aed?q=group:comp.std.c%2B%2B+insubject:seed_seq">http://groups.google.com/group/comp.std.c++/browse_thread/thread/e34cbee1932efdb8/aad523dccec12aed?q=group:comp.std.c%2B%2B+insubject:seed_seq</a><br><br>If you follow the SFMT link to the software, the software also uses r3 and r4 in a manner inconsistent with N3092. I believe N3092 should be changed to be consistent with the defect report. | Change 8c to end:<br><br>... update begin[k + p] by xoring it with r3, update begin[k + q] by xoring it with r4, and ...   | ACCEPTED   |
| US 134          | 26.5.8.5.2<br>[rand.dist.samp.pconst],<br>26.5.8.5.3<br>[rand.dist.samp.plinear] |  | te                           | These two distributions have a member function called densities() which returns a vector<double>. The distribution is templated on RealType. The distribution also has another member called intervals() which returns a vector<RealType>. Why doesn't densities return vector<RealType> as well? If RealType is long double, the computed densities property isn't being computed to the precision the client desires. If RealType is float, the densities vector is taking up twice as much space as the client desires.   | Change the piecewise constant and linear distributions to hold / return the densities in a vector<result_type>.<br><br>If this is not done, at least correct[rand.dist.samp.pconst]/13 which describes the return of densities as a vector<result_type>. | ACCEPTED   |
| US 135          | 26.5.8.5.3   | 10   | te                           | This paragraph says: Let $b_k = x_{min} + k \cdot \delta$ for $k = 0, \dots, n$ , and $w_k = fw(b_k + \delta)$ for $k = 0, \dots, n$ .<br><br>However I believe that $fw(b_k)$ would be far more desirable. I strongly suspect that this is nothing but a type-o.  | Change p10 to read:<br><br>Let $b_k = x_{min} + k \cdot \delta$ for $k = 0, \dots, n$ , and $w_k = fw(b_k)$ for $k = 0, \dots, n$ .  | ACCEPTED   |
| JP 47           | 26.7.1   | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br><br>However there are some notations "(Z)" consisting of only   | Change "(35)" to "(Table 35)".   | ACCEPTED   |

<sup>1</sup> MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

<sup>2</sup> Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|--|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)                  | Paragraph/ Figure/Table/ Note (e.g. Table 1)                 | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |  |  |                              | a number Z in parentheses to confer Clause or Table number Z.   |  |  |
| JP 48           | 26.7.2   | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(35)" to "(Table 35)".   | ACCEPTED   |
| JP 49           | 26.7.4   | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(36)" to "(Table 36)".   | ACCEPTED   |
| US 136          | 26.8   |  | Te                           | Floating-point test functions are incorrectly specified.  | See Appendix 1 - Additional Details  | ACCEPTED   |
| CA 9            | 27.2.3p2<br>30.3.1.2p6<br>30.3.1.5p7<br>30.6.4p7<br>30.6.9p5 | 27.2.3p2<br>30.3.1.2p6<br>30.3.1.5p7<br>30.6.4p7<br>30.6.9p5 | te                           | Imposed happens-before edges should be in synchronizes-with<br><br>Each use of the words "happens-before" should be replaced with the words "synchronizes-with" in the following sentences:<br><br>27.2.3p2<br><br>30.3.1.2p6   | Each use of the words "happens-before" should be replaced with the words "synchronizes-with" in the following sentences:<br><br>27.2.3p2<br><br>30.3.1.2p6<br><br>30.3.1.5p7 | ACCEPTED<br><br>See paper <a href="#">N3196</a>    |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 | 30.6.10.1p23                                | 30.6.10.1p23                                 |                              | <p>30.3.1.5p7</p> <p>30.6.4p7</p> <p>30.6.9p5</p> <p>30.6.10.1p23</p> <p>Rationale: Happens-before is defined in 1.10p11 in a way that (deliberately) does not make it explicitly transitively closed. Adding edges to happens-before directly, as in 27.2.3p2 etc., does not provide transitivity with sequenced-before or any other existing happens-before edge. This lack of transitivity seems to be unintentional.</p> | <p>30.6.4p7</p> <p>30.6.9p5</p> <p>30.6.10.1p23</p>   |  |
| GB 122          | 27, 30                                      |  | Te                           | See (D) in attachment Appendix 1 - Additional Details  | Request the concurrency working group to determine if changes are needed  | ACCEPTED with MODIFICATIONS<br>See LWG 1442, and LWG 1443  |
| GB 123          | 27.5.3.2                                    | Table 124                                    | Te                           | Several rows in table 124 specify a Return type of 'OFF_T', which does not appear to be a type defined in this standard.   | Resolve outstanding references to the removed type 'OFF_T'.   | ACCEPTED with MODIFICATIONS<br>See LWG 1414, and LWG 1444  |
| US 137          | 27.7  |  | Te                           | Several iostreams member functions are incorrectly specified.  | See Appendix 1 - Additional Details   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3168</a>   |
| US 138          | 27.7  |  | te                           | For istreams and ostream, the move-constructor does not move-construct, the move-assignment operator does not move-assign, and the swap function does not swap because these operations do not manage the rdbuf() pointer. Useful applications of these operations are prevented both by their incorrect semantics and because   | In short: reverse the resolution of issue 900, then change the semantics to move and swap the rdbuf() pointer. Add a new protected constructor that takes an rvalue reference to a stream and a pointer to a streambuf, a new protected assign() operator that takes the same arguments, and a new protected partial_swap() function that doesn't | REJECTED<br>The Library Working Group reviewed n3179 and concluded that this change alone was not sufficient, as it would require changes to |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
|                 |   |  |                              | they are protected.   | swap rdbuf().<br>See Appendix 1 - Additional Details   | some of the derived stream types in the library. The preference is not make such a broad fix, and retain the current semantic. It would be difficult to rename the new functions introduced in the C++0x revision of the standard at a later date. Thus there is no consensus to make this change.<br>See paper <a href="#">N3179</a> |
| US 139          | 27.7  | 1.1.3  | te                           | Resolve issue LWG 1328 one way or the other, but preferably in the direction outlined in the proposed resolution, which, however, is not complete as-is: in any case, the sentry must not set ok_ = false if is.good() == false, otherwise istream::seekg, being an unformatted input function, does not take any action because the sentry object returns false when converted to type bool. Thus, it remains impossible to seek away from end of file.      |  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3168</a>  |
| US 140          | 27.8  |  | te                           | It should be possible to construct a stringstream with a specific allocator.  | Add an allocator_type and overloaded constructors that take an Allocator argument to basic_stringbuf, basic_istream, basic_ostringstream, and basic_stringstream. The semantics of allocator propagation should be the same as if the stringbuf contained an embedded basic_string using the same allocator. | REJECTED<br><br>There is no consensus to adopt this change for this revision.   |
| GB 124          | 27.8.1.3                                    | 3  | Te                           | N3092 27.8.1.3 Member functions contains this text specifying the postconditions of basic_stringbuf::str(basic_string): "Postconditions: If mode & ios_base::out is true, pbase() points to the first underlying character and epptr() >= pbase() + s.size() holds; in addition, if mode & ios_base::in is true, pptr() == pbase() + s.data() holds, otherwise pptr() == pbase() is true. [...]"<br>Firstly, there's a simple mistake: It should be pbase() + |  | ACCEPTED with MODIFICATIONS<br>See LWG 1448   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)                       | (7)  |
|-----------------|---|--|------------------------------|--|---------------------------|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB | Secretariat observations on each comment submitted |
|                 |   |  |                              | <p>s.length(), not pbase() + s.data().</p> <p>Secondly, it doesn't match existing implementations. As far as I can tell, GCC 4.5 does not test for mode &amp; ios_base::in in the second part of that sentence, but for mode &amp; (ios_base::app   ios_base_ate), and Visual C++ 9 for mode &amp; ios_base::app. Besides, the wording of the C++0x draft doesn't make any sense to me. I suggest changing the second part of the sentence to one of the following:</p> <p>Replace ios_base::in with (ios_base::ate   ios_base::app), but this would require Visual C++ to change (replacing only with ios_base::ate would require GCC to change, and would make ios_base::app completely useless with stringstreams):</p> <p>in addition, if mode &amp; (ios_base::ate   ios_base::app) is true, pptr() == pbase() + s.length() holds, otherwise pptr() == pbase() is true.</p> <p>Leave pptr() unspecified if mode &amp; ios_base::app, but not mode &amp; ios_base::ate (implementations already differ in this case, and it's always possible to use ios_base::ate to get the effect of appending, so it's not necessary to require any implementation to change):</p> <p>in addition, if mode &amp; ios_base::ate is true, pptr() == pbase() + s.length() holds, if neither mode &amp; ios_base::ate nor mode &amp; ios_base::app is true, pptr() == pbase() holds, otherwise pptr() &gt;= pbase() &amp;&amp; pptr() &lt;= pbase() + s.length() (which of the values in this range is unspecified).</p> <p>Slightly stricter:</p> <p>in addition, if mode &amp; ios_base::ate is true, pptr() == pbase() + s.length() holds, if neither mode &amp; ios_base::ate nor mode &amp; ios_base::app is true, pptr() == pbase() holds, otherwise pptr() == pbase()    pptr() == pbase() + s.length() (which of these two values is unspecified).</p> <p>A small table might help to better explain the three cases.</p> <p>BTW, at the end of the postconditions is this text: "egptr() == eback() + s.size() hold". Is there a preference for</p> |                           |  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                   |
|                 |   |  |                              | basic_string::length or basic_string::size? It doesn't really matter, but it looks a bit inconsistent.  |  |  |
| CA 4            | 27.8.2                                      | various                                      | te                           | Subclause 27.9.2 [c.files] specifies that < cintypes > has declarations for abs() and div(); however, the signatures are not present in this subclause. The signatures proposed under TR1 ([tr.c99.inttypes]) are not present in FCD (unless if intmax_t happened to be long long). It is unclear as to which, if any of the abs() and div() functions in [c.math] are meant to be declared by < cintypes >. This subclause mentions imaxabs() and imaxdiv(). These functions, among other things, are not specified in FCD to be the functions from Subclause 7.8 of the C Standard. Finally, < cintypes > is not specified in FCD to include < cstdint > (whereas < inttypes.h > includes < stdint.h > in C). | Please clarify.  | ACCEPTED with MODIFICATIONS<br><br>See LWG 1449                      |
| JP 14           | 28.4  |  | TL                           | Support of char16_t/char32_t is insufficient. The < regex > does not have typedefs for char16_t/char32_t.<br><br>The reason we need this typedefs is, because anybody may define exactly same type with different typedef names. Doesn't < locale > offer enough operations which is required by regex implementation?  | Add typedefs below<br>typedef basic_regex<char16_t> u16regex;<br>typedef basic_regex<char32_t> u32regex;   | REJECTED<br><br>No consensus to adopt this change for this revision. |
| GB 127          | 28.5.2                                      |  | Te                           | The Bitmask Type requirements in 17.5.2.1.3 p3 say that all elements on a bitmask type have distinct values, but 28.5.2 defines regex_constants::match_default and regex_constants::format_default as elements of the bitmask type regex_constants::match-flag_type, both with value 0. This is a contradiction.  | One of the bitmask elements should be removed from the declaration and should be defined separately, in the same manner as ios_base::adjustfield, ios_base::basefield and ios_base::floatfield are defined by 27.5.2.1.2p2 and Table 120. These are constants of a bitmask type, but are not distinct elements, they have more than one value set in the bitmask. regex_constants::format_default should be specified as a constant with the same value as regex_constants::match_default. | REJECTED<br><br>No consensus to make a change at this time.          |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
| JP 50           | 28.5.2                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 136" to "Table 136".  | ACCEPTED   |
| JP 51           | 28.5.3                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 137" to "Table 137".  | ACCEPTED   |
| US 141          | 28.8  |  | te                           | std::basic_regex should have an allocator for all the reasons that a std::string does. For example, I can use boost::interprocess to put a string or vector in shared memory, but not a regex.  | Add allocators to regexes   | REJECTED<br><br>There is no consensus for a change at this time.   |
| GB 125          | 28.10.3                                     | 2  | Te                           | The term "target sequence" is not defined.  | Replace "target sequence" with "string being searched/matched"  | REJECTED<br><br>The standard is correct as written. The term 'target sequence' is a common term used in Regular Expressions. A simple search will provide examples of its use. |
| GB 126          | 28.10.3                                     |  | Te                           | It's unclear how match_results should behave if it has been default-constructed. The sub_match objects returned by operator[], prefix and suffix cannot point to the  | Add to match_results::operator[], match_results::prefix and match_results::suffix: Requires: !empty() | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | end of the sequence that was searched if no search was done. The iterators held by unmatched sub_match objects might be singular.   |   | See paper <a href="#">N3158</a>                    |
| JP 52           | 28.11.2                                     | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 139" to "Table 139".                                      | ACCEPTED   |
| JP 95           | 28.11.3                                     | 1  | E                            | The section number "(24.1.4)" for "Bidirectional Iterator" is wrong. The correct one is "(24.2.6)". In addition, it is written as normal text, but it should be embedded as a link to the section.  | Change "(24.1.4)" to "(24.2.6)" and make it a link to section (24.2.6). | ACCEPTED   |
| JP 53           | 28.11.3                                     | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 140" to "Table 140".                                      | ACCEPTED   |
| JP 54           | 28.13                                       | 6  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 135" to "Table 135".                                      | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1)             | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
| GB 128          | 29  |  | Ge                           | WG14 has made some late changes to their specification of atomics, and care should be taken to ensure that we retain a common subset of language/library syntax to declare headers that are portable to both languages. Ideally, such headers would not require users to define their own macros, especially not macros that map to keywords (which remains undefined behaviour) | Depends on result of the review of WG14 work, which is expected to be out to ballot during the time wg21 is resolving its own ballot comments. Liaison may also want to file comments in WG14 to ensure compatibility from both sides.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| CH 22           | 29  |  | te                           | WG14 currently plans to introduce atomic facilities that are intended to be compatible with the facilities of clause 29. They should be compatible.  | Make sure the headers in clause 29 are defined in a way that is compatible with the planned C standard.  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| GB 129          | 29  | Table 143  | Te                           | Table 143 lists the typedefs for various atomic types corresponding to the various standard integer typedefs, such as atomic_int_least8_t for int_least8_t, and atomic_uint_fast64_t for uint_fast64_t. However, there are no atomic typedefs corresponding to the fixed-size standard typedefs int8_t, int16_t, and so forth.   | Add the following entries to table 143:<br><br>atomic_int8_t => int8_t (optional),<br>atomic_int16_t => int16_t (optional),<br>atomic_int32_t => int32_t (optional),<br>atomic_int64_t => int64_t (optional),<br>atomic_uint8_t => uint8_t (optional),<br>atomic_uint16_t => uint16_t (optional),<br>atomic_uint32_t => uint32_t (optional),<br>atomic_uint64_t => uint64_t (optional)<br><br>These typedefs should be available if the corresponding typedefs from are available. | REJECTED<br><br>There was no consensus to make this change.        |
| CA 16           | 29.1p1<br>29.3p8                            | 29.1p1<br>footnote 343<br><br>29.3p8<br><br>footnote 344 | ed                           | Radioactivity<br><br>Footnotes 343 and 344 from 29.1p1 and 29.3p8 read:<br><br>"Atomic objects are neither active nor radioactive" and<br><br>"Among other implications, atomic variables shall not decay".<br><br>We suggest that these be removed - the first is pretty  | Footnotes 343 and 344 from 29.1p1 and 29.3p8 should be removed.  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   |  |                              | clearly a joke, but it's not obvious that the second doesn't have some technical meaning.   |   |  |
| US 142          | 29.1  | P2 table 141                                 | ed                           | Missing 29.8 Fences   | Add 29.8 Fences   | ACCEPTED   |
| US 143          | [atomics.syn] 29.2                          | before 1                                     | ed                           | There is no free function atomic_compare_exchange_strong for volatile atomic integral types; there is one for non-volatile types.   | Add atomic_compare_exchange_strong for volatile integral types to the synopsis. | ACCEPTED   |
| US 144          | [atomics.syn] 29.2                          | before 1                                     | ed                           | The synopsis lists the macros ATOMIC_INTEGRAL_LOCK_FREE and ATOMIC_ADDRESS_LOCK_FREE; the Lock-free Property requirements don't have ATOMIC_INTEGRAL_LOCK_FREE, but have 8 macros for the various integral types.   | Change 29.2 [atomics.syn] to match 29.4 [atomics.lockfree].                     | ACCEPTED   |
| US 145          | 29.2  |  | ed                           | missing atomic_compare_exchange_strong(volatile ....) and atomic_compare_exchange_strong_explicit( ...//no volatile *?<br><br>> bool atomic_compare_exchange_weak(volatile atomic_itype*, integral*, integral);<br>> bool atomic_compare_exchange_weak(atomic_itype*, integral*, integral);<br>> bool atomic_compare_exchange_strong(atomic_itype*, integral*, integral);<br>> bool atomic_compare_exchange_weak_explicit(volatile atomic_itype*, integral*,<br>> integral, memory_order, memory_order);<br>> bool atomic_compare_exchange_weak_explicit(atomic_itype*, integral*,<br>> integral, memory_order, memory_order);<br>> bool atomic_compare_exchange_strong_explicit(volatile atomic_itype*, integral*,<br>> integral, memory_order, memory_order); | Repair as suggested   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                                 |
| GB 130          | 29.2  |  | Ed                           | The synopsis for the <atomic> header lists the macros ATOMIC_INTEGRAL_LOCK_FREE and ATOMIC_ADDRESS_LOCK_FREE. The ATOMIC_INTEGRAL_LOCK_FREE macro has been replaced with a set of macros for each integral type, as listed in 29.4 | Replace "#define ATOMIC_INTEGRAL_LOCK_FREE unspecified" with<br>#define ATOMIC_CHAR_LOCK_FREE implementation-defined<br>#define ATOMIC_CHAR16_T_LOCK_FREE implementation-defined<br>#define ATOMIC_CHAR32_T_LOCK_FREE implementation-defined<br>#define ATOMIC_WCHAR_T_LOCK_FREE implementation-defined<br>#define ATOMIC_SHORT_LOCK_FREE implementation-defined<br>#define ATOMIC_INT_LOCK_FREE implementation-defined<br>#define ATOMIC_LONG_LOCK_FREE implementation-defined<br>#define ATOMIC_LLONG_LOCK_FREE implementation-defined | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3278</a>                 |
| US 146          | 29.2  | syn 29.4                                     | ed                           | The ATOMIC_..._LOCK_FREE macros have not had changes applied.  | Change to match 29.4/0.  | ACCEPTED   |
| US 147          | 29.2  | syn 29.7                                     | ed                           | The declaration of ATOMIC_VAR_INIT should be referenced to section 29.6 [atomics.types.operations], not 29.7.  | Change it to 29.6.   | ACCEPTED   |
| US 148          | 29.2  | syn 29.7                                     | ed                           | The definition of ATOMIC_VAR_INIT should be 'implementation defined' rather than 'see below'.  | Change it to 'implementation defined'.   | REJECTED<br><br>No, definitely not implementation defined. "See below "is correct. |
| US 149          | 29.2  | syn 29.5.1                                   | ed                           | The synopsis is missing the atomic_init function declarations for the bool, integral and address types.  | Copy them from 29.5.1.   | ACCEPTED   |
| US 150          | 29.2  | syn 29.5.1                                   | ed                           | There are missing function prototypes bool atomic_compare_exchange_strong(volatile atomic_itype*, integral*, integral); and integral atomic_fetch_add(volatile   | Add them.  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)   |
|-----------------|---|--|------------------------------|--|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted            |
|                 |   |  |                              | atomic_itype*, integral);  |   |   |
| US 151          | 29.2  | syn 29.5.1                                   | ed                           | There is a duplicate function declaration of integral atomic_fetch_or(volatile atomic_itype*, integral);   | Remove the volatile qualifier from the second declaration.  | ACCEPTED  |
| US 152          | 29.3  | para 1                                       | ed                           | The table shows no disinct meaning for memory_order_seq_cst.   | Add another bullet: "- memory_order_seq_cst: See below."  | REJECTED<br>It's named in the second and fourth bullet items. |
| GB 131          | 29.3  | 8  | Te                           | See (H) in attachment Appendix 1 - Additional Details  | Request the concurrency working group to determine if changes are needed. Consider changing the use of "sequence" in 29.3 | REJECTED<br>There is no concensus for a change at this time.  |
| CA 21           | 29.3p8<br>1.9p13                            | 29.3p8<br>1.9p13                             | Te                           | Overlapping evaluations are allowed<br>29.3p8 states:<br><br>"An atomic store shall only store a value that has been computed from constants and program input values by a finite sequence of program evaluations, such that each evaluation observes the values of variables as computed by the last prior assignment in the sequence."<br><br>... but 1.9p13 states:<br><br>"If A is not sequenced before B and B is not sequenced before A, then A and B are unsequenced.<br><br>[ Note: The execution of unsequenced evaluations can overlap. -end note ]" | Please clarify.   |   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              | <p>Overlapping executions can make it impossible to construct the sequence described in 29.3p8. We are not sure of the intention here and do not offer a suggestion for change, but note that 29.3p8 is the condition that prevents out-of-thin-air reads.</p> <p>For an example, suppose we have a function invocation f(e1,e2). The evaluations of e1 and e2 can overlap. Suppose that the evaluation of e1 writes y and reads x whereas the evaluation of e2 reads y and writes x, with reads-from edges as below (all this is within a single thread).</p> <pre> e1                e2  Wrlx y--      --Wrlx x       rf\  /rf         X       /  \ Rrlx x&lt;-      -&gt;Rrlx y </pre> <p>This seems like it should be allowed, but there seems to be no way to produce a sequence of evaluations with the property above.</p> <p>In more detail, here the two evaluations, e1 and e2, are being executed as the arguments of a function and are consequently not sequenced-before each other. In practice we'd expect that they could overlap (as allowed by 1.9p13), with the two writes taking effect before the two reads. However, if we have to construct a linear order of evaluations, as in 29.3p8, then the execution above is not permitted. Is that really intended?</p> |   |  |
| US 153          | [atomics.lockfree]                          | before 1                                     | ed                           | The macros are all specified as "implementation-deifned"; they should be "see below", since the required values are spelled out.  | Change the definitions of the macros in 29.4 [atomics.lockfree] from "implementation-defined" to "see below". | REJECTED<br>Implementation defined requires that the |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 | 29.4  |  |                              |   |  | implementation document what it does. Removing this requirement would be a technical change. |
| US 154          | [atomics.lockfree]<br>29.4                  | before 1                                     | te                           | There is no ATOMIC_BOOL_LOCK_FREE macro.  | Add ATOMIC_BOOL_LOCK_FREE to 29.4 [atomics.lockfree] and to 29.2 [atomics.syn]   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3278</a>                           |
| US 155          | 29.4  | para 3                                       | ed                           | The 'via' 'by' word pairing is awkward.   | Replace 'by' with 'via' in 'communication via memory ... and by memory'.   | ACCEPTED   |
| CA 1            | 29.4, 29.6<br>29.7                          | various                                      | te                           | All ATOMIC_... macros should be prefixed with STD_ as in STD_ATOMIC_... to indicate they are STD macros as other standard macros. The rationale that they all seem too long seems weak. | This covers the following macros which we suggest prepending with STD_:<br><br>29.4:<br><i>#define ATOMIC_CHAR_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_CHAR16_T_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_CHAR32_T_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_WCHAR_T_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_SHORT_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_INT_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_LONG_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_LLONG_LOCK_FREE implementation-defined</i><br><i>#define ATOMIC_ADDRESS_LOCK_FREE implementation-defined</i><br><br>29.6:<br><i>#define ATOMIC_VAR_INIT(value) see below</i> | REJECTED<br><br>There is no consensus for a change.  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                 |
|                 |   |  |                              |  | 29.7:<br>#define ATOMIC_FLAG_INIT see below   |  |
| US 156          | [atomics.types .integral]<br>29.5.1         | before 1                                     | ed                           | The list of member functions for atomic_bool has four versions of compare_exchange_weak taking one memory_order argument; the last two should be compare_exchange_strong.  | Change the last two member functions comapare_exchange_weak taking two memory_order arguments to compare_exchange_strong. | ACCEPTED   |
| JP 55           | 29.5.1                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  | Change "table 142" to "Table 142".  | ACCEPTED   |
| GB 132          | 29.5.1                                      |  | Te                           | The atomic_itype types and atomic_address have two overloads of operator= --- one is volatile qualified, and the other is not. atomic_bool only has the volatile qualified version:<br>bool operator=(bool) volatile;<br>On a non-volatile-qualified object this is ambiguous with the deleted copy-assignment operator<br>atomic_bool& operator=(atomic_bool const&) = delete;<br>due to the need for a single standard conversion in each case when assigning a bool to an atomic_bool as in:<br>atomic_bool b;<br>b=true;<br>The conversions are atomic_bool& -> atomic_bool volatile& vs bool -> atomic_bool | Add the "bool operator=(bool);" overload to atomic_bool in 29.5.1   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| US 157          | [atomics.types .integral]<br>29.5.1         | before 1                                     | ed                           | atomic_bool has a volatile assignment operator but not a non-volatile operator The other integral types have both..  | Add a non-volatile assignment operator to atomic_bool.  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|--|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)      | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
| US 158          | 29.5.1   | para 0                                       | ed                           | There is a space before the second & in the declaration <code>atomic_itype&amp; operator=(const atomic_itype &amp;) = delete;</code>  | Remove the space.  | ACCEPTED   |
| US 159          | 29.5.1   |  | Editorial                    | Last 2 should be <code>compare_exchange_strong</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order, memory_order) volatile;</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order, memory_order);</code><br><code>&gt; bool compare_exchange_strong(bool&amp;, bool, memory_order, memory_order) volatile;</code><br><code>&gt; bool compare_exchange_strong(bool&amp;, bool, memory_order, memory_order);</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order = memory_order_seq_cst) volatile;</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order = memory_order_seq_cst);</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order = memory_order_seq_cst) volatile;</code><br><code>&gt; bool compare_exchange_weak(bool&amp;, bool, memory_order = memory_order_seq_cst);</code> | Fix last 2   | ACCEPTED   |
| US 160          | [ <code>atomic.types.integral</code> ]<br>29.5.1 | 1  | te                           | The last sentence of 29.5.1 [ <code>atomics.types.integral</code> ]/1 says "Table 143 shows typedefs to atomic integral classes and the corresponding typedefs." That's nice, but nothing says these are supposed to be part of the implementation, and they are not listed in the synopsis.  | Remove Table 143 and the last sentence of 29.5.1 [ <code>atomics.types.integral</code> ]/1.  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| US 161          | [ <code>atomic.types.address</code> ]<br>29.5.2  | before 1                                     | te                           | <code>atomic_address</code> has <code>operator+=</code> and <code>operator-=</code> , but no <code>operator++</code> or <code>operator--</code> . The template specialization <code>atomic&lt;Ty*&gt;</code> has all of them.   | Add <code>operator++(int) volatile</code> , <code>operator++(int)</code> , <code>operator++() volatile</code> , <code>operator++()</code> , <code>operator--(int) volatile</code> , <code>operator--(int)</code> , <code>operator--() volatile</code> , and <code>operator--()</code> to <code>atomic_address</code> . | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| US 162          | [ <code>atomics.types.address</code> ]<br>29.5.2 |  | te                           | The <code>compare_exchange_weak</code> and <code>compare_exchange_strong</code> member functions that take <code>const void*</code> arguments lead to a silent removal of <code>const</code> , because the load member function and other accessors return the stored value as a <code>void*</code> .   | Remove the <code>const void*</code> overloads of <code>compare_exchange_weak</code> and <code>compare_exchange_strong</code>   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|--|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)                          | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted                 |
| US 163          | [atomics.type.address],<br>[atomics.types.generic]<br>29.5.2, 29.5.3 |  | te                           | Requiring atomic<Ty*> to be derived from atomic_address breaks type safety:<br><br>atomic<double*> ip;<br><br>char ch;<br><br>atomic_store(&ip, &ch);<br><br>*ip.load() = 3.14159;<br><br>The last line overwrites ch with a value of type double   | Remove the requirement that atomic<Ty*> be derived from atomic_address.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| US 164          | [atomics.types.address]<br>29.5.2                                    | before 1                                     | te                           | atomic_address has member functions compare_exchange_weak and compare_exchange_strong that take arguments of type const void*, in addition to the void* versions. If these member functions survive, there should be corresponding free functions.  | Add atomic_compare_exchange_weak and atomic_compare_exchange_strong free functions taking pointers to volatile and non-volatile atomic_address objects and const void* arguments.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |
| JP 56           | 29.5.3   | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "table 142 or table 143" to "Table 142 or Table 143".  | ACCEPTED   |
| GB 133          | 29.5.3   |  | Te                           | The free functions that operate on atomic_address can be used to store a pointer to an unrelated type in an atomic<T*> without a cast. e.g.<br><br>int i;<br>atomic<int*> ai(&i);<br>string s;<br>atomic_store(&ai,&s);   | Overload the atomic_store, atomic_exchange and atomic_compare_exchange_[weak/strong] operations for atomic<T*> to allow storing only pointers to T:<br><br>template<typename T><br>void atomic_store(atomic<T*>&,T*);<br><br>template<typename T> | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3193</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | <pre> void atomic_store(atomic&lt;T*&gt;&amp;,void*) = delete;  template&lt;typename T&gt; void atomic_store_explicit(atomic&lt;T*&gt;&amp;,T*,memory_order);  template&lt;typename T&gt; void atomic_store_explicit(atomic&lt;T*&gt;&amp;,void*,memory_order) = delete;  template&lt;typename T&gt; T* atomic_exchange(atomic&lt;T*&gt;&amp;,T*);  template&lt;typename T&gt; T* atomic_exchange(atomic&lt;T*&gt;&amp;,void*) = delete;  template&lt;typename T&gt; T* atomic_exchange_explicit(atomic&lt;T*&gt;&amp;,T*,memory_order);  template&lt;typename T&gt; T* atomic_exchange_explicit(atomic&lt;T*&gt;&amp;,void*,memory_order) = delete;  template&lt;typename T&gt; T* atomic_compare_exchange_weak(atomic&lt;T*&gt;&amp;,T**,T*);  template&lt;typename T&gt; T* atomic_compare_exchange_weak(atomic&lt;T*&gt;&amp;,void**,void*) = delete; </pre> |  |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB                             | Proposed change by the MB  | Secretariat observations on each comment submitted |
|                 |   |  |                              |  | <pre> template&lt;typename T&gt; T* atomic_compare_exchange_weak_explicit(atomic &lt;T*&gt;&amp;,T**,T*,memory_order);  template&lt;typename T&gt; T* atomic_compare_exchange_weak_explicit(atomic &lt;T*&gt;&amp;,void**,void*,memory_order) = delete;  template&lt;typename T&gt; T* atomic_compare_exchange_strong(atomic&lt;T*&gt;&amp;, T**,T*);  template&lt;typename T&gt; T* atomic_compare_exchange_strong(atomic&lt;T*&gt;&amp;, void**,void*) = delete;  template&lt;typename T&gt; T* atomic_compare_exchange_strong_explicit(atomi c&lt;T*&gt;&amp;,T**,T*,memory_order);  template&lt;typename T&gt; T* atomic_compare_exchange_strong_explicit(atomi c&lt;T*&gt;&amp;,void**,void*,memory_order) = delete;                     </pre> |  |
| US 165          | 29.5.3                                      | Paragraph 23                                 | ed                           | "is the same that same as that of" is not grammatical (and is not clear) |  | ACCEPTED with MODIFICATIONS<br>See LWG 1470        |
| US 166          | 29.6  | para 2                                       | ed                           | The first three bullets seem to be missing terminal punctuation.         | Add semicolons to ends of the bullets.   | ACCEPTED   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

|                    |                           |
|--------------------|---------------------------|
| Date: 7 April 2011 | Document: SC22 WG21 N3289 |
|--------------------|---------------------------|

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)   |
|-----------------|---|--|------------------------------|---|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted  |
| US 167          | 29.6  | para 3                                       | ed                           | The first three bullets seem to be missing terminal punctuation.  | Add semicolons to ends of the bullets.   | REJECTED<br><br>Seems to be a duplicate of US 166 but with the wrong section number.                |
| US 168          | 29.6  | para 4                                       | te                           | The definition of the default constructor needs exposition.   | 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]  | ACCEPTED<br><br>See paper <a href="#">N3196</a>   |
| US 169          | 29.6  | para 5                                       | ed                           | The definition of ATOMIC_VAR_INIT should be 'implementation defined' rather than 'see below'.                                     | Change it to 'implementation defined'.   | REJECTED<br><br>No, we should not require implementors to document the details of how they do this. |
| US 170          | 29.6  | para 6                                       | ed                           | The definition of atomic_init should be grouped with the value constructor.   | Move the atomic_init definition to just after the constructor definition.  | REJECTED<br><br>The current order reflects the order of declarations. No need to change it.         |
| GB 134          | 29.6  | 5  | Ed                           | Some of the declarations of is_lock_free seem to be missing return types.   | Replace:<br>A::is_lock_free() const volatile;<br>A::is_lock_free() const;<br>With:<br>bool A::is_lock_free() const volatile;<br>bool A::is_lock_free() const;  | ACCEPTED  |
| US 171          | 29.6  | para 6                                       | te                           | The atomic_init definition "Non-atomically assigns the value" is not quite correct, as the atomic_init purpose is initialization. | 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]" | ACCEPTED<br><br>See paper <a href="#">N3196</a>   |
| US 172          | 29.6  | para 9, 13, 17, 20                           | ed                           | The order specifications are incomplete because the non_implicit functions do not have such parameters.                           | Add a new sentence: "If the program does not specify an order, it shall be memory_order_seq_cst." Or perhaps: "The non_implicit non-member functions shall affect memory as though they were_implicit with   | REJECTED<br><br>There is no need to make a change. See 29.6 [atomics.types.operations]              |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  | memory_order_seq_cst."   | paragraph 2, bullet 4.   |
| US 173          | 29.6  | para 20                                      | ed                           | Misspelling.   | Replace "operations ate atomic" with "operations are atomic".  | ACCEPTED   |
| US 174          | 29.6  | para 22                                      | ed                           | The first note is about effects, not returns.  | Move this note to just after the Effects paragraph.  | REJECTED<br><br>The note is correct as written. This would require a significant rewrite, because the Returns clause refers to the "result of the comparison" from the Effects clause. |
| US 175          | 29.6  | para 23                                      | ed                           | The first sentence is grammatically incorrect.   | Replace the sentence with two: "The weak compare-and-exchange operations may fail spuriously. That is, it may return false while leaving the contents of memory pointed to by expected the same as it was before the operation." | ACCEPTED with MODIFICATIONS<br>See LWG 1474  |
| CH 23           | 29.6  | p23  | ed                           | The first sentence has non-English syntax.   | Change to "The weak compare-and-exchange operations may fail spuriously, that is, return false while leaving the contents of memory pointed to by expected unchanged."   | ACCEPTED with MODIFICATIONS<br>See LWG 1475  |
| US 176          | 29.6  | para 23                                      | ed                           | Unintended paragraph break.  | Proposal: Remove the paragraph break between "will be in a loop." and "When a compare-and-exchange is in a loop,".   | REJECTED<br><br>Correct as written. Two different subjects, two paragraphs.  |
| US 177          | [atomics.types .operations]<br>29.6         | 23   | ed                           | The first sentence of this paragraph doesn't make sense.   | Figure out what it's supposed to say, and say it.  | ACCEPTED with MODIFICATIONS<br>See LWG 1476  |
| GB 135          | 29.6  | 23   | Ed                           | The first sentence of 29.6p23 was changed by n2992 but now makes no sense: "that is, return false while leaving the contents of memory pointed to by expected before the operation is the same that same as that of the object and | Fix the Remark to say whatever was intended.   | ACCEPTED with MODIFICATIONS<br>See LWG 1477  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)   |
|-----------------|---|--|------------------------------|--|--|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted          |
|                 |   |  |                              | the same as that of expected after the operation." There's a minor editorial difference between n2992 ("is that same as that" vs "is the same that same as that") but neither version makes sense. Also, the remark talks about "object" which should probably be "object or this" to cover the member functions which have no object parameter. |  |   |
| GB 136          | 29.6  |  | Te                           | See (K) in attachment Appendix 1 - Additional Details  | GB requests normative clarification in 29.6p4 that concurrent access constitutes a race, as already done on p6 and p7.   | ACCEPTED with MODIFICATIONS See LWG 1478                    |
| US 178          | 29.7  | para 7                                       | ed                           | The sentence "The order argument shall not be memory_order_acquire nor memory_order_acq_rel." is awkwardly phrased.  | Change the sentence to "The order argument shall be neither memory_order_acquire nor memory_order_acq_rel."  | ACCEPTED with MODIFICATIONS Changed "nor" to "or".          |
| US 179          | 29.8  | para 5, 6                                    | te                           | The fence functions should be extern "C", for C compatibility.   | Add extern "C" to their declarations in 29.8 and in 29.2.  | ACCEPTED with MODIFICATIONS See LWG 1479                    |
| GB 137          | 29.8  | 6  | Te                           | Thread fence not only establish synchronizes with relationships, there are semantics of fences that are expressed not in terms of synchronizes with relationships (for example see 29.3p5). These semantics also need to apply to the use of atomic_signal_fence in a restricted way.  | Change 29.8p6 to "Effects: equivalent to atomic_thread_fence(order), except that the resulting ordering constraints are established only between a thread and a signal handler executed in the same thread." | ACCEPTED  |
| US 180          | 30.1  | para 1                                       | ed                           | The introductory sentence is missing futures.  | Replace "communicate conditions between threads" with "communicate conditions and values between threads".   | ACCEPTED  |
| GB 138          | 30.2  |  | Te                           | The FCD combines the requirements for lockable objects with those for the standard mutex objects. These should be separate. This is LWG issue 1268.  | See attached Appendix 1 - Additional Details   | ACCEPTED with MODIFICATIONS See paper <a href="#">N3197</a> |
| US 181          | 30.2.4                                      | para 2                                       | te                           | The timeout operations are under-specified.  | Define precise semantics for timeout_until and timeout_for. See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS                                 |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted   |
|                 |   |  |                              |  |   | See paper <a href="#">N3191</a>  |
| US 182          | [thread.req.native] 30.2.3                  |  | te                           | native_handle and native_handle_type should be removed. It is implementation-defined whether these names are present in the various thread support classes, and if present, it is implementation-defined what the name native_handle_type refers to. This is exactly what the implementor namespace is for. There is no benefit to programmers from providing a way to portably detect that an implementation provides non-portable extensions. The standard should not reserve these names, with unspecified semantics; if it does, the names will never become portable because implementations will differ on what they mean. | Remove [thread.req.native] 30.2.3 and remove all mentions of native_handle and native_handle_type.                                      | REJECTED<br><br>No consensus to adopt this change for this revision.   |
| DE 23           | 30.3  |  | te                           | Predefined macros usually start and end with two underscores, see 16.8 and FDIS 29124 = WG21 N3060 clause 7. __STDCPP_THREADS should blend in.   | Change the macro name to __STDCPP_THREADS_.   | ACCEPTED   |
| US 183          | 30.3.1                                      |  | te                           | There is no way to join a thread with a timeout.   | Add join_for and join_until. Or decide one should never join a thread with a timeout since pthread_join doesn't have a timeout version. | REJECTED<br><br>Thjere is no consensus to make this change. The poposed changed is regarded as an extension beyond the scope of C++0x. |
| US 184          | 30.3.1.1                                    | para 2                                       | te                           | It is unclear when a thread::id ceases to be meaningful. The sentence "The library may reuse the value of a thread::id of a terminated thread that can no longer be joined." implies that some terminated threads can be joined. It says nothing about detached threads.   | Require a unique thread::id for every thread that is (1) detached and not terminated or (2) has an associated std::thread object.       | REJECTED<br><br>There is no concensus for a change.  |
| JP 97           | 30.3.1.5                                    | 9  | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.   | Change to:<br>Throws: system_error  | ACCEPTED   |
| JP 98           | 30.3.1.5                                    | 14   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead   | Change to:<br>Throws: system_error  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                                     |
|                 |   |  |                              | of std::system_error.  |  |  |
| CH 24           | 30.3.2                                      | p1   | te                           | What would be the value this_thread::get_id() when called from a detached thread?  | Add some text to clarify that get_id() still returns the same value even after detaching.  | REJECTED<br><br>The standard is correct as written. See 30.3.2 [thread.thread.this]/1. |
| CH 25           | 30.3.2                                      | p8 and p11                                   | te                           | Clock related operations are currently not required not to throw. So "Throws: Nothing." is not always true.  | Either require clock related operations not to throw (in 20.10) or change the Throws clauses in 30.3.2. Also possibly add a note that abs_time in the past or negative rel_time is allowed.  | ACCEPTED with MODIFICATIONS<br>See LWG 1487  |
| US 185          | 30.4  |  | te                           | Cooperate with WG14 to improve interoperability between the C++0x and C1x threads APIs. In particular, C1x mutexes should be conveniently usable with a C++0x lock_guard. Performance overheads for this combination should be considered.   | Remove C++0x timed_mutex and timed_recursive_mutex if that facilitates development of more compatible APIs.  | REJECTED<br><br>There is no consensus to adopt this change for this revision.          |
| CH 26           | 30.4  |  | te                           | Specifications of unlock member functions and unlock mutex requirements are inconsistent wrt to exceptions and pre- and postconditions.  | unlock should specify the precondition that the current thread "owns the lock", this will make calls without holding the locks "undefined behavior". unlock in [mutex.requirements] should either be noexcept(true) or be allowed to throw system_error like unique_lock::unlock, or the latter should be nothrow(true) and have the precondition owns==true.<br>Furthermore unique_lock's postcondition is wrong in the case of a recursive mutex where owns might stay true, when it is not the last unlock needed to be called. | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3197</a>                     |
| CH 27           | 30.4.1                                      | p18  | te                           | The mutex requirements force try_lock to be noexcept(true). However, where they are used by the generic algorithms, those relax this requirement and say that try_lock may throw. This means the requirement is too stringent, also a non-throwing try_lock does not allow for a diagnostic such as system_error that lock() will give us. | delete p18, adjust 30.4.4 p1 and p4 accordingly  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3197</a>                     |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                                     |
| JP 99           | 30.4.1                                      | 11   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.  | Change to:<br>Throws: system_error   | ACCEPTED   |
| US 186          | 30.4.1                                      | 14   | te                           | try_lock does not provide a guarantee of forward progress because it is allowed to spuriously fail.   | The standard mutex types must not fail spuriously in try_lock. See Appendix 1 - Additional Details | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3209</a>                     |
| US 187          | 30.4.1                                      | 14   | ed                           | Paragraph mentions compare_exchange, which no longer exists.  | Change "compare_exchange" to "compare_exchange_weak".  | ACCEPTED with MODIFICATIONS<br><br>Changed compare_exchange to "compare and exchange". |
| JP 57           | 30.4.1                                      | 14   | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(29)" to "(Clause 29)".  | ACCEPTED   |
| US 188          | 30.4.1                                      | 20, 21                                       | te                           | Mutex requirements should not be bound to threads   | See Appendix 1 - Additional Details  | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3197</a>                     |
| JP 58           | 30.4.1.1                                    | 3  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.   | Change "(9)" to "(Clause 9)".  | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted                             |
|                 |   |  |                              | However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z.  |   |  |
| US 189          | 30.4.1.1<br>30.4.1.2                        |  | te                           | mutex and recursive_mutex should have an is_locked() member function. is_locked allows a user to test a lock without acquiring it and can be used to implement a lightweight try_try_lock.   | Add a member function:<br><br>bool is_locked() const;<br><br>to std::mutex and std::recursive_mutex. These functions return true if the current thread would not be able to obtain a mutex. These functions do not synchronize with anything (and, thus, can avoid a memory fence). | REJECTED<br><br>There was no consensus to adopt that change for this revision. |
| JP 59           | 30.4.1.2                                    | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(9)" to "(Clause 9)".   | ACCEPTED   |
| JP 60           | 30.4.2.1                                    | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context.<br>However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(9)" to "(Clause 9)".   | ACCEPTED   |
| JP 61           | 30.4.2.2                                    | 2  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and   | Change "(9)" to "(Clause 9)".   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)  | (7)  |
|-----------------|---|--|------------------------------|--|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
|                 |   |  |                              | subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. |  |  |
| JP 70           | 30.4.3.2                                    | 18   | E                            | Constant width font should be used for "std::system_error"s in the paragraph as described in Syntax notation (1.6).  | Change the font for "std::system_error" to constant width type.<br>Throws: <b>system_error</b> when an exception is required   | ACCEPTED   |
| JP 100          | 30.4.3.2.2                                  | 18   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.   | Change to:<br>Throws: system_error   | ACCEPTED   |
| US 190          | 30.4.5.2                                    | para 2, 3                                    | te                           | The term "are serialized" is never defined.  | Remove the sentence with "are serialized" from paragraph 2. Add "Calls to call_once on the same once_flag object shall not introduce data races (17.6.4.8)." to paragraph 3. | ACCEPTED with MODIFICATIONS<br>See LWG 1494                        |
| JP 101          | 30.4.5.2                                    | 4  | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.   | Change to:<br>Throws: system_error   | ACCEPTED   |
| US 191          | 30.5  |  | te                           | The condition variable wait_for returning cv_status is insufficient.   | Return a duration of timeout remaining instead.<br>See Appendix 1 - Additional Details   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3191</a> |
| GB 139          | 30.5  | 7  | Ed                           | The text says "... ownership of the lock as the current thred exits, ...", with "thread" misspelled.   | Replace "thred" with "thread"  | ACCEPTED   |
| GB 140          | 30.5  | 9  | Ed                           | The text says "... waiting threds ..." with "threads" misspelled.  | Replace "threds" with "threads".   | ACCEPTED   |
| CH              | 30.5.1                                      |  | te                           | Requiring wait_until makes it impossible to implement condition_variable correctly using respective objects  | Remove the wait_until functions or make them at least conditionally supported.   | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1)   | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted  |
| 28              |   |  |                              | provided by the operating system (i.e. implementing the native_handle() function) on many platforms (e.g. POSIX, Windows, MacOS X) or using the same object as for the condition variable proposed for C. |   | See paper <a href="#">N3191</a>   |
| JP 102          | 30.5.1                                      | 3  | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.                  | Change to:<br>Throws: system_error  | ACCEPTED  |
| CH 30           | 30.5.1 and 30.5.2                           | p13, last bullet, and corresponding paragraphs in all wait functions                 | te                           | If lock.lock() throws an exception, the postcondition can not be generally achieved.  | Either state that the postcondition might not be achieved, depending on the error condition, or state that terminate() is called in this case.  | ACCEPTED with MODIFICATIONS<br>See LWG 1497   |
| JP 103          | 30.5.1                                      | 15   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.                  | Change to:<br>Throws: system_error  | ACCEPTED  |
| CH 31           | 30.5.1 and 30.5.2                           | p19, third bullet, and corresponding paragraphs in all wait_for/wait_until functions | ed                           | The sentences contain superfluous "or"s.  | Change "The function will unblock when signaled by a call to notify_one() or a call to notify_all(), if abs_time <= Clock::now(), or spuriously." to "The function will unblock when signaled by a call to notify_one(), a call to notify_all(), if abs_time <= Clock::now(), or spuriously." | REJECTED<br><br>This is correct as written. There are three conditions: "when signalled", "if", "spuriously". The first has two parts, separated by "or". |
| JP 104          | 30.5.1                                      | 22   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.                  | Change to:<br>Throws: system_error  | ACCEPTED  |
| US              | 30.5.1                                      | para 26  | ed                           | The identifier cv_status::no_timeout is not in code font.   | Change it to code font.   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)   |
|-----------------|---|--|------------------------------|---|---|---|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted                        |
| 192             |   |  |                              |   |   |   |
| CH 29           | 30.5.1 and 30.5.2                           | p34 and p28, respectively                    | te                           | It is unclear if a spurious wake-up during the loop and re-entering of the blocked state due to a repeated execution of the loop will adjust the timer of the blocking with the respect to the previously specified rel_time value.   | Make it clear (e.g. by a note) that when re-executing the loop the waiting time when blocked will be adjusted with respect to the elapsed time of the previous loop executions. | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3191</a>            |
| US 193          | 30.5.1, 30.5.2                              |  | te                           | Condition variables preclude a wakeup optimization.   | Change condition_variable to allow such optimization. See Appendix 1 - Additional Details   | REJECTED<br>There is no consensus to adopt this change for this revision. |
| JP 62           | 30.5.1                                      | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(9)" to "(Clause 9)".   | ACCEPTED  |
| CH 32           | 30.5.2                                      |  | te                           | Given that the lock type can be something the underlying doesn't know 'native_handle()' is probably unimplementable on essentially all platforms.   | Consider the removal of 'native_handle()'.  | ACCEPTED  |
| JP 105          | 30.5.2                                      | 12   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.  | Change to:<br>Throws: system_error  | ACCEPTED  |
| JP 106          | 30.5.2                                      | 18   | E                            | In Throw clause, both "Throws: system_error" and "Throws: std::system_error" are used. They should be in a unified way, and we propose to use system_error instead of std::system_error.  | Change to:<br>Throws: system_error  | ACCEPTED  |
| CH              | 30.5.2                                      | before p25                                   | ed                           | Template function wait_until is missing class Clock template parameter.   | Change "template <class Lock, class Duration, class Predicate>" to "template <class Lock, class   | ACCEPTED  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.



ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted             |
| 33              |   |  |                              |   | Clock, class Duration, class Predicate>".   |  |
| JP 96           | 30.5.3                                      | 2  | E                            | Inconsistency between 30.4 paragraph 1 and 30.4.3 paragraph 2.<br>In 30.4 paragraph 1:<br>namespace std {<br>...<br>constexpr defer_lock_t defer_lock {};<br>constexpr try_to_lock_t try_to_lock {};<br>constexpr adopt_lock_t adopt_lock {};<br>}<br>In 30.4.3 paragraph 2:<br>namespace std {<br>...<br>extern const defer_lock_t defer_lock {};<br>extern const try_to_lock_t try_to_lock {};<br>extern const adopt_lock_t adopt_lock {};<br>}<br>The writer seems to have forgotten to rewrite latter cases, so 30.4.3 paragraph 2 should be modified as this proposal. | Change "extern const" to "constexpr".<br>namespace std {<br>...<br>constexpr defer_lock_t defer_lock {};<br>constexpr try_to_lock_t try_to_lock {};<br>constexpr adopt_lock_t adopt_lock {};<br>} | ACCEPTED   |
| US 194          | 30.6  |  | te                           | The specification for managing associated asynchronous state is confusing, sometimes omitted, and redundantly specified.  | Define terms-of-art for releasing, making ready, and abandoning an associated asynchronous state. Use those terms where appropriate. See Appendix 1 - Additional Details                          | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3192</a> |
| CH 34           | 30.6.1                                      | p1   | ed                           | The paragraph is misleading and incorrect wrt to the current specification, since an async call with launch::sync will execute in the same thread.  | Change the paragraph to '30.6 describes components that a C++ program can use to retrieve the result (value or exception) of a function that has run in a (potentially different) thread.'        | ACCEPTED   |
| US 195          | 30.6.4                                      | para 8                                       | te                           | The intent and meaning of the paragraph is not apparent.  |   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3278</a> |
| CH              | 30.6.4ff                                    |  | ed/te                        | The term "associated asynchronous state" is long, ugly  | Change all occurrences of "associated   | ACCEPTED with  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2  | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|--|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)                                      | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted             |
| 35              |  |  |                              | and misleading terminology. When introduced we agreed upon that we should come up with a better name. Here it is: "liaison state". Since the state is hidden and provides synchronization of a future with its corresponding promise, we believe "liaison state" is a much better and shorter name (liaison ~ (typically hidden) relationship) | asynchronous state" to "liaison state".   | MODIFICATIONS<br>See LWG 1503                                  |
| US 196          | 30.6.5   | para 21, 25                                  | te                           | The term "are serialized" is not defined.  | Replace "are serialized" with "shall not introduce a data race (17.6.4.8)".   | ACCEPTED with MODIFICATIONS<br>See LWG 1504                    |
| US 197          | 30.6.5   | para 21, 25                                  | te                           | There is no defined synchronization between promise::set_value and future::get.  | Replace "[Note: and they synchronize and serialize with other functions through the referred associated asynchronous state. --end note]" with the normative "They synchronize with (1.10) any operation on a future object with the same associated asynchronous state marked ready." | ACCEPTED with MODIFICATIONS<br>See LWG 1505, and 1507          |
| US 198          | 30.6.5   | 22   | te                           | promise::set_exception can be called with a null pointer, but none of the descriptions of the get() functions for the three types of futures say what happens for this case.   | Add the following sentence to the end of 30.6.5/22: The behavior of a program that calls set_exception with a null pointer is undefined.  | ACCEPTED   |
| US 199          | [futures.promise]<br>30.6.5  | 26ff, 29ff                                   | te                           | promise::XXX_at_thread_exit functions have no synchronization requirements. Specifying synchronization for these member functions requires coordinating with the words in 30.6.5/21 and 25, which give synchronization requirements for promise::set_value and promise::set_exception.   | Change 30.6.5/21 to mention set_value_at_thread_exit and set_exception_at_thread_exit; with this text, replace 30.6.5/25 and add two new paragraphs, after 30.6.5/28 and 30.6.5/31.   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3278</a> |
| US 200          | 30.6.6   | para 26                                      | ed                           | The paragraph is missing the "Returns:" label.   | Add the label.  | ACCEPTED   |
| US 201          | [futures.unique_future],<br>[futures.shared_future],<br>[futures.atomic_future], |  | te                           | packaged_task provides operator bool() to check whether an object has an associated asynchronous state. The various future types provide a member function valid() that does the same thing. The names of these members should be the same.  | Replaced the name packaged_task::operator bool() with packaged_task::valid() in the synopsis (30.6.10 [futures.task]/2) and the member function specification (before 30.6.10.1 [futures.task.members]/15).   | ACCEPTED with MODIFICATIONS<br>See paper <a href="#">N3194</a> |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)   | (7)  |
|-----------------|---|--|------------------------------|---|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1)             | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB   | Secretariat observations on each comment submitted                 |
|                 | [futures.task]<br>30.6.6<br>30.6.7<br>30.6.8<br>30.6.10 |  |                              |   |   |  |
| US 202          | [futures.atomic_future]<br>30.6.8                       | 18   | te                           | The note in this paragraph says "unlike future, calling get more than once on the same atomic_future object is well defined and produces the result again." There is nothing in future that says anything negative about calling get more than once.  | Remove this note, or add words to the requirements for future that reflect what this note says.                     | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3194</a> |
| US 203          | [futures.atomic_future]<br>30.6.8                       |  | te                           | Both future and shared_future specify that calling most member functions on an object for which valid() == false produces undefined behavior. There is no such statement for atomic_future.   | Add a new paragraph after 30.6.8 [futures.atomic_future]/2 with the same words as 30.6.7 [futures.shared_future]/3. | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3194</a> |
| US 204          | 30.6.8  | Paragraph 7-8                                | te                           | According to the definition of atomic_future, all members of atomic_future are synchronizing except constructors. However, it would probably be appropriate for a move constructor to be synchronizing on the source object. If not, the postconditions on paragraphs 7-8, might not be satisfied. This may be applicable if a collection of futures are being doled out to a set of threads that process their value.  | Make the move constructor for atomic future lock the source   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3194</a> |
| US 205          | [futures.async]<br>30.6.9                               | 3  | te                           | The third sentence says "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." The next sentence says "If the invocation is not deferred, the join() on the created thread..." Blocking until a thread completes is not necessarily a join. | Decide whether the requirement is to block until finished or to call join, and rewrite to match.                    | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3194</a> |
| CH 36           | 30.6.9 and 30.6.1                                       | <future> synopsis                            | te                           | Providing only three different possible values for the enum launch and saying that launch::any means either   | Change in 30.6.1 'enum class launch' to allow further implementation defined values and provide                     | ACCEPTED with MODIFICATIONS  |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011 Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5  | (6)   | (7)  |
|-----------------|---|--|------------------------------|--|---|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB   | Proposed change by the MB   | Secretariat observations on each comment submitted |
|                 |   | and p3, respectively                         |                              | launch::sync or launch::async is very restricting. This hinders future implementors to provide clever infrastructures that can simply be used by a call to async(launch::any,...). Also there is no hook for an implementation to provide additional alternatives to launch enumeration and no useful means to combine those (i.e. interpret them like flags). We believe something like async(launch::sync   launch::async, ...) should be allowed and can become especially useful if one could say also something like async(launch::any & ~launch::sync, ...) respectively. This flexibility might limit the features usable in the function called through async(), but it will allow a path to effortless profit from improved hardware/software without complicating the programming model when just using async(launch::any,...) | the following bit-operators on the launch values (operator , operator&, operator~ delivering a launch value).<br>Note: a possible implementation might use an unsigned value to represent the launch enums, but we shouldn't limit the standard to just 32 or 64 available bits in that case and also should keep the launch enums in their own enum namespace. Change [future.async] p3 according to the changes to enum launch. change --launch::any to "the implementation may choose any of the policies it provides." Note: this can mean that an implementation may restrict the called function to take all required information by copy in case it will be called in a different address space, or even, on a different processor type. To ensure that a call is either performed like launch::async or launch::sync describe one should call async(launch::sync launch::async,...) | See paper <a href="#">N3188</a>                    |
| JP 107          | 30.6.9                                      | 3  | E                            | Typo, "." should be ">".<br>decay_copy(std::forward<Args>(args))...  | Correct typo.<br>decay_copy(std::forward<Args>(args))...  | ACCEPTED   |
| JP 108          | 30.6.9                                      | 3  | E                            | <Arg> should be <Args>.<br>launch::sync — Stores decay_copy(std::forward<F>(f)) and<br>decay_copy(std::forward<Args>(args))...   | Change to:<br>launch::sync — Stores<br>decay_copy(std::forward<F>(f)) and<br>decay_copy(std::forward<Args>(args))...  | ACCEPTED   |
| US 206          | [futures.task.members]<br>30.6.10.1         | 27, 28                                       | ed                           | The text of paragraph 27 says that reset() moves the function object, but the text of paragraph 28 talks about exceptions thrown by the copy constructor.  | Change "copy constructor" to "move constructor" in 30.6.10.1/28, bullet 2.  | ACCEPTED   |
| US 207          | [futures.task.members]<br>30.6.10.1         | 1-8  | te                           | The constructor that takes R*(ArgTypes...) is not needed; the constructor that takes a callable type works for this argument type. More generally, the constructors  | Review the constructors for packaged_task and provide the same ones as function, except where inappropriate.  | ACCEPTED with MODIFICATIONS<br>See LWG 1514        |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

ISO/IEC FCD 14882 Ballot Comments and Responses

Date: 7 April 2011

Document: SC22 WG21 N3289

| 1               | 2   | (3)  | 4                            | 5   | (6)  | (7)  |
|-----------------|---|--|------------------------------|---|--|--|
| MB <sup>1</sup> | Clause No./ Subclause No./ Annex (e.g. 3.1) | Paragraph/ Figure/Table/ Note (e.g. Table 1) | Type of comment <sup>2</sup> | Comment (justification for change) by the MB  | Proposed change by the MB  | Secretariat observations on each comment submitted                 |
|                 |   |  |                              | for packaged_task should parallel those for function.   |  |  |
| US 208          | [futures.task.members]<br>30.6.10.1         | 24-26  | te                           | packaged_task::make_ready_at_thread_exit has no synchronization requirements.   | Figure out what the synchronization requirements should be and write them.   | ACCEPTED with MODIFICATIONS<br><br>See paper <a href="#">N3278</a> |
| GB 141          | Appendix A [gram] paragraph 1               |  | Ed                           | The links for disambiguation rules go to 6.8, 7.1 and 10.2. Section 8.2 covers ambiguity in declarators, so should be added   | Added a link to 8.2 [dcl.ambig.res] to Appendix A p1   | ACCEPTED   |
| JP 63           | A.1   | 1  | E                            | Representations of reference link are not unified. Most reference links to clause (table) number, say X, are in the form "Clause X" ("Table X") capitalized, and subsection Y.Y.Y is referenced with its number only in the form "Y.Y.Y". Whether they are parenthesized or not depends on the context. However there are some notations "(Z)" consisting of only a number Z in parentheses to confer Clause or Table number Z. | Change "(clause 9)" to "(Clause 9)".<br>Change "(clause 14)" to "(Clause 14)".   | ACCEPTED   |
| FI 6            | D.2 [depr.static]                           | Paragraph 1                                  | te                           | The use of static in namespace scope should not be deprecated. Anonymous namespaces are not a sufficient replacement for the functionality.   | Strike [depr.static] completely.   | ACCEPTED   |
| GB 142          | D10   |  | Ge                           | auto_ptr does not appear in the <memory> synopsis and [depr.auto_ptr] doesn't say which header declares it. Conversely, the deprecated binders bind1st etc. are in the <functional> synopsis, this is inconsistent  | Either auto_ptr should be declared in the <memory> synopsis, or the deprecated binders should be removed from the <functional> synopsis and appendix D should say which header declares the binders and auto_ptr | ACCEPTED   |
| JP 109          | Annex B                                     |  | E                            | Although implementation limits for at_quick_exit() is mentioned in 18.5 paragraph 5, it is not on the list of Implementation quantities.  | Add the following line.<br>— Functions registered by at_quick_exit()[32].  | ACCEPTED   |
| JP 110          | Index                                       |  | E                            | "local scope" has been renamed to "block scope", but the reference to "local scope" still remains in Index.<br><br>block scope; see local scope, 37   | Change to:<br>local scope; see block scope, 37   | ACCEPTED   |

1 MB = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 Type of comment: ge = general te = technical ed = editorial

NOTE Columns 1, 2, 4, 5 are compulsory.

**ISO/IEC FCD 14882 Ballot Comments and Responses**

|                    |                           |
|--------------------|---------------------------|
| Date: 7 April 2011 | Document: SC22 WG21 N3289 |
|--------------------|---------------------------|

| 1                     | 2  | (3)   | 4  | 5   | (6)                              | (7)   |
|-----------------------|--|---|--|---|----------------------------------|---|
| <b>MB<sup>1</sup></b> | <b>Clause No./<br/>Subclause No./<br/>Annex<br/>(e.g. 3.1)</b> | <b>Paragraph/<br/>Figure/Table/<br/>Note<br/>(e.g. Table 1)</b> | <b>Type<br/>of<br/>com-<br/>ment<sup>2</sup></b> | <b>Comment (justification for change) by the MB</b> | <b>Proposed change by the MB</b> | <b>Secretariat observations<br/>on each comment submitted</b> |
|                       |  |   |  | "local scope" should refer to "block scope".        |                                  |   |

1 **MB** = Member body (enter the ISO 3166 two-letter country code, e.g. CN for China; comments from the ISO/CS editing unit are identified by \*\*)

2 **Type of comment:** **ge** = general **te** = technical **ed** = editorial

**NOTE** Columns 1, 2, 4, 5 are compulsory.

## **APPENDIX 1**

### **ADDITIONAL DETAILS**

#### **ISO/IEC FCD 14882, C++0X, National Body Comments**

**FCD 14882**

**ADDITIONAL DETAILS TO US BALLOT COMMENTS**



FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 23:**

**3.4.5: Global class templates should not hide member templates.**

**Technical details:**

[basic.lookup.classref] p1 says:

In a class member access expression (5.2.5), if the . or -> token is immediately followed by an identifier followed by a <, the identifier must be looked up to determine whether the < is the beginning of a template argument list (14.3) or a less-than operator. The identifier is first looked up in the class of the object expression. If the identifier is not found, it is then looked up in the context of the entire postfix-expression and shall name a class template. If the lookup in the class of the object expression finds a template, the name is also looked up in the context of the entire postfix-expression and

- if the name is not found, the name found in the class of the object expression is used, otherwise
- if the name is found in the context of the entire postfix-expression and does not name a class template, the name found in the class of the object expression is used, otherwise
- if the name found is a class template, it shall refer to the same entity as the one found in the class of the object expression, otherwise the program is ill-formed.

This means that the following program is ill-formed:

```
#include <set>
using std::set;
struct X {
    template <typename T> void set(const T& value);
};
void foo() {
    X x;
    x.set<double>(3.2);
}
```

That's confusing and unnecessary. The compiler has already done the lookup in X's scope, and the obviously-correct resolution is that one, not the identifier from the postfix-expression's scope. Issue 305 fixed a similar issue for destructor names, but missed member functions.

**Proposed resolution:**

Strike the end of paragraph 1 starting with "If the lookup in the class of the object expression finds a template," and including all three bullets.

**US 26****Clause 3.7.4**

# C++ Sized Deallocation

ISO/IEC JTC1 SC22 WG21 - 2010-05-19 - National Body Comment by Google

Lawrence Crowl, crowl@google.com, Lawrence@Crowl.org

[Problem](#)

[Solution](#)

[Wording](#)

[3.7.4 Dynamic storage duration \[basic.stc.dynamic\]](#)

[3.7.4.2 Deallocation functions \[basic.stc.dynamic.deallocation\]](#)

[5.3.5 Delete \[expr.delete\]](#)

[12.5 Free store \[class.free\]](#)

[17.6.3.6 Replacement functions \[replacement.functions\]](#)

[18.6 Dynamic memory management \[support.dynamic\]](#)

[18.6.1.1 Single-object forms \[new.delete.single\]](#)

[18.6.1.2 Array forms \[new.delete.array\]](#)

## Problem

Within the Final Committee Draft, programmers may define a static member function `operator delete` that takes a size parameter indicating the size of the object to be deleted. The equivalent global `operator delete` is not available. This omission has unfortunate performance consequences.

Modern memory allocators often allocate in size categories, and, for space efficiency reasons, do not store the size of the object near the object. Deallocation then requires searching for the size category store that contains the object. This search can be expensive, particularly as the search data structures are often not in memory caches.

## Solution

Permit implementations and programmers to define sized versions of the global `operator delete`. The compiler shall the sized version in preference to the unsized version when the sized version is available.

There are two potential problems with this solution.

- When deleting an incomplete type, there is not size available. In this case, the unsized version must be used. This observation implies that calls to one version must be effectively equivalent to calls to the other version. Excepting the specific deallocation function called, we believe that any programs that would change behavior already have undefined behavior within the standard.
- Existing programs that redefine the global unsized version do not also define the sized version. When an implementation introduces a sized version, the replacement would be incomplete. In this case, the only viable option seems to be for the implementation to emit a diagnostic when the programmer provides an unsized replacement but does not provide a sized replacement. The workaround is to define a sized version that simply calls the unsized version.

Note, however, that the converse case is not a problem. The programmer may define a both an unsized and a sized version even when the underlying implementation only provides a sized version. The reason is that the two versions must be functionally equivalent.

As a consequence of the second problem, we expect implementations to be somewhat conservative in the introduction of the pre-defined sized versions. On the other hand, programmers may aggressively define the sized versions.

## Wording

The proposed wording changes are relative to the Final Committee Draft, [N3092](#).

### 3.7.4 Dynamic storage duration [basic.stc.dynamic]

Edit within paragraph 2 as follows.

.... The following allocation and deallocation functions (18.6) are implicitly declared in global scope in each translation unit of a program.

```
void* operator new(std::size_t)
throw(std::bad_alloc);
void* operator new[](std::size_t)
throw(std::bad_alloc);
void operator delete(void*) throw();
void operator delete[](void*) throw();
```

Furthermore, the implementation may implicitly declare the following functions in global scope in each translation unit of a program.

```
void operator delete(void*, std::size_t)
throw();
void operator delete[](void*, std::size_t)
throw();
```

If the implementation provides these functions, and if a translation unit provides a definition for an unsized version but not for a sized version, the program is ill-formed.

These implicit declarations introduce only the function names `operator new`, `operator new[]`, `operator delete`, `operator delete[]`. ....

### 3.7.4.2 Deallocation functions [basic.stc.dynamic.deallocation]

Edit paragraph 2 as follows.

Each deallocation function shall return `void` and its first parameter shall be `void*`. A deallocation function can have more than one parameter. If there is a declaration of global `operator delete` with exactly two parameters, the second of which has type `std::size_t` (18.2), then that function is a usual (non-placement) deallocation function. Similarly, if there is a declaration of global `operator delete[]` with exactly two parameters, the second of which has type `std::size_t`, then this function is a usual deallocation function. If a class `T` has a member deallocation function named `operator delete` with exactly one parameter, then that function is a usual (~~non-placement~~) deallocation function. If class `T` does not declare such an `operator delete` but does declare a member deallocation function named `operator delete` with exactly two parameters, the second of which has type `std::size_t` (18.2), then this function is a usual deallocation function. Similarly, if a class `T` has a member deallocation function named `operator delete[]` with exactly one parameter, then that function is a usual (non-placement) deallocation function. If class `T` does not declare such an `operator delete[]` but does declare a member deallocation function named `operator delete[]` with exactly two parameters, the second of which has type `std::size_t`, then this function is a usual deallocation function. A deallocation function can be an instance of a function template. Neither the first parameter nor the return type shall depend on a template parameter. [*Note*: that is, a deallocation function template shall have a first parameter of type `void*` and a return type of `void` (as specified above). —*end note*] A deallocation function template shall have two or more function parameters. A template instance is never a usual deallocation function, regardless of its signature.

### 5.3.5 Delete [expr.delete]

Paragraph 1 remains unchanged, though note the restrictions on the delete operand.

.... The operand shall have a pointer to object type, or a class type having a single non-explicit conversion function (12.3.2) to a pointer to object type. The result has type `void`. [*Footnote*: This implies that an object cannot be

deleted using a pointer of type `void*` because `void` is not an object type.  
—*end footnote*] ....

paragraph 2 remains unchanged, though note the restriction on inheritance with respect to the `delete` operand.

.... If it is not a null pointer value, in the first alternative (*delete object*), the value of the operand of `delete` shall be a pointer to a non-array object or a pointer to a subobject (1.8) representing a base class of such an object (Clause 10). If not, the behavior is undefined. In the second alternative (*delete array*), the value of the operand of `delete` shall be the pointer value which resulted from a previous array `new`-expression. [*Footnote*: For non-zero-length arrays, this is the same as a pointer to the first element of the array created by that `new`-expression. Zero-length arrays do not have a first element. —*end footnote*] If not, the behavior is undefined. [*Note*: this means that the syntax of the *delete-expression* must match the type of the object allocated by `new`, not the syntax of the *new-expression*. —*end note*] ....

Paragraph 3 remains unchanged, though note the further restriction on inheritance.

In the first alternative (*delete object*), if the static type of the object to be deleted is different from its dynamic type, the static type shall be a base class of the dynamic type of the object to be deleted and the static type shall have a virtual destructor or the behavior is undefined. In the second alternative (*delete array*) if the dynamic type of the object to be deleted differs from its static type, the behavior is undefined.

Paragraph 5 remains unchanged.

If the object being deleted has incomplete class type at the point of deletion and the complete class has a non-trivial destructor or a deallocation function, the behavior is undefined.

Remove paragraph 9 as follows.

When the keyword `delete` in a *delete-expression* is preceded by the unary `::` operator, the global deallocation function is used to deallocate the storage.

Add a new paragraph in place of paragraph 9 as follows.

If a *delete-expression* begins with a unary `::` operator, the deallocation function's name is looked up in global scope. Otherwise, the lookup considers class-specific deallocations (12.5 [class.free]). If no class-specific deallocation is found, the deallocation function's name is looked

up in global scope. If the lookup selects a placement deallocation function, the program is ill-formed.

Add a new paragraph as follows.

If deallocation function lookup finds both a usual deallocation function with one parameter and a usual deallocation function with two parameters, and then if the object being deleted has incomplete class type, the selected deallocation function shall be the one with two parameters. Otherwise, the selected deallocation function shall be the function with one parameter.

Add a new paragraph as follows. This paragraph is identical to the existing 12.5/5.

When a *delete-expression* is executed, the selected deallocation function shall be called with the address of the block of storage to be reclaimed as its first argument and (if the two-parameter style is used) the size of the block as its second argument. [*Footnote*: If the static type of the object to be deleted is different from the dynamic type and the destructor is not virtual the size might be incorrect, but that case is already undefined, as stated above. —*end footnote*]

## 12.5 Free store [class.free]

Edit paragraph 4 as follows.

Class-specific deallocation function lookup is a part of general deallocation function lookup (5.3.5 [expr.delete]) and occurs as follows. If a *delete-expression* begins with a unary `::` operator, the deallocation function's name is looked up in global scope. Otherwise, if the *delete-expression* is used to deallocate a class object whose static type has a virtual destructor, the deallocation function is the one selected at the point of definition of the dynamic type's virtual destructor (12.4). [*Footnote*: A similar provision is not needed for the array version of `operator delete` because 5.3.5 requires that in this situation, the static type of the object to be deleted be the same as its dynamic type. —*end footnote*] Otherwise, if the *delete-expression* is used to deallocate an object of class `T` or array thereof, the static and dynamic types of the object shall be identical and the deallocation function's name is looked up in the scope of `T`. If this lookup fails to find the name, ~~the name is looked up in the global scope.~~ the class-specific deallocation function lookup has failed and general deallocation function lookup (5.3.5 [expr.delete]) continues. If the result of the lookup is ambiguous or inaccessible, or if the lookup selects a placement deallocation function, the program is ill-formed.

Remove paragraph 5 as follows. This paragraph moves to 5.3.5/9++.

~~When a *delete-expression* is executed, the selected deallocation function shall be called with the address of the block of storage to be reclaimed as its first argument and (if the two-parameter style is used) the size of the block as its second argument. [*Footnote*: If the static type of the object to be deleted is different from the dynamic type and the destructor is not virtual the size might be incorrect, but that case is already undefined; see 5.3.5.—*end footnote*]~~

### 17.6.3.6 Replacement functions [replacement.functions]

Edit paragraph 2 as follows.

A C++ program may provide the definition for any of eight dynamic memory allocation function signatures declared in header `<new>` (3.7.4, ~~Clause 18~~ 18.4 [support.dynamic]):

- `operator new(std::size_t)`
- `operator new(std::size_t, const std::nothrow_t&)`
- `operator new[](std::size_t)`
- `operator new[](std::size_t, const std::nothrow_t&)`
- `operator delete(void*)`
- `operator delete(void*, const std::nothrow_t&)`
- `operator delete[](void*)`
- `operator delete[](void*, const std::nothrow_t&)`

Furthermore, a C++ program may provide the definition for any of the four dynamic memory deallocation function signatures that implementations may choose to declare in header `<new>`:

- `operator delete(void*, std::size_t)`
- `operator delete(void*, std::size_t, const std::nothrow_t&)`
- `operator delete[](void*, std::size_t)`
- `operator delete[](void*, std::size_t, const std::nothrow_t&)`

## 18.6 Dynamic memory management [support.dynamic]

At the end of the synopsis add the following.

The implementation may, but need not, provide the following additional group of functions.

```
operator delete(void* ptr, std::size_t size)  
throw();  
operator delete(void* ptr, std::size_t size,  
const std::nothrow_t&) throw();
```

```
operator delete[](void* ptr, std::size_t size)  
throw();  
operator delete[](void* ptr, std::size_t size,  
const std::nothrow_t&) throw();
```

### 18.6.1.1 Single-object forms [new.delete.single]

At the end of the section, add a new paragraph as follows.

```
operator delete(void* ptr, std::size_t size) throw();  
operator delete(void* ptr, std::size_t size, const  
std::nothrow_t&) throw();
```

Add a new paragraph as follows.

These functions behave as their corresponding version  
without the std::size\_t size parameter, with the additional  
constraints:

Requires: size shall equal that used to  
allocate ptr.

Required behavior: Calls to the sized and  
unsized versions shall be interchangeable.

### 18.6.1.2 Array forms [new.delete.array]

At the end of the section, add a new paragraph as follows.

```
operator delete[](void* ptr, std::size_t size) throw();  
operator delete[](void* ptr, std::size_t size, const  
std::nothrow_t&) throw();
```

Add a new paragraph as follows.

These functions behave as their corresponding version  
without the std::size\_t size parameter, with the additional  
constraints:

Requires: size shall equal that used to  
allocate ptr.

Required behavior: Calls to the sized and  
unsized versions shall be interchangeable.



FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 28**  
**Clause 4.4**

**Comment 4.4/3: A const member function pointer could safely be applied to a non-const object without violating const correctness.**

**Technical details:**

While a cv-qualified member function can be called on a less-cv-qualified object-expression (see [class.this] 9.3.2/4), a cv-qualified member function cannot be assigned to less-cv-qualified member function pointer even though it is no more dangerous to const correctness than the rule in 9.3.2/4. This code sample illustrates:

```
struct X {
  void func1() const;
  void func2();
};
void f() {
  void (X::*fptr1)() = &X::func1;      // ill-formed, should be well-formed
  void (X::*fptr2)() = &X::func2;      // well-formed
  void (X::*fptr3)() const = &X::func1; // well-formed
  void (X::*fptr4)() const = &X::func2; // ill-formed, must remain so
}
```

**Proposed resolution:**

Informally, we want to add an implicit conversion for pointer-to-member-function types from "void (T::\*)() const" to "void (T::\*)()". This is slightly tricky because the cv-qualifier-seq on a function declaration isn't actually a cv-qualified function type (see 8.3.5).

One possible wording change: strike 9.3.2/4, and add to 4.4/3 "A prvalue of type 'pointer to cv-qualified member function' can be converted to a prvalue of type 'pointer to less-cv-qualified member function'."

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 29**

Jason Merrill  
2010-03-09  
Revision 3

**CORE ISSUE 974. Default arguments for lambdas**

**Section:** 5.1.2 [expr.prim.lambda]    **Status:** open    **Submitter:** Jason Merrill    **Date:** 4 September, 2009  
**Priority:** 2

(From [15012](#) and [15148](#).)

There does not appear to be any technical difficulty that would require the restriction in 5.1.2 [expr.prim.lambda] paragraph 5 against default arguments in *lambda-expressions*.

**Suggested Resolution (March, 2010)**

Strike from 5.1.2/5:

~~Default arguments (8.3.6) shall not be specified in the parameter declaration clause of a lambda-declarator~~

## US 30

### 975. Restrictions on return type deduction for lambdas

**Section:** 5.1.2 [expr.prim.lambda]   **Status:** open   **Submitter:** Jason Merrill   **Date:** 4 September, 2009   **Priority:** 2

(From messages [15012](#), [15148](#), [15152](#), and [15170](#).)

There does not appear to be any technical difficulty that would require the current restriction that the return type of a lambda can be deduced only if the body of the lambda consists of a single return statement. In particular, multiple return statements could be permitted if they all return the same type.

#### Drafting note

It is unfortunate that there is no way of writing directly the type deduced from an expression by auto or lambda return type deduction; decltype has significantly different results.

#### Suggested Resolution (March, 2010)

Change 5.1.2/4 from:

..... If a *lambda-expression* does not include a *trailing-return-type*, it is as if the *trailing-return-type* denotes the following type:

- if the *compound-statement* ~~if of the form~~

```
{ return attribute-specifieropt expression ; }
```

the type of the returned expression after lvalue-to-rvalue conversion (4.1), array-to-pointer conversion (4.2), and function-to-pointer conversion (4.3);

- otherwise, ~~void~~.

[ *Example:*

```
auto x1 = [](int i){ return i; }; // OK: return type is int
auto x2 = []{ return { 1, 2 }; }; // error: the return type is void
(a
                                     // braced-init-list is not an
expression)
-- end example ]
```

to:

..... If a *lambda-expression* does not include a *trailing-return-type*, it is as if the *trailing-return-type* denotes the following type:

- If there are no return statements in the compound-statement or all return statements return void, void;
- otherwise, if all return statements are of the form return expression ; and for all return statements the type of the returned expression after lvalue-to-rvalue conversion (4.1), array-to-pointer conversion (4.2), and function-to-pointer conversion (4.3) is equivalent, that type;
- otherwise, the program is ill-formed.

[ Example:

```
auto x1 = [](int i){ return i; }; // OK: return type is int
auto x2 = []{ return { 1, 2 }; }; // error (a braced-init-list is
not an expression)
```

```
template <class T> void f () {
    [](T t, bool b){
        if (b)
            return t.fn();
        else
            return t.fn();
    }; // OK: return type is type of
t.fn()
    [](T t, bool b){
        if (b)
            return t.fn1();
        else
            return t.fn2();
    }; // error: the type of t.fn1() is
not equivalent to the type of t.fn2() (14.6.6.1)
}
```

-- end example ]

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 34**

**Comment 5.3.4, 5.3.5: Allocation functions are missing happens-before requirements and guarantees.**

**Technical details:**

When the same unit of storage is allocated and deallocated repeatedly, operations on it can't be allowed to race between the allocator and the user program. But I don't see any mention of happens-before in the descriptions of allocation and deallocation functions.

**Proposed resolution (not wording yet):**

- The call to an allocation function returning a pointer P must happen-before the matching deallocation call with P as a parameter. Otherwise the behavior is undefined. I don't know whether receiving P with `memory_order_consume` fits this requirement. `memory_order_relaxed` does not.
- If some memory is passed to a deallocation function, the implementation must ensure that the deallocation call happens-before any allocation call that returns the same memory address.

US 49

Jason Merrill  
2010-05-25  
Revision 1

# Direct Binding correction

## The Problem

The FCD rules do not specify direct binding for this example:

```
int i;  
int main()  
{  
    int&& ir = static_cast<int&&>(i);  
    ir = 42;  
    return (i != 42);  
}
```

We ought to do direct binding for reference-compatible xvalues. It also seems that the array rvalue case ought to be folded in with the class rvalue case, since the only way to get an array rvalue is to refer to an array member of a class rvalue. I considered introducing the term "crvalue" for that subset of prvalues, but couldn't find any other places to use it. I also considered extending the direct binding treatment to scalar members of a class rvalue, but that seems like a larger change than necessary.

## Proposed Wording

Change 8.5.3 [dcl.init.ref] paragraph 5 as follows:

...

- Otherwise, if the initializer expression is an xvalue, class prvalue or array prvalue, and "cv1 T1" is reference-compatible with "cv2 T2", the reference is bound to the object represented by the rvalue (see 3.10) or to a subobject within that object.
- Otherwise, if T2 is a class type and
  - the initializer expression is an rvalue and "cv1 T1" is reference-compatible with "cv2 T2", or
  - T1 is not reference-related to T2 and the initializer expression can be implicitly converted to an rvalue of type "cv3 T3" where "cv1 T1" is reference-compatible with "cv3 T3" (this conversion is selected by enumerating the applicable conversion functions (13.3.1.6) and choosing the best one through overload resolution (13.3)),

then the reference is bound to the initializer expression rvalue in the first case and to the object that is the result of the conversion in the second case (or, in either case, to the appropriate base class subobject of the object).

[ *Example:*

```
    struct A { };
    struct B : A { } b;
    extern B f();
    const A& rca = f();           // bound to the A subobject of
the B rvalue.
    A&& rcb = f();               // same as above
    struct X {
        operator B();
    } x;
    const A& r = x;              // bound to the A subobject of
the result of the conversion
```

-- *end example* ]

- If the initializer expression is an rvalue, with T2 an array type, and "cv1 T1" is reference-compatible with "cv2 T2," the reference is bound to the object represented by the rvalue (see 3.10). ...

[ *Example:*

```
    struct A { };
    struct B : A { } b;
    extern B f();
    const A& rca = f();           // bound to the A subobject of
the B rvalue.
    A&& rcb = f();              // same as above
    struct X {
        operator B();
    } x;
    const A& r = x;             // bound to the A subobject of
the result of the conversion
-- end example ]
```

- If the initializer expression is an rvalue, with T2 an array type, and T1 is reference-compatible with T2, the reference is bound to the object represented by the rvalue (see 3.10). ....



## US 66

Author: Jason Merrill

2010-05-10

Revision 2

# List-construction fallback

## The Problem

In core-16131, Daniel Krügler offered the following example:

a) If I have no initializer-list c'tors, I can write

```
struct A{} a;

struct S {
    S(A, A); // #1
    S(int, double, bool); // #2
};

S s1{a, a}; // OK, calls #1
S s2{12, 3.1, false}; // OK, calls #2
```

b) If I now add an initializer-list c'tor like this:

```
#include <initializer_list>

struct A{} a;
struct B{};

struct S {
    S(A, A); // #1
    S(int, double, bool); // #2
    S(std::initializer_list<B>); // #3
};

S s1{a, a}; // Error, does not match #3
S s2{12, 3.1, false}; // Error, does not match #3
```

This struck him as unfortunate; if there are list constructors but none are viable, we should subsequently try to match a non-list constructor. This is especially problematic if S has an additional constructor:

```
#include <initializer_list>

struct A{} a;
struct B{};
struct C {
    C(A, A);
};

struct S {
    S(A, A); // #1
    S(std::initializer_list<B>); // #3
    S(C); // #4
```

```
};
```

```
S s1{a, a}; // Matches #4 instead of #1
```

This was not a problem in N2385, but my changes to formalize the rules ended up oversimplifying the handling of classes with list constructors in a way that created this problem.

The right answer seems to be to do overload resolution for list-initialization of a class in two phases: first look for a suitable list constructor, and if none is found then look for a suitable non-list constructor.

## Proposed Wording

Change 8.5.4 [dcl.init.list] as follows:

....

- Otherwise, if T is a class type, constructors are considered. ~~If T has an initializer list constructor, the argument list consists of the initializer list as a single argument; otherwise, the argument list consists of the elements of the initializer list.~~ The applicable constructors are enumerated (13.3.1.7) and the best one is chosen through overload resolution (13.3.1.7, 13.3). If a narrowing conversion (see below) is required to convert any of the arguments, the program is ill-formed.

....

Change 13.3.1.7 [over.match.list] as follows:

When objects of non-aggregate class type T are list-initialized (8.5.4), overload resolution selects the constructor ~~in two phases~~ as follows, where T is the cv-unqualified class type of the object being initialized:

- ~~If T has an initializer list constructor (8.5.4),~~ Initially, the candidate functions are the initializer-list constructors (8.5.4) of the class T and the argument list consists of the initializer list as a single ~~argument;~~ argument.
- ~~otherwise,~~ If no viable initializer-list constructor is found, overload resolution is performed again, where the candidate functions are all the constructors of the class T and the argument list consists of the elements of the initializer list.
- ~~For direct list initialization, the candidate functions are all the constructors of the class T.~~
- ~~For~~ In copy-list-initialization, ~~the candidate functions are all the constructors of T.~~ However, if an explicit constructor is chosen, the initialization is ill-formed. [ Note: This differs from other situations (13.3.1.3, 13.3.1.4), where only converting constructors are considered for copy-initialization. This restriction only applies if this initialization is part of the final result of overload resolution -- end note ]

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 76**

**Comment 14.8.2/9: extern template prevents inlining functions not marked inline.**

**Technical details:**

Existing compilers often inline functions that weren't explicitly declared inline. To inline a function, they need to instantiate its template. Yet the FCD [temp.explicit]p9 says, "Except for inline functions and class template specializations, explicit instantiation declarations have the effect of suppressing the implicit instantiation of the entity to which they refer." This means that adding an explicit instantiation declaration can affect performance, even though the user only intended to suppress out-of-line copies of functions. N1987 doesn't seem to expect any changes in inlining behavior.

**Proposed resolution:**

- Remove [temp.explicit] / p9.
- Replace the note in [temp.explicit] / p10 with: "The intent of this rule is to allow compilers to avoid emitting out-of-line copies of template functions, while still allowing them to inline those functions. The compiler may assume that another translation unit will supply the body."

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 88**

**Clause 20.2.5 [allocator.requirements]**

Allocator Interface Not Backward Compatible

Date: 2010-05-21

Author: P.J. Plauger

Clause 20.2.5, Allocator requirements, adds functionality to the C++03 (old) allocator interface. It also endeavors to maintain compatibility with old allocators, so that old allocators can be used with new users of allocators, such as containers, function objects, regular expressions, shared pointers, string streams, etc. To supply the missing functionality for old allocators, C++0X also adds **allocator traits**, in the form of a new template class `allocator_traits`. This template class uses various template metaprogramming techniques to determine whether it can obtain a given type or function from an allocator, or whether it should supply a default instead. If code that uses allocators is changed to work through the intermediary of `allocator_traits`, it can access (most of) the functions and types defined in the old allocator interface.

Unfortunately, the new allocator requirements go too far. They permit a new allocator to omit many of the types and functions required by old allocators. `allocator_traits` will fill in the blanks. The new allocator requirements also eliminate some functions and types required of old allocators, presumably on the assumption that programmers don't need them any more and won't miss them. In so doing, however, the new allocator requirements fail to maintain the other half of the allocator interface -- they permit, even encourage, programmers to write **new** allocators that don't work with **old** users of allocators. Indeed, every existing container, function object, etc. listed above has to be reviewed and probably revised to work in the new world of abbreviated allocators and allocator traits.

Things that use allocators in the Standard C++ library are not a problem. Implementers are obliged to update the standard containers, etc. to match the new requirements. User-defined things that use the library-supplied default allocator are also not a problem. The default allocator retains even those features of old allocators that are no longer required by the new allocator requirements. But in general:

- Old allocators work with old allocator users, as we would expect.
- New allocators work with new allocator users, as we would expect.
- Old allocators work with new allocator users, if the new allocator users use `allocator_traits`, as needed.
- Old allocator users work with the new default allocator.
- But otherwise, old allocator users **do not** work with new allocators.

In designing the new allocator/user interface, committee discussions focused overwhelmingly on the problem of retaining backward compatibility between old allocators and the new containers supplied by the Standard C++ library. There was strong sentiment that old code should not have to be modified to work with the new library, even though the new library demanded more of allocators than before. The committee could have simply required that old allocators have to work only with the standard containers ("just make it work"), but it chose to go farther. By publishing the `allocator_traits` interface, C++0X provides tools for user-supplied containers, etc. to work with either old or new allocators.

The same consideration was not extended nearly as far to old containers, etc. If they want to survive in the C++0X world of new allocators, they will typically have to be rewritten. The rewrite is largely mechanical, but a careful review has to be done. To the best of my knowledge, the committee has not commented on this issue. There may be no sentiment for preserving old allocator users, but that weakens the case for indulging in heroics to preserve old allocators.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

### Proposal 0

If the committee is willing to mandate at least **some** rewrite of old code, it could simply suggest that all allocators be changed, if necessary, to (behave as if they) have the default allocator as a public base class, as in:

```
template<class T, class... Rest>
class Myalloc
    : public std::allocator<T>
{.....}
```

This solves the problem of defaulting the new functions and types that are missing in the old allocators. It ensures that old containers, etc. work with new allocators. And it provides for future additions to the allocator interface. Deriving from `std::allocator<T>` has long been a common idiom. The author of `Myalloc` merely has to supply those things that change from the default allocator, and those things that cannot be inherited from a base class (such as constructors and `rebind`). Allocators of this form do not have to change to conform to this new requirement.

The fix is to remove the defaults from Table 41 in 20.2.5, and to restore the requirements dropped from the old allocator specification. `allocator_traits` can also be removed, or it can be retained for those who believe that abbreviating allocators will prove to be a good idea.

### Proposal 1

If the preceding Proposal is too sweeping, at the very least the C++ Standard should fix the backward compatibility problem outlined above:

- Restore to the allocator requirements the functions and types that were dropped (primarily `reference`, `const_reference`, and `address`).
- Drop the default behavior for all allocator requirements inherited from C++2003.

Detailed changes for Proposal 1 are summarized below.

A salutary effect of adopting this proposal is to reduce the amount of new code that has to be added to `<memory>`. In its current form, `allocator_traits` adds roughly 200 lines of code to this header, which must be included in every program that uses a library container, **including `string`**. The code removed merely gives the programmer the ability to abbreviate allocators, solving a non-problem that can usually be addressed, if needed, by the one-line addition shown above.

Note that a library implementation can still choose to provide all the defaults currently listed in `allocator_traits`, as a conforming extension. If it turns out that there really is a market for abbreviated allocators, the evidence will be apparent after a couple of years of field experience. The C++ Standard can then be modified with a Technical Corrigendum or Amendment. But if we allow abbreviations from the start, it's much harder to later disallow abbreviations in user code.

### Proposal 2

We can simplify `allocator_traits` even farther by noting that only two kinds of allocators really need to be supported, old and new. As before, there's no compelling need for various levels of newness. `allocator_traits` need only have code that recognizes the presence of **any** new feature to conclude that **all** new features are present.

Programmers sophisticated enough to write allocators know to add a base class, as above, when they are content with the defaults for some features. Adding a hundred-odd lines of code to practically **every** compile penalizes simple programs that, say, just use `string`, all to save sophisticated programs from writing perhaps one line of text. This is a clear violation of the rule that you don't pay for what you don't use.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

The fix is simple:

In 20.2.5/2 Allocator requirements, after the sentence "Table 42 specifies a default for a given expression."  
ADD "If an allocator supplies any requirement that has a default, it shall supply all requirements that have a default."

**Proposal 1 Details**

In 20.2.5 Allocator requirements:

ADD to Table 41 the line:

mt a value of type T&

REMOVE from Table 42 defaults for pointer, const\_pointer, size\_type, difference\_type, rebind, allocate with hint, and max\_size.

Add to Table 42 the lines:

```
X::reference (no Return) Assertion/note: T&
X::const_reference (no Return) Assertion/note: const T&
p.address(mt) X::pointer &mt
p.address(t) X::const_pointer &t
a.construct(p, t) (not used) Effect: Constructs an object of type T at p.
a.destroy(p) (not used) Effect: Destroys the object at p
```

REMOVE from 20.2.5/2 the sentence "A user specialization of allocator\_traits may provide different defaults and may provide defaults for different requirements than the primary template."

In 20.2.5/3, DELETE: "If Allocator is a class template instantiation of the form SomeAllocator<T, Args>, where Args is zero or more type arguments, and Allocator does not supply a rebind member template, the standard allocator\_traits template uses SomeAllocator<U, Args> in place of Allocator::rebind<U>::other by default. For allocator types that are not template instantiations of the above form, no default is provided."

In 20.3.5/5, CHANGE the example from:

```
template <class Tp>
struct SimpleAllocator {
    typedef Tp value_type;
    SimpleAllocator(ctor args);
    template <class T> SimpleAllocator(const SimpleAllocator<T>& other);
    Tp *allocate(std::size_t n);
    void deallocate(Tp *p, std::size_t n);
};
```

TO:

```
template <class Tp>
struct SimpleAllocator
    : public std::allocator<Tp> {
    template <class T>
        SimpleAllocator(const SimpleAllocator<T>& other);
    template <class T>
        SimpleAllocator& (const SimpleAllocator<T>& other)
        { return *this; }
    template <class T>
        struct rebind {
```

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

```
typedef SimpleAllocator<T>& other; }  
};  
The allocator can add just those functions it wishes to override,  
such as allocate and deallocate, to obtain behavior that differs  
from std::allocator<Tp>.
```

In 20.9.4 Allocator traits, CHANGE:

```
typedef see below pointer;  
typedef see below const_pointer;  
typedef see below void_pointer;  
typedef see below const_void_pointer;  
  
typedef see below difference_type;  
typedef see below size_type;  
.....  
template <class T> using rebind_alloc = see below;
```

TO:

```
typedef Alloc::pointer pointer;  
typedef Alloc::const_pointer const_pointer;  
typedef pointer_traits<pointer>::rebind<void> void_pointer;  
typedef pointer_traits<pointer>::rebind<const void> const_void_pointer;  
  
typedef Alloc::difference_type difference_type;  
typedef Alloc::size_type size_type;  
.....  
template <class T> using rebind_alloc = Alloc::rebind<T>::other;
```

In 20.9.4.1 Allocator traits member types, REMOVE:

typedef see below pointer;

Type: Alloc::pointer if such a type exists; otherwise, value\_type\*.

typedef see below const\_pointer;

Type: Alloc::const\_pointer if such a type exists; otherwise, pointer\_traits<pointer>::rebind<const value\_type>.

typedef see below void\_pointer;

Type: Alloc::void\_pointer if such a type exists; otherwise, pointer\_traits<pointer>::rebind<void>.

typedef see below const\_void\_pointer;

Type: Alloc::const\_void\_pointer if such a type exists; otherwise, pointer\_traits<pointer>::rebind<const void>.

typedef see below difference\_type;

Type: Alloc::difference\_type if such a type exists; otherwise, ptrdiff\_t.

typedef see below size\_type;

Type: Alloc::size\_type if such a type exists; otherwise, size\_t.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

.....

template <class T> using rebind\_alloc = see below;

Template alias: Alloc::rebind<T>::other if such a type exists; otherwise, Alloc<T, Args> if Alloc is a class template instantiation of the form Alloc<U, Args>, where Args is zero or more type arguments; otherwise, the instantiation of rebind\_alloc is ill-formed.

In 20.9.4.2 Allocator traits static member functions, CHANGE:

static pointer allocate(Alloc& a, size\_type n, const\_void\_pointer hint);

Returns: a.allocate(n, hint) if that expression is well-formed; otherwise, a.allocate(n).

.....

template <class T, class... Args>  
static void construct(Alloc& a, T\* p, Args&&... args);

Effects: calls a.construct(p, std::forward<Args>(args)...) if that call is well-formed; otherwise, invokes ::new (static\_cast<void\*>(p)) T(std::forward<Args>(args)...)..

template <class T>  
static void destroy(Alloc& a, T\* p);

Effects: calls a.destroy(p) if that call is well-formed; otherwise, invokes p->~T().

static size\_type max\_size(Alloc& a);

Returns: a.max\_size() if that expression is well-formed; otherwise, numeric\_limits<size\_type>::max().

TO:

static pointer allocate(Alloc& a, size\_type n, const\_void\_pointer hint);

Returns: a.allocate(n, hint).

.....

template <class T, class... Args>  
static void construct(Alloc& a, T\* p, Args&&... args);

Effects: calls a.construct(p, std::forward<Args>(args)...) if that call is well-formed; otherwise, invokes ::new (static\_cast<void\*>(p)) T(std::forward<Args>(args)...)..

[Note: Alloc always defines a signature equivalent to construct(pointer, const value\_type&). -- end note]

template <class T>  
static void destroy(Alloc& a, T\* p);

Effects: calls a.destroy(p) if that call is well-formed; otherwise, invokes p->~T().

[Note: Alloc always defines a signature equivalent to destroy(pointer). -- end note]



FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

```
static size_type max_size(Alloc& a);
```

Returns: a.max\_size().

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 96**

**20.3.5.2 [pairs.pair] Move/forward confusion in pair and tuple construction and assignment**

**Section:** 20.3.5.2 [pairs.pair], 20.4.2.1 [tuple.cnstr], and 20.4.2.2/6 [tuple.assign]

**Submitter:** Pablo Halpern

**Discussion**

Note: This issue overlaps with LWG 1326.

There are actually two issues that are intertwined, so it makes sense to resolve them together.

The first issue is that the terms MoveConstructible, MoveAssignable, CopyAssignable are being misused in [pair] and [tuple] to describe heterogeneous construction or assignment, even though the terms are defined only in terms of a single type.

The second issue is that std::move is being used where std::forward is required. In particular, std::move will erroneously convert an lvalue-reference to an rvalue-reference whereas std::forward will not. Also, the terms “move-constructs” and “move-assigns” are being used in contexts where std::forward should be used, implying the wrong semantic.

**Proposed Resolution**

Change [pairs.pair]/6 as follows:

```
template<class U, class V> pair(pair<U, V>&& p);
6   Effects: The constructor initializes first with std::moveforward<U>(p.first) and second with
      std::moveforward<V>(p.second).
```

Change [pairs.pair] paragraphs 12-15 as follows:

```
pair& operator=(pair&& p);
12  Effects: Assigns to first with std::moveforward<T1>(p.first) and to second with std::
      moveforward<T2>(p.second).
13  Returns: *this.
```

```
template<class U, class V> pair& operator=(pair<U, V>&& p);
14  Effects: Assigns to first with std::moveforward<U>(p.first) and to second with std::
      moveforward<V>(p.second).
15  Returns: *this.
```

Change [tuple.cnstr] paragraph 6 as follows:

```
template <class... UTypes>
explicit tuple(UTypes&&... u);
6   Requires: Each type in Types shall satisfy the requirements of MoveConstructible (Table 34) be
      constructible from the corresponding type in UTypes&& . sizeof...(Types) == sizeof...(UTypes).
```

Change [tuple.cnstr] paragraphs 11-20 as follows:

```
tuple(tuple&& u);
```

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

- 10 *Requires:* Each type in Types shall ~~shall~~ satisfy the requirements of MoveConstructible (Table 34).
- 11 *Effects:* ~~Move~~-constructs each element of \*this with the corresponding element of [std::forward<Types>\(u\)](#).
- template <class... UTypes> tuple(const tuple<UTypes...>& u);
- 12 *Requires:* Each type in Types shall be constructible from the corresponding type in UTypes&&. sizeof...(Types) == sizeof...(UTypes).
- 13 *Effects:* Constructs each element of \*this with the corresponding element of [std::forward<UTypes>\(u\)](#).
- 14 [ *Note:* enable\_if can be used to make the converting constructor and assignment operator exist only in the cases where the source and target have the same number of elements. —end note ]
- template <class... UTypes> tuple(tuple<UTypes...>&& u);
- 15 *Requires:* Each type in Types shall ~~shall the requirements of MoveConstructible (Table 34)~~[be constructible](#) from the corresponding type in UTypes&&. sizeof...(Types) == sizeof...(UTypes).
- 16 *Effects:* ~~Move~~-constructs each element of \*this with the corresponding element of [std::forward<UTypes>\(u\)](#).  
[ *Note:* enable\_if can be used to make the converting constructor and assignment operator exist only in the cases where the source and target have the same number of elements. —end note ]
- template <class U1, class U2> tuple(const pair<U1, U2>& u);
- 17 *Requires:* The first type in Types shall be constructible from U1 and the second type in Types shall be constructible from U2. sizeof...(Types) == 2.
- 18 *Effects:* Constructs the first element with u.first and the second element with u.second.
- template <class U1, class U2> tuple(pair<U1, U2>&& u);
- 19 *Requires:* The first type in Types shall ~~shall the requirements of MoveConstructible (Table 34)~~[be constructible](#) from U1&& and the second type in Types shall be ~~move~~-constructible from U2&&. sizeof...(Types) == 2.
- 20 *Effects:* Constructs the first element with [std::-moveforward<U1>\(u.first\)](#) and the second element with [std::-moveforward<U2>\(u.second\)](#).

Change [tuple.assign] paragraph 6 as follows:

- 6 *Effects:* ~~Move~~-assigns each element of [std::forward<Types>\(u\)](#) to the corresponding element of \*this.

Change [tuple.assign] paragraphs 11 and 12 as follows:

- 11 *Requires:* Each type in Types shall ~~satisfy the requirements of MoveConstructible (Table 34)~~[be assignable](#) from the corresponding type in UTypes&&. sizeof...(Types) == sizeof...(UTypes).
- 12 *Effects:* ~~Move~~-assigns each element of [std::forward<UTypes>\(u\)](#) to the corresponding element of \*this.

Change [tuple.assign] paragraph 14 as follows:

- 14 *Requires:* The first type in Types shall ~~shall satisfy the requirements of MoveAssignable (Table 36)~~[be assignable](#) from U1 and the second type in Types shall ~~shall satisfy the requirements of MoveAssignable (Table 36)~~[be assignable](#) from U2. sizeof...(Types) == 2.

Change [tuple.assign] paragraphs 18 and 19 as follows:

- 18 *Requires:* The first type in Types shall be Assignable from U1&& and the second type in Types shall be Assignable from U2&&. sizeof...(Types) == 2.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

- 19 *Effects:* Assigns std::~~move~~forward<U1>(u.first) to the first element of \*this and std::~~move~~forward<U2>(u.second) to the second element of \*this.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

US 106

[pointer.traits] pointer\_traits lacks size\_type

Section: 20.9.3 [pointer.traits]

Submitter: Pablo Halpern

Discussion:

The pointer\_traits template has a difference\_type member, but not a corresponding size\_type member. This asymmetry causes issues for allocators. Currently, the default type for allocator\_traits<A>::difference\_type is ptrdiff\_t and the default type for allocator\_traits<A>::size\_type is size\_t. However, it would be more useful and natural for allocator\_traits<A>::difference\_type to default to pointer\_traits<A::pointer>::difference\_type and allocator\_traits<A>::size\_type to default to pointer\_traits<A::pointer>::size\_type. The former is currently possible but the latter is not because of the absence of size\_type in pointer\_traits.

Proposed Resolution:

Add a new type to [pointer.traits]/1:

```
typedef see below difference_type;  
typedef see below size_type;
```

And a paragraph after [pointer.traits]/2:

```
typedef see below difference_type;  
2          Type: Ptr::difference_type if such a type exists; otherwise, std::ptrdiff_t.  
typedef see below size_type;  
          Type: Ptr::size_type if such a type exists; otherwise, std::size_t.
```

In section [allocator.requirements], Table 42, replace the defaults for difference\_type and size\_type:

|                    |                        |  |  |
|--------------------|------------------------|--|--|
| X::size_type       | unsigned integral type | a type that can represent the size of the largest object in the allocation model.          | <u>size_t</u> <u>pointer_traits&lt;pointer&gt;::size_type</u>          |
| X::difference_type | signed integral type   | a type that can represent the difference between any two pointers in the allocation model. | <u>ptrdiff_t</u> <u>pointer_traits&lt;pointer&gt;::difference_type</u> |

Change the definitions of difference\_type and size\_type in [allocator.traits] paragraphs 5 and 6:

```
typedef see below difference_type;
```

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

5           Type: Alloc::difference\_type if such a type exists; otherwise,  
~~ptr\_diff\_t~~pointer\_traits<pointer>::difference\_type.

typedef see *below* size\_type;

6           Type: Alloc::size\_type if such a type exists; otherwise, ~~size\_t~~pointer\_traits<pointer>::size\_type.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 107**

**Clauses 20.1, 20.9.1, 20.9.6**

Scoped Allocator Adaptor Inappropriate for <memory>

Date: 2010-05-21

Author: P.J. Plauger

Template class `scoped_allocator_adaptor` (20.9.6) requires a nontrivial piece of code to implement. The reference implementation (referenced in N2982) runs to nearly 800 lines. C++0X has already added significantly to the header <memory> with features such as `shared_ptr` and `unique_ptr`, but scoped allocators still represent a significant addition to a widely used header.

Dinkumware has already received push back from customers who dislike the longer compile times brought on by C++0X features that they don't even use. This may appear to be carping, given the speed of today's computers, but it is not. If an overnight build doesn't complete overnight, scheduling problems arise. We have repartitioned our headers more than once, in response to this feedback. But that technique goes only so far.

WG21 can probably justify placing `shared_ptr` and `unique_ptr` in <memory>, since the expectation is that these template classes are also likely to be widely used, and by a broad range of programmers. But it is hard to make the same case for `scoped_allocator_adaptor`. The latter requires a sophisticated knowledge of allocators and how they might interact with containers and container elements.

Template class `scoped_allocator_adaptor` has the virtue of being self-contained. It has no required uses within the Standard C++ library, and there are no references to it from other parts of the C++ Standard. Hence, it is both easy and beneficial to move this template class to a new header.

The fix is to add the new header <scoped\_allocator>:

IN 20.1 General, Table 30, ADD the line:

```
20.xx Scoped allocators    <scoped_allocator>
```

IN 20 General Utilities Library, INSERT/APPEND a new section 20.xx, Class `scoped_allocator_adaptor`

FROM 20.9.1 Header <memory> synopsis, MOVE to 20.xx the new 20.xx.1 Header <scoped\_allocator> synopsis:

```
// 20.9.6, scoped allocator adaptor
template <class OuterAlloc, class... InnerAlloc>
    class scoped_allocator_adaptor;
template <class OuterA1, class OuterA2, class... InnerAllocs>
bool operator==(const scoped_allocator_adaptor<OuterA1, InnerAllocs...>& a,)
const scoped_allocator_adaptor<OuterA2, InnerAllocs...>& b);
template <class OuterA1, class OuterA2, class... InnerAllocs>
bool operator!=(const scoped_allocator_adaptor<OuterA1, InnerAllocs...>& a,)
const scoped_allocator_adaptor<OuterA2, InnerAllocs...>& b);
```

MOVE 20.9.6 Scoped allocator adaptor to the new 20.xx.2.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 115**

**23.2.1 [container.requirements.general] Terminology for constructing container elements**

**Section:** 23.2.1 [container.requirements.general]

**Submitter:** Pablo Halpern

**Discussion**

The resolution to issue 704 added a number of requirements to the container operations in order to properly constrain the elements of the containers. Unfortunately, the current wording effectively *redefines* the terms CopyConstructible and MoveConstructible and the phrase “constructible with *args*” so that they have different meanings in the containers section than in the rest of the standard. This use of terminology is not only confusing and vague, it is also not applied correctly through the section. There are some cases, in fact, when the term CopyConstructible is used in its original meaning, but the reader would have no way to know that.

The best solution is to choose an entirely new and more precise set of terms and apply them consistently and correctly in the containers section. I nominate the terms *X can copy-insert T*, *X can move-insert T*, and *X can construct-insert T with args* as replacements for the above terms, where X is the container type, A is X's allocator type and T is its element type.

**Proposed Resolution:**

Replace [container.requirements.general]/15 as follows:

- 15 ~~The descriptions of the requirements of the type T in this section use the terms CopyConstructible, MoveConstructible, constructible from \*, and constructible from args. These terms are equivalent to the following expression using the appropriate arguments:~~

Given a container type X having an allocator type A and a value type T and given an lvalue m of type A, a pointer p of type T\*, a value v of type T, or a value rv of type rvalue-of-T, the following terms are defined. (If X is not allocator-aware, the terms below are defined as if A were std::allocator<T>.):

X can copy-insert T means that the following expression is well-formed:

allocator\_traits<A>::construct(m, p, v);

X can move-insert T means that the following expression is well-formed:

allocator\_traits<A>::construct(m, p, rv);

A can construct-insert T from args for zero or more arguments, args, means that the following expression is well-formed:

allocator\_traits<A>::construct(m, p, args);

[ Note: The default of construct in std::allocator will call

::new((void\*) p) T(args)

but specialized allocators may choose a different definition. – end note ]



FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

A review of the rest of section 23 shows that one can substitute the above terms in all cases where the phrase "T is CopyConstructible", etc.. A complete resolution will need to spell out each individual case, as sometimes the wording varies as in "T shall be CopyConstructible" or "value\_type is constructible from," etc.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 117**

**Clause 23.3.3**

Problems with `forward_list::erase_after`

Date: 2010-05-21

Author: P.J. Plauger

The iterator returns from the two overloads of `forward_list::erase_after` were removed as part of the adoption of N2988. The rationale given was, in its entirety:

"I believe that `erase_after` should return `void`. Returning its argument does nothing (potentially for a price), and confuses the programmer. I think it's better that the programmer notice that `forward_list` is different from other containers (and better not to pass a value through a function unnecessarily)."

The first and last sentences are, by their own admission ("I believe", "I think"), unfounded conjectures.

The second sentence is misleading, since only the second overload -- `erase_after(iterator first, iterator last)` - happens to return one of its arguments. The first overload -- `erase_after(iterator before)` -- does not. Both return an iterator designating the first element after the element(s) removed, or `end()` if the last element is removed.

The parenthetic remark in the second sentence happens not to be true, since the iterator to be returned arises naturally in the process of erasing the elements.

The **real price** that this change has exacted is to destroy backward compatibility with the Committee Draft, implementations of which are now widely available to programmers. It takes more than conjectures to justify making such a breaking change.

It is also worth noting that the Dinkumware implementation of `forward_list` calls `erase_after` thirteen times for internal purposes. In five of these cases, the code makes use of the return value. This is admittedly a small sample of all the code now using `forward_list`, but enough to refute the notion that the return value of `erase_after` "does nothing" or merely "confuses the programmer." Whether it is "better" to enforce a different style of programming is a matter of taste.

The fix is to restore the wording in the Committee Draft:

IN 23.3.3/3, CHANGE:

```
void erase_after(const_iterator position);  
void erase_after(const_iterator position, iterator last);
```

TO:

```
iterator erase_after(const_iterator position);  
iterator erase_after(const_iterator position, iterator last);
```

IN 23.3.3.5/18-21, CHANGE:

```
void erase_after(const_iterator position);
```

Requires: The iterator following position is dereferenceable.

Effects: Erases the element pointed to by the iterator following position.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

```
void erase_after(const_iterator position, iterator last);
```

Requires: All iterators in the range (position,last) are dereferenceable.

Effects: Erases the elements in the range (position,last).

TO:

```
iterator erase_after(const_iterator position);
```

Requires: The iterator following position is dereferenceable.

Effects: Erases the element pointed to by the iterator following position.

Returns: An iterator pointing to the element following the one that was erased, or end() if no such element exists.

```
iterator erase_after(const_iterator position, iterator last);
```

Requires: All iterators in the range (position,last) are dereferenceable.

Effects: Erases the elements in the range (position,last).

Returns: last

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 118**

**Comment 23.5: Some unordered associative container operations have undesirable complexities when the container is implemented using singly linked lists.**

**Technical details**

A common implementation technique for hash tables is as a vector of buckets, where each bucket is a separate singly linked list. The TR1 unordered associative container specification was intended to allow that implementation, among others.

This representation is desirable in some ways, but it does have a known problem: if the current element happens to be the last in a bucket, incrementing it to get to the next element requires a linear scan over the bucket array to find the next non-empty bucket. If the load factor (the ratio of elements to buckets) is very low, this means that iterator increment is expensive.

In itself this is unfortunate but tolerable ---iteration through hash tables is less common than element insertion or lookup, and there is long experience with STL hash table implementations (such as the SGI implementation) that have this property. FCD unordered associative containers, however, make the problem worse. The `erase()` member function is defined to return an iterator pointing to the element after the one that's erased, which means that a common hash table operation can result in invoking the potentially expensive iterator increment

**Proposed resolution**

None. In principle there are several possibilities, including:

- Change the signature of `erase` so that it returns `void`. This would return us to the situation with SGI STL hash tables, where iterator increment is potentially expensive but programmers can realistically avoid it in most cases.
- Impose a minimum load factor, resizing the tables when necessary. This might involve changing the iterator invalidation rules.
- Require implementations to use a doubly linked list implementation instead. This would be undesirable because it would greatly increase the space requirements, for essentially technical reasons.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

US 120

**Comment 25.2.12/1: `is_permutation` is underspecified for anything but the simple case where both ranges have the same value type and the comparison function is an equivalence relation.**

**Technical details**

There are no requirements on `is_permutation`'s template parameters beyond those implied by their names, and by the fact that the returns clause is written in terms of `std::equal`. It thus appears at first sight that `is_permutation` has the same constraints as `equal`, and that no further constraints are needed. This apparent simplicity is misleading.

The returns clause says that `is_permutation(f1,l1,f2)` is true iff there is some permutation of the elements in `[f1, l1)`, which we can call `[f1',l1')`, such that `equal(f1',l1',f2)` is true. However, it doesn't say how to find that permutation or prove that no such permutation exists. That's obviously the crux of the algorithm, and we need to think about the constraints that are needed for it to work. So let's think about a definition that describes how we actually compute whether `is_permutation(f1,l1,f2)` is true or false.

The n2986 reference implementation describes the actual algorithm that's intended: loop through `[f1, l1)`. For each value `v`, count how many times it appears in `[f1, l1)` and in `[f2, l2)`. If the counts differ, then `[f1, l1)` can't be a permutation of `[f2, l2)`. If the counts are the same for every `v` and the ranges have the same length, then it is. There are several variations on this technique, but all the variations I know of involve choosing some value `v` in one of the input ranges and looking at other instances of the same value in the same range.

This means that, unlike the case of `std::equal`, it isn't enough to require the existence of comparison between a value of type `iterator_traits<ForwardIterator1>::value_type` and one of type `iterator_traits<ForwardIterator2>::value_type`. At a minimum, we also need to perform comparisons between two values of type `iterator_traits<ForwardIterator1>::value_type` and/or two values of type `iterator_traits<ForwardIterator2>::value_type`.

Even that isn't enough, though --it's enough to get `is_permutation` to compile, but not enough for it to give sane results. It's easy to come up with examples of two types `V1` and `V2` such that there are sensible comparisons between `V1` and itself, `V2` and itself, and `V1/V2`, but where `is_permutation` would give nonsensical results. Consider, for example:

```
struct A { int val; char tag; };
struct B { int val; string tag; };

bool operator==(A x, A y) { return x.val == y.val && x.tag == y.tag; }
bool operator==(B x, B y) { return x.val == y.val && x.tag == y.tag; }
bool operator==(A x, B y) { return x.val == y.val; }
```

Now suppose we call `is_permutation` on two sequences `c1` and `c2`, where `c1` is `[(1, 'a'), (2, 'b'), (1, 'c')]` and `c2` is `[(2, "x"), (1, "y"), (1, "z")]`. Do we want the result to be true, or false? Taken literally the returns clause in the FCD says that in this example the return value should be true, but the algorithm described above, and the n2986 reference implementation, will return false. (And I don't know of any practical algorithm that would do otherwise.) It will examine `(1, a)`, find that it appears once in `c1` and twice in `c2`, and return false.

The fundamental problem with this example is that we have three elements, `x`, `y`, and `z`, such that `x == y`, `y == z`, and `x != z`. That should look familiar. If `operator==` is allowed to violate the requirements of an equivalence relation, one can construct equally pathological examples using only a single value type. This algorithm has to compare multiple elements against each other, so we won't get sensible results unless `operator==` obeys sensible axioms.

**Proposed resolution**

Informally: restrict `is_permutation` to the case where it is well specified. More formally:  
Add a new paragraph to `[alg.is_permutation]`, before the existing paragraph 1:

```
Requires: ForwardIterator1 and ForwardIterator2 shall have the same value type. The
comparison function shall be an equivalence relation.
```

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

[There are other possible fixes. We could allow heterogeneous comparisons but constrain them more. Or, more radically, we could get rid of `is_permutation` entirely and describe in words what equality comparison should mean for hash multisets/multimaps, or even get rid of or change the definition of hash multiset/multimap equality. We recommend the resolution above because the former option would be complicated and easy to get wrong, while the latter two are larger changes than necessary.]

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 136**

**Clause 26.8**

Problems with Floating-point Test Functions

Date: 2010-05-21

Author: P.J. Plauger

In 26.8 C Library, Table 116 lists as "Templates" a number of functions defined as macros in the C Standard:

**Templates:**

```
fpclassify  isgreaterequal  islessequal  isnan      isunordered  
isfinite    isinf          islessgreater  isnormal  signbit  
isgreater   isless
```

This is inaccurate. Each macro is better described in C++ as three overloads, with operand(s) of type float, double, and long double. Moreover, as with the math functions described in this same section, each of the functions has "sufficient additional overloads" to simulate the effect of the C type-generic functions.

The fix is to change 26.8 as follows:

In Table 116, CHANGE "Functions:" TO "Math Functions:" and CHANGE "Templates:" TO "Classification/comparison Functions:".

BEFORE paragraph 10 ("Moreover, there shall be...") INSERT:

The classification/comparison functions behave the same as the C macros with corresponding names defined in 7.12.3, Classification macros, and 7.12.14, Comparison macros in the C Standard. Each function is overloaded for the three floating-point types, as follows:

```
namespace std {  
int fpclassify(float x);  
bool isfinite(float x);  
bool isinf(float x);  
bool isnan(float x);  
bool isnormal(float x);  
bool signbit(float x);
```

```
bool isgreater(float x, float y);  
bool isgreaterequal(float x, float y);  
bool isless(float x, float y);  
bool islessequal(float x, float y);  
bool islessgreater(float x, float y);  
bool isunordered(float x, float y);
```

```
int fpclassify(double x);  
bool isfinite(double x);  
bool isinf(double x);  
bool isnan(double x);  
bool isnormal(double x);  
bool signbit(double x);
```

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

```
bool isgreater(double x, double y);  
bool isgreaterequal(double x, double y);  
bool isless(double x, double y);  
bool islessequal(double x, double y);  
bool islessgreater(double x, double y);  
bool isunordered(double x, double y);
```

```
int fpclassify(long double x);  
bool isfinite(long double x);  
bool isinf(long double x);  
bool isnan(long double x);  
bool isnormal(long double x);  
bool signbit(long double x);
```

```
bool isgreater(long double x, long double y);  
bool isgreaterequal(long double x, long double y);  
bool isless(long double x, long double y);  
bool islessequal(long double x, long double y);  
bool islessgreater(long double x, long double y);  
bool isunordered(long double x, long double y);  
} // namespace std
```

DELETE paragraph 11 ("The templates defined...") and paragraph 12 ("The templates behave...").



FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 137**

**Clause: 27.7**

Problems with Iostreams Member Functions

Date: 2010-05-21

Author: P.J. Plauger

There are several problems with member functions in `basic_istream` and `basic_ostream`:

-- `putback` is obliged to fail at end of file. Moreover, the member function doesn't clear `eofbit`, as it should. The current wording says:

Effects: Behaves as an unformatted input function (as described in 27.7.1.3, paragraph 1). After constructing a sentry object, if `!good()` calls `setstate(failbit)` which may throw an exception, and return.

Both problems can be solved by first clearing `eofbit`:

CHANGE 27.7.1.3/34 (`putback`) first sentence from:

Behaves as an unformatted input function (as described in 27.7.1.3, paragraph 1).

TO:

Behaves as an unformatted input function (as described in 27.7.1.3, paragraph 1), except that the function first clears `eofbit`.

-- A similar problem exists with `unget`, with a similar fix:

CHANGE 27.7.1.3/36 (`unget`) first sentence from:

Behaves as an unformatted input function (as described in 27.7.1.3, paragraph 1).

TO:

Behaves as an unformatted input function (as described in 27.7.1.3, paragraph 1), except that the function first clears `eofbit`.

-- The first overload of `seekg` ends with the sentence:

In case of failure, the function calls `setstate(failbit)` (which may throw `ios_basefailure`).

This sentence is missing in the description of the second overload, for no good reason that I can detect. The fix is to add the sentence:

CHANGE 27.7.1.3/43 (`seekg`) by ADDING at the end of the paragraph:

In case of failure, the function calls `setstate(failbit)` (which may throw `ios_basefailure`).

-- The `basic_ostream` seek functions, `seekp` and `tellp` say nothing about constructing a sentry object. But these functions need such protection just as much as `seekg` and `tellg`, which do. The fix is to add a sentence:

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

AFTER 27.7.2.5 basic\_ostream seek members, ADD:

Each seek member function begins execution by constructing an object of class sentry. It returns by destroying the sentry object.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 138**

**27.7 [iostream.format] Move and swap for I/O streams**

**Section: 27.7 [iostream.format]**

**Submitter:** Pablo Halpern

**Discussion:**

For `basic_istream`, `basic_ostream`, and `basic_iostream`, the move constructor does not do a move construction, the move-assignment operation does not do move-assignment and swap does not perform a swap. Moreover, these functions are protected, precluding their use in reasonable code.

The resolution to issue 900 (and related issue 911) assumes that these functions would never be called from client code. However, these classes are not abstract. They can be instantiated and there are use-cases for such instances. For example, one can create a `filebuf` outside of an `fstream`, then associate it with an `ostream`:

```
filebuf fb("name");  
  
ostream fstr(&fb);
```

The above `ostream` is a full-fledged object that should be movable, and copyable. In that case, the move and copy operations should move and copy the *whole* object, including the `rdbuf()` member.

However, moving the `rdbuf()` member poses a problem for derived classes like `fstream`, that contain embedded `streambuf` objects and which want to ensure that the base class portion of the copy container a pointer to a copy of the `streambuf`, not a pointer to the original `streambuf`. To simplify programming of these operations, we can have new constructors and a new function that perform the same actions as copy-construction, move-construction, and move-assignment, but do not move or copy the `rdbuf()` pointer.

**Proposed Resolution:**

The `move()` functions declared in `basic_ios` provides a good start. `move()` functions could be added to `basic_istream`, `basic_ostream`, and `basic_iostream` in order to support the same idiom. Alternatively, each class can have special constructors and assignment operators specific to the three operations. For example, `basic_istream` could have:

```
basic_istream(basic_istream&& rhs, basic_streambuf<...> *sb);  
  
partial_move(basic_istream&& rhs, basic_streambuf<...> *sb);  
  
partial_swap(basic_istream& other, basic_streambuf<...> *this_sb,  
             basic_streambuf<...> *other_sb);
```

These functions are useful in general, and should be public. Additionally, the regular move constructor, move-assignment operator, and swap functions should be made public and changed so that they move and swap the `streambuf` pointer as well as the rest of the stream state.

US 181

Clauses 30.2.4, 30, 20.10

# C++ Timeout Specification

ISO/IEC JTC1 SC22 WG21 - 2010-05-19 - National Body Comment by Google

Lawrence Crowl, crowl@google.com, Lawrence@Crowl.org

[Problem](#)

[Solution](#)

[Wording](#)

[20.10 Time utilities \[time\]](#)

[20.10.1 Clock requirements \[time.clock.req\]](#)

[20.10.5.1 Class `system\_clock` \[time.clock.system\]](#)

[20.10.5.2 Class `monotonic\_clock` \[time.clock.monotonic\]](#)

[20.10.5.3 Class `high\_resolution\_clock` \[time.clock.hires\]](#)

[20.10.5.4 Class `steady\_clock` \[time.clock.steady\]](#)

[30.2.4 Timing specifications \[thread.req.timing\]](#)

[30.3.2 Namespace `this\_thread` \[thread.thread.this\]](#)

[30.4.2 TimedMutex requirements \[thread.timedmutex.requirements\]](#)

[30.5.1 Class `condition\_variable` \[thread.condition.condvar\]](#)

[30.5.2 Class `condition\_variable\_any` \[thread.condition.condvarany\]](#)

[30.6.6 Class template `future` \[futures.unique\\_future\]](#)

[30.6.7 Class template `shared\_future` \[futures.shared\\_future\]](#)

[30.6.8 Class template `atomic\_future` \[futures.atomic\\_future\]](#)

## Problem

The meaning of clocks and timeouts is poorly defined within the Final Committee Draft when those clocks may be adjusted. Clocks can be adjusted by hours and many network protocols expect responses within seconds. A task regularly scheduled for midnight GMT should execute at midnight even though the clock has been adjusted to eliminate accumulated error. Failure of the standard to be precise about this distinction makes programs effectively unportable.

The root of the problem is that the current definition leaves open disparate implementations that even as implementations increase their quality, separate implementations will not converge on the same behavior.

There will necessarily be some delay in the interrupt response, function return, and scheduling of a thread waking from a timeout. Implementations can reasonably strive to approach zero delay for these activities. So, we call this delay the "quality of implementation".

Separately, there will be some delay due to contention for processor and memory resources. This delay is more under the control of the application programmer and systems administrator than it is under the implementation. So, we call this delay the "quality of management". The tradeoff between resources and responsiveness is necessarily application-dependent.

We can express the actual time of a timeout as the sum of the intended time, the quality of implementation and the quality of management. The remaining problem is to map the given timeout specifications to a common intended time.

If there are no adjustments to the clock time, the intended time may be straightforwardly determined from the manner of specification. In this case, we assume that any difference in durations between reported times and SI units is small, and thus constitutes a measure of the quality of implementation.

The problem arises with the specification of the intended timeout when the clock is adjusted in the middle of the timeout. There are two plausible strategies, the timeout is sensitive to adjustments in the clock time, or it is not. The *timeout\_until* functions have straightforward definitions when they are sensitive to adjustments. The *timeout\_for* functions have straightforward definitions when they are insensitive to adjustments.

## Solution

Define *timeout\_until* to respect reported clock time points and define *timeout\_for* to respect real time durations.

A consequence of these definitions is that *timeout\_until* and *timeout\_for* are not functionally redundant. That is, *timeout\_until*(Clock::now()+3\_seconds) is not equivalent to *timeout\_for*(3\_seconds) when the clock is adjusted in the interval.

The implementation of the *timeout* definition necessarily depends on a *steady clock*, one that cannot be adjusted. A monotonic clock is not sufficient. While one could be implicit in the standard, below we make one explicit.

Given a steady clock, the monotonic clock seems to be of marginal utility. Still, we have preserved it in this proposal.

[Note that considered the practice of setting clocks forward in unit tests, and believe that continuing to do so even in *timeout\_for* operations would be reasonable because the tests operate in a virtual world, not in the real world. The definition of time in that world need not be the same as the time in the real world.]

## Wording

The proposed wording changes are relative to the Final Committee Draft, [N3092](#).

### 20.10 Time utilities [time]

Edit within the header synopsis as follows.

```
// Clocks
class system_clock;
class monotonic_clock;
class steady_clock;
class high_resolution_clock;
```

#### 20.10.1 Clock requirements [time.clock.req]

Edit within table 54 as follows.

|                               |                             |  |
|-------------------------------|-----------------------------|--|
| ....                          | ....                        | ....   |
| <code>C1::is_monotonic</code> | <code>const bool</code>     | true if <code>t1 &lt;= t2</code> is always true, otherwise false. [ <i>Note</i> : A clock that can be adjusted backwards is not monotonic. — <i>end note</i> ] |
| <code>C1::is_steady</code>    | <code>const bool</code>     | true if the clock value cannot be adjusted and that the duration between ticks is close to the tick period, otherwise false.                                   |
| <code>C1::now()</code>        | <code>C1::time_point</code> | Returns a <code>time_point</code> object representing the current point in time.   |

Add a new paragraph 3.

[Note: The relative difference in durations between those reported by the given clock and the SI definition is a measure of the quality of implementation. —end note]

#### 20.10.5.1 Class `system_clock` [time.clock.system]

Edit paragraph 1 as follows.

Objects of class `system_clock` represent wall clock time from the system-wide realtime clock.

```
class system_clock {
```

```

public:
    typedef see below
rep;
    typedef ratio<unspecified, unspecified>
period;
    typedef chrono::duration<rep, period>
duration;
    typedef chrono::time_point<system_clock>
time_point;
    static const bool is_monotonic =
unspecified;
    static const bool is_steady =
unspecified;

    static time_point now();

    // Map to C API
    static time_t      to_time_t(const
time_point& t);
    static time_point  from_time_t(time_t t);
};

```

### 20.10.5.2 Class `monotonic_clock` [`time.clock.monotonic`]

Edit paragraph 2 as follows.

~~The class `monotonic_clock` is conditionally supported.~~

```

class monotonic_clock {
public:
    typedef unspecified
rep;
    typedef ratio<unspecified, unspecified>
period;
    typedef chrono::duration<rep, period>
duration;
    typedef chrono::time_point<unspecified>
time_point;
    typedef chrono::time_point<monotonic clock>
time_point;
    static const bool is_monotonic =
true;
    static const bool is_steady =
unspecified;

    static time_point now();
};

```

### 20.10.5.3 Class `high_resolution_clock` [`time.clock.hires`]

Edit paragraph 1 as follows.

Objects of class `high_resolution_clock` represent clocks with the shortest tick period. `high_resolution_clock` may be a synonym for `system_clock` or `monotonic_clock`.

```
class high_resolution_clock {
public:
    typedef unspecified
rep;
    typedef ratio<unspecified, unspecified>
period;
    typedef chrono::duration<rep, period>
duration;
    typedef chrono::time_point<unspecified>
time_point;
    typedef
chrono::time_point<high_resolution_clock>
time_point;
    static const bool is_monotonic =
unspecified;
    static const bool is_steady =
unspecified;

    static time_point now();
};
```

#### 20.10.5.4 Class `steady_clock` [time.clock.steady]

Add a new section 20.10.5.4 Class `steady_clock` [time.clock.steady] as follows.

Add a new paragraph.

Objects of class `steady_clock` represent clocks for which values of `time_point` advance at a steady rate relative to real time. That is, the clock may not be adjusted.

```
class steady_clock {
public:
    typedef unspecified
rep;
    typedef ratio<unspecified, unspecified>
period;
    typedef chrono::duration<rep, period>
duration;
    typedef chrono::time_point<steady_clock>
time_point;
    static const bool is_monotonic =
true;
    static const bool is_steady =
true;

    static time_point now();
```



};

### 30.2.4 Timing specifications [thread.req.timing]

Add a new paragraph after paragraph 1 as follows.

Implementations necessarily have some delay in returning from a timeout. Any overhead in interrupt response, function return, and scheduling induces a "quality of implementation" delay, expressed as duration  $D_i$ . Ideally, this delay would be zero. Further, any contention for processor and memory resources induces a "quality of management" delay, expressed as duration  $D_m$ . The delay durations may vary from timeout to timeout, but in all cases shorter is better.

Edit paragraph 2 as follows.

The member functions whose names end in `_for` take an argument that specifies a relative time duration. These functions produce *relative timeouts*. Implementations ~~should~~ shall use a monotonic steady clock to measure time for these functions. [~~Note: Implementations are not required to use a monotonic clock because such a clock may not be available. —end note~~][Footnote: All implementations for which standard time units are meaningful must necessarily have a steady clock within their hardware implementation. —end footnote] Given a duration argument  $D_t$ , the real-time duration of the timeout is  $D_t + D_i + D_m$ .

Add a new paragraph after paragraph 2 as follows.

The member functions whose names end in `_until` take an argument that specifies a time point. These functions produce *absolute timeouts*. Implementations shall use the clock specified in the time point to measure time for these functions. Given a clock time point argument  $C_t$ , the clock time point of the return from timeout is  $C_t + D_i + D_m$  unless, during the timeout, the clock has been adjusted to a time  $C_a > C_t$ , in which case the return time is  $C_a + D_i + D_m$ .

### 30.3.2 Namespace `this_thread` [thread.thread.this]

Edit paragraph 6, regarding `sleep_until`, as follows.

*Effects:* Blocks the calling thread ~~at least until the time for the absolute~~ timeout (30.2.4 [thread.req.timing]) specified by `abs_time`.

Edit paragraph 9, regarding `sleep_for`, as follows.

*Effects:* Blocks the calling thread ~~for at least the time~~ for the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`.

### 30.4.2 TimedMutex requirements [thread.timedmutex.requirements]

Edit paragraph 4, regarding `try_lock_for`, as follows.

*Effects:* The function attempts to obtain ownership of the mutex within the ~~time~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`. If the time specified by `rel_time` is less than or equal to 0, the function attempts to obtain ownership without blocking (as if by calling `try_lock()`). The function shall return within the ~~time~~ timeout specified by `rel_time` only if it has obtained ownership of the mutex object. [*Note:* As with `try_lock()`, there is no guarantee that ownership will be obtained if the lock is available, but implementations are expected to make a strong effort to do so. —*end note*]

Edit paragraph 10, regarding `try_lock_until`, as follows.

*Effects:* The function attempts to obtain ownership of the mutex by the ~~time~~ absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time`. If `abs_time` has already passed, the function attempts to obtain ownership without blocking (as if by calling `try_lock()`). The function shall return before the ~~time~~ timeout specified by `abs_time` only if it has obtained ownership of the mutex object. [*Note:* As with `try_lock()`, there is no guarantee that ownership will be obtained if the lock is available, but implementations are expected to make a strong effort to do so. —*end note*]

### 30.5.1 Class condition\_variable [thread.condition.condvar]

Edit within paragraph 19, regarding `wait_until`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` ~~or~~ a call to `notify_all()`, if `abs_time`  $\leq$  `clock::now()` expiration of the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time`, or spuriously.
- ....

Edit paragraph 21, regarding `wait_until`, as follows.

*Returns:* `cv_status::timeout` if ~~the function unblocked because `abs_time` was reached~~ the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time` expired, otherwise `cv_status::no_timeout`.

Edit within paragraph 25, regarding `wait_for`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` or a call to `notify_all()`, ~~by the elapsed time `rel_time` passing~~ expiration of the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`, or spuriously.
- ....

Edit paragraph 26, regarding `wait_for`, as follows.

*Returns:* `cv_status::timeout` if ~~the function unblocked because `rel_time` elapsed~~ the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` expired, otherwise `cv_status::no_timeout`.

Edit within paragraph 34, regarding the predicate `wait_for`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` or a call to `notify_all()`, ~~by the elapsed time `rel_time` passing~~ expiration of the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`, or spuriously.
- ....
- The loop terminates when `pred()` returns true or when the ~~time duration~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed.
- ....

### 30.5.2 Class `condition_variable_any` [thread.condition.condvarany]

Edit within paragraph 15, regarding `wait_until`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` or a call to `notify_all()`, ~~if `abs_time`  $\leq$~~

`clock::now()` expiration of the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time`, or spuriously.

- ....

Edit paragraph 17, regarding `wait_until`, as follows.

*Returns:* `cv_status::timeout` if ~~the function unblocked because `abs_time` was reached~~ the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time` expired, otherwise `cv_status::no_timeout`.

Edit within paragraph 20, regarding `wait_for`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` ~~or a call to `notify_all()`, by the elapsed time `rel_time` passing~~ expiration of the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`, or spuriously.
- ....

Edit paragraph 21, regarding `wait_for`, as follows.

*Returns:* `cv_status::timeout` if ~~the function unblocked because `rel_time` elapsed~~ the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` expired, otherwise `cv_status::no_timeout`.

Edit within paragraph 28, regarding the predicate `wait_for`, as follows.

*Effects:*

- ....
- The function will unblock when signaled by a call to `notify_one()` ~~or a call to `notify_all()`, by the elapsed time `rel_time` passing~~ expiration of the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time`, or spuriously.
- ....
- The loop terminates when `pred()` returns true or when the ~~time duration~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed.
- ....

### 30.6.6 Class template `future` [`futures.unique_future`]

Edit within paragraph 22, regarding `wait_for`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed expired.

Edit within paragraph 23, regarding `wait_for`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~time period~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed expired.
- ....

Edit within paragraph 25, regarding `wait_until`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until ~~the current time exceeds~~ the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time` has expired.

Edit within paragraph 26, regarding `wait_until`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~the time point~~ absolute timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has ~~been reached~~ expired.
- ....

### **30.6.7 Class template `shared_future` [futures.shared\_future]**

Edit within paragraph 27, regarding `wait_for`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed expired.

Edit within paragraph 28, regarding `wait_for`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~time period~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has elapsed expired.

- ....

Edit within paragraph 30, regarding `wait_until`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until ~~the current time exceeds~~ the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time` has expired.

Edit within paragraph 31, regarding `wait_until`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~the time point~~ absolute timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has been reached expired.
- ....

### 30.6.8 Class template `atomic_future` [futures.atomic\_future]

Edit within paragraph 23, regarding `wait_for`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until the relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has ~~elapsed~~ expired.

Edit within paragraph 24, regarding `wait_for`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~time period~~ relative timeout (30.2.4 [thread.req.timing]) specified by `rel_time` has ~~elapsed~~ expired.
- ....

Edit within paragraph 27, regarding `wait_until`, as follows.

*Effects:* blocks until the associated asynchronous state is ready or until ~~the current time exceeds~~ the absolute timeout (30.2.4 [thread.req.timing]) specified by `abs_time` has expired.

Edit within paragraph 28, regarding `wait_until`, as follows.

*Returns:*

- ....
- `future_status::timeout` if the function is returning because the ~~the time point~~ absolute timeout (30.2.4 [thread.req.timing]) specified by `rel_time` ~~has been reached~~ expired.
- ....

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 186**

**30.4.1 [thread.mutex.requirements/A] try\_lock does not provide a guarantee of forward progress**

**Section:** 30.4.1 [thread.mutex.requirements]

**Submitter:** Pablo Halpern (with help from Charles Leiserson, Edya Laden Mozes, and Tao B. Schardl)

**Discussion**

**Summary**

The current definition of `try_lock()` allows the attempt to fail spuriously, even if another thread has not acquired the lock. This definition does not provide a guarantee of forward progress without elaborate countermeasures by the programmer. We consider this to be a serious deficiency.

**Background**

The reasons to allow spurious failures are described in *Foundations of the C++ Concurrency Memory Model*, by Boehm and Adve, PLDI 2008: <http://portal.acm.org/citation.cfm?id=1375591> (requires ACM membership). The authors describe a situation like the following:

**THREAD 1**

```
x = 42;  
  
m.lock();
```

**THREAD 2**

```
while (m.try_lock()) { }  
  
assert(x == 42);
```

In order to maintain the “sequential consistency for data-race-free programs” model and still have the `assert` succeed, it is necessary to have an extra fence on entering `lock()` for some CPU architectures. Otherwise, the `try_lock()` may fail before the `lock()` has completed and, therefore, before the memory fence within `lock()` has executed. To avoid this cost, it was decided that this case should be allowed to fail. One way to allow such a failure is to allow spurious false return from `try_lock()`.

It is interesting to note that spurious failures are unlikely to occur in practice, on any architecture. POSIX mutexes, for example, are not permitted to fail spuriously. The effect of *allowing* spurious failures, however, is that the above code is permitted to fail, allowing some implementations to avoid an extra fence on `lock()`.

**Problems with the current definition**

The FCD's definition of `try_lock` does not guarantee that forward progress will be made by any thread using `try_lock`. Consider code that contains no `lock()` calls, but any number of `try_lock()` calls that could be executed in parallel. Whenever threads contend on a `try_lock()`, they may both fail and spin indefinitely, causing livelock. In fact, a degenerate but conforming implementation of `try_lock()` would simply return `false`, thus guaranteeing livelock! (Such an implementation of `try_lock()` would be very efficient because it contains no fences! :-).).

Livelock is a surprising result, and is one of those traps that is easy enough to fall into without the standard lending assistance. To defend against livelock, the programmer would be forced to implement elaborate logic, such as a probabilistic back-off protocol, which most programmers are not capable of writing correctly. In contrast, if `try_lock()` could not fail spuriously, then one thread would be guaranteed to make progress on



## FCD 14882 ADDITIONAL DETAILS TO US COMMENTS

the critical section. Code that did not have to defend against this kind of livelock would be simpler and less error-prone.

Allowing spurious failures is a premature optimization. People will have to write defensive code forever in the future, no matter how fast locks and memory fences become in the future (and the performance of these synchronization primitives has been improving). There is no doubt that lock acquisition can hurt performance, and that an extra fence could make a lock 20-30% slower, but if the nature of a program is such that the cost of acquiring an uncontended lock becomes a major bottleneck, then the solution is probably not a cheaper standard lock, but a cheaper algorithm – or else it is time to write your own lock with weaker guarantees. Our feeling is that the performance penalty for a well-behaved, predictable lock will make little difference to most multithreaded programs.

### **Proposed Resolution**

We propose to ban spurious failures of `try_lock` and to require sequential consistency in the system-provided mutex types.

To permit the user to supply higher-performance mutexes in the requirements clause, we can permit one or both of the following in the Mutex requirements:

1. Mutexes with spurious failures on `try_lock`
2. Mutexes whose locking functions permit reordering of ordinary memory operations into the critical region (i.e., after the point of synchronization with a failed `try_lock`).

Each approach has as its downsides. Spurious failures are reasonably easy to reason about, but very difficult to correct-for in cases where they might cause live-lock. Allowing memory operations to be reordered across a lock is not unheard-of, but it does violate the sequential consistency rules that the standard tries hard to preserve. Either approach will give implementers of custom mutexes the flexibility they need to reduce the cost of a lock, so it is not necessary to provide both options (and therefore suffer both sets of downsides).

As per the PLDI paper, allowing spurious failures is not an implementation strategy but rather a simple way to explain the behavior of certain locks without breaking sequential consistency. Effectively, an equivalence relationship is posited such that the simpler language of spurious failures can be used instead of the more complex happens-before language. However, this equivalence is not proven and, in fact, it appears not to hold in all cases. Specifically, spurious failures could result in livelock in situations where reordering memory operations would not. According to the PLDI paper, “this is the first solution that eliminates the need to provide a fence before a lock, while still maintaining a simple definition of a race.” Indeed, a `try_lock` that may fail spuriously is a new invention, not vetted in common use. It is not clear how many algorithms that depend on reliable `try_lock` behavior would be impacted, but it is clear that the number is not zero. To our knowledge, there is no widely-used threading library that allows spurious failures in `try_lock` as a matter of policy.

Perhaps the most important argument against allowing for spurious failures is that the work-around for code that requires a reliable `try_lock` is beyond the skill of most programmers, whereas the work-around for code that relies on sequential consistency simply involves adding a well-placed fence. If a `lock()` call allows memory operations to be reordered into the critical section, but disallows spurious `try_lock()` failures, then the code example above could be made to work reliably by adding a fence before the `lock()`:

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**THREAD 1**

```
x = 42;
```

```
atomic_thread_fence(memory_order_release);
```

```
m.lock();
```

**THREAD 2**

```
while (m.try_lock()) { }
```

```
assert(x == 42);
```

For these reasons, we propose wording that would **ban** spurious failure of `try_lock` for both system-supplied and user-supplied mutex types. For improved performance, a user-supplied mutex would be permitted to use a subtler synchronization protocol to avoid extra fences. Such a protocol might violate the "sequential consistency for data-race-free programs" model and would, therefore, be harder to reason about than mutexes that don't take such liberties. But such reasoning doesn't rely on incomplete equivalences which may yield incorrect conclusions. Critically, unlike spurious failures, there would not be the same threat of livelock. To simplify matters for the reader, however, we propose a non-normative note suggesting that such weak ordering of memory operations before a lock may manifest as an apparent spurious failure. It should be noted that the example code in the PLDI paper uses locks in a strange and not recommended fashion. Confusion about why it may fail should not be a driving force for choosing a sub-optimal solution.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

US 188

**30.4.1 [thread.mutex.requirements/B] Mutex requirements should not be bound to threads**

**Section:** 30.4.1 [thread.mutex.requirements]

**Submitter:** Pablo Halpern (with help from Charles Leiserson)

**Discussion:**

The set of requirements called “Mutex” are intended to describe mutual exclusion locks as a set of generic requirements that permit a reasonably wide range of different mutual exclusion mechanisms. However, these requirements unnecessarily narrow the range of locks that can be called Mutexes because they define mutex ownership in terms of threads. There are, however, other agents besides threads (e.g. tasks or processes) that could meaningfully hold a mutex, and the definition of Mutex should be broad enough to encompass those types of mutex.

Some types of mutex that would not be owned by a thread:

- Models involving dynamic scheduling of tasks may have a task migrate from thread to thread. A mutex may be acquired and released within a task, but may be rescheduled in between.
- There are models involving parallelism within locked regions. There could be a type of mutex that does not protect the threads within the parallel region against one another, but rather protects the entire parallel region against other concurrent actions.
- Parallel libraries like TBB (Threading Building Blocks) may need a type of mutex that can be inherited by a child task.
- Similarly, one could conceive of a mutex that would be inherited by a packaged task or `async()` call.
- There are parallelism models like Cilk for which thread ownership is not well defined.
- A coarse-grained mutex might be owned by a process, rather than with an individual thread.
- A library might allow the owner of a mutex to delegate the release of that mutex to another entity.

None of the above would meet the requirements of a Mutex by the current definition, yet programmers would probably be surprised if they were unable to use them with the standard `lock_guard`, `unique_lock`, and generic locking algorithms. Restricting a mutex to thread-based ownership is fine for the standard-supplied mutex types, but should not be required for all mutex types.

**Proposed resolution:**

Change [thread.mutex.requirements] as follows:

- 1 [A concurrent agent is an entity \(such as a thread, process, or packaged task\) that may perform work in parallel with other concurrent agents. The calling agent is determined by context, e.g., the calling thread, the calling process, or the packaged task that contains the call, etc..](#) A mutex object facilitates protection against data races and allows ~~thread~~-safe synchronization of data between ~~threads~~-concurrent agents. An agent ~~thread~~ owns a mutex from the time it successfully calls one of the lock functions until it calls unlock. Mutexes may be either recursive or non-recursive, and may grant simultaneous ownership to one or many ~~threads~~-concurrent agents. The mutex types supplied by the standard library provide exclusive ownership semantics [for threads](#): only one thread may own the mutex at a time. Both recursive and non-recursive mutexes are supplied. [\[ Note: Some mutexes are “agent-oblivious” in that they work for any concurrent-agent model because they do not determine or store the agent’s ID \(e.g., an ordinary spin-lock\). – end note \]](#)
- 2 This section describes requirements on template argument types used to instantiate templates defined in the C++ standard library. The template definitions in the C++ standard library refer to the named Mutex requirements whose details are set out below. In this description, `m` is an object of a Mutex type.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

- 3 A Mutex type shall be DefaultConstructible and Destructible. If initialization of an object of a Mutex type fails, an exception of type `std::system_error` shall be thrown. A Mutex type shall not be copyable nor movable.
- 4 The error conditions for error codes, if any, reported by member functions of a Mutex type shall be:
- `resource_unavailable_try_again` — if any native handle type manipulated is not available.
  - `operation_not_permitted` — if the `thread-calling agent` does not have the privilege to perform the operation.
  - `device_or_resource_busy` — if any native handle type manipulated is already locked.
  - `invalid_argument` — if any native handle type manipulated as part of mutex construction is incorrect.
- 5 The implementation shall provide lock and unlock operations, as described below. The implementation shall serialize those operations. [ *Note*: Construction and destruction of an object of a Mutex type need not be thread-safe; other synchronization should be used to ensure that Mutex objects are initialized and visible to other threads. —*end note* ]
- 6 The expression `m.lock()` shall be well-formed and have the following semantics:
- 7 *Effects*: Blocks the calling `thread-agent` until ownership of the mutex can be obtained for the calling `thread-agent`.
- 8 *Postcondition*: The calling `thread-agent` owns the mutex.
- 9 *Return type*: void
- 10 *Synchronization*: Prior `unlock()` operations on the same object shall synchronize with (1.10) this operation.
- 11 *Throws*: `std::system_error` when an exception is required (30.2.2).
- 12 *Error conditions*:
- `operation_not_permitted` — if the `thread-calling agent` does not have the privilege to perform the operation.
  - `resource_deadlock_would_occur` — if the implementation detects that a deadlock would occur.
  - `device_or_resource_busy` — if the mutex is already locked and blocking is not possible.
- 13 The expression `m.try_lock()` shall be well-formed and have the following semantics:
- 14 *Effects*: Attempts to obtain ownership of the mutex for the calling `thread-agent` without blocking. If ownership is not obtained, there is no effect and `try_lock()` immediately returns. An implementation may fail to obtain the lock even if it is not held by any other `thread concurrent agent`. [ *Note*: This spurious failure is normally uncommon, but allows interesting implementations based on a simple compare\_exchange\_weak (29). —*end note* ]
- 15 *Return type*: bool
- 16 *Returns*: true if ownership of the mutex was obtained for the calling `thread-agent`, otherwise false.
- 17 *Synchronization*: If `try_lock()` returns true, prior `unlock()` operations on the same object synchronize with (1.10) this operation. [ *Note*: Since `lock()` does not synchronize with a failed

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

subsequent `try_lock()`, the visibility rules are weak enough that little would be known about the state after a failure, even in the absence of spurious failures. —end note ]

18       *Throws:* Nothing.

19       The expression `m.unlock()` shall be well-formed and have the following semantics:

20       *Requires:* The calling `thread-agent` shall own the mutex.

21       *Effects:* Releases the calling `thread-agent`'s ownership of the mutex.

22       *Return type:* void

23       *Synchronization:* This operation synchronizes with (1.10) subsequent lock operations that obtain ownership on the same object.

24       *Throws:* Nothing.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 191**

**Comment 30.5: Condition variable wait\_for returning cv\_status is insufficient.**

**Technical details:**

The predicate wait is defined as a loop over the non-predicate wait. The predicate wait\_until is defined as a loop over the non-predicate wait\_until. The predicate wait\_for is not so implemented. The problem is that the return value from the non-predicate wait\_for operation provides insufficient information to synthesize the behavior of the predicate wait\_for. This problem indicates that the wait\_for facility is under-provisioned.

**Proposed resolution:**

Rather than return the enumeration cv\_status, the non-predicate wait\_for should return the duration remaining in the timeout. Then the predicate wait\_for has a constructive definition testing the duration against zero and waiting as needed for the successively smaller return values.

For consistency, the wait\_until operation could return the time point of the return for comparison to the specified abs\_time.

FCD 14882  
ADDITIONAL DETAILS TO US COMMENTS

**US 193**

**Comment 30.5.1: Condition variables preclude a wakeup optimization.**

**Technical details:**

When a thread calls `condition_variable::notify_all`, it causes all threads blocked in `condition_variable::wait(Predicate)` to wake up, reacquire the lock, evaluate their predicates, and then often go back to sleep. It's more efficient to evaluate the predicates in the notifying thread, and only wake up the threads for which the predicate is true. However, `wait(Predicate)` doesn't allow the predicate to run in another thread, and it doesn't allow missed wakeups when the predicate would return true when evaluated in the waiting thread but returns false when evaluated in the notifying thread. Further, the predicate must run under the wait's mutex, but `notify_all` may be called with or without that mutex held.

I believe `condition_variable_any` cannot implement this optimization because it doesn't require that all waiters hold the same mutex, so there's no way for `notify_all` to guarantee that it won't acquire a mutex recursively when trying to guard the predicate call.

**Proposed resolution:**

1. To `condition_variable::notify_one` and `condition_variable::notify_all`, add:

Requires: if any threads are concurrently executing a call to `wait*(lock, ...)`, `lock.mutex()` must be locked by the calling thread. Note: this allows them to run a wait's predicate.

2. To each of the `condition_variable::wait` variants that takes a predicate, add to the "Effects" a line saying that

"pred may be called in any thread that calls any notify or wait member function of this `condition_variable`, either inside the call to `notify()` or at the start of the next `mutex.unlock()` call, before the mutex is unlocked. If it returns false at any of these points, it is unspecified whether the preceding call to `notify*()` causes this wait to unblock."

US 194  
Clause 30.6

# Managing C++ Associated Asynchronous State

ISO/IEC JTC1 SC22 WG21 - 2010-05-19 - National Body Comment by Google

Lawrence Crowl, crowl@google.com, Lawrence@Crowl.org

[Problem](#)

[Solution](#)

[Wording](#)

[30.6.4 Associated asynchronous state \[futures.state\]](#)

[30.6.5 Class template promise \[futures.promise\]](#)

[30.6.6 Class template future \[futures.unique\\_future\]](#)

[30.6.7 Class template shared\\_future \[futures.shared\\_future\]](#)

[30.6.8 Class template atomic\\_future \[futures.atomic\\_future\]](#)

[30.6.10.1 packaged\\_task member functions \[futures.task.members\]](#)

## Problem

Within the Final Committee Draft, the specification for managing associated asynchronous state [futures.state] is confusing, sometimes omitted, and redundantly specified.

## Solution

Define terms-of-art for *releasing*, *making ready*, and *abandoning* an associated asynchronous state. Use those terms where appropriate.

## Wording

The wording is relative to the FCD.

**30.6.4 Associated asynchronous state [futures.state]**



Edit paragraph 5 as follows.

~~When the last reference to an associated asynchronous state is given up, any resources held by that associated asynchronous state are released. An asynchronous return object or an asynchronous provider *release their associated asynchronous state* as follows. If the return object or provider contains the last reference to that state, destroys that state. Destroys the reference to that state.~~

After paragraph 5, add a new paragraph as follows.

An asynchronous provider *makes ready an associated asynchronous state by marking that state ready and then unblocking any threads waiting for the associated state to become ready. An asynchronous provider *abandons an associated asynchronous state* as follows. If that state is not ready, the provider stores an exception object of type `future_error` with an error condition of `broken_promise` within that state and then makes ready that state. The provider then releases that state.*

### 30.6.5 Class template promise [futures.promise]

Edit paragraph 7, regarding the destructor, as follows.

~~*Effects:* if the associated asynchronous state of `*this` is not ready, stores an exception object of type `future_error` with an error condition of `broken_promise`. Any threads blocked in a function waiting for the asynchronous state associated with `*this` to become ready are unblocked. Destroys `*this` and releases its reference to its associated asynchronous state if any. If this is the last reference to that associated asynchronous state, destroys that state. abandons any associated asynchronous state ([futures.state]).~~

Edit paragraph 8, regarding the move assignment operator, as follows.

*Effects:* abandons any associated asynchronous state ([futures.state]) and then as if `promise<R>(std::move(rhs)).swap(*this)`.

Remove paragraph 9, as it is now redundant with the effects.

~~*Postcondition:* `rhs` has no associated asynchronous state. `*this` has the associated asynchronous state of `rhs` prior to the assignment.~~

Edit paragraph 18, regarding `set_value`, as follows.

*Effects:* atomically stores `r` in the associated asynchronous state and sets that state to ready. Any threads blocked in a call of a blocking function of

~~any future that refers to the same associated asynchronous state as `*this` are unblocked. makes ready that state ([futures.state]).~~

Edit paragraph 22, regarding `set_exception`, as follows.

*Effects:* atomically stores `p` in the associated asynchronous state and ~~sets that state to ready. Any threads blocked in a call of a blocking function of any future that refers to the same associated asynchronous state as `*this` are unblocked.~~ makes ready that state ([futures.state]).

Edit paragraph 26, regarding `set_value_at_thread_exit`, as follows.

*Effects:* Stores `r` in the associated asynchronous state without making ~~ready the associated asynchronous that state ready~~ immediately. Schedules ~~the associated asynchronous that state~~ to be made ready when the current thread exits, after all objects of thread storage duration associated with the current thread have been destroyed.

Edit paragraph 29, regarding `set_exception_at_thread_exit`, as follows.

*Effects:* Stores `p` in the associated asynchronous state without making ~~ready the associated asynchronous that state ready~~ immediately. Schedules ~~the associated asynchronous that state~~ to be made ready when the current thread exits, after all objects of thread storage duration associated with the current thread have been destroyed.

### 30.6.6 Class template future [futures.unique\_future]

Edit paragraph 10, regarding the destructor, as follows.

*Effects:*

- ~~gives up the reference to its~~ releases any associated asynchronous state ([futures.state]).
- destroys `*this`.

Edit paragraph 11, regarding the move assignment operator, as follows.

*Effects:*

- ~~if `*this` referred to an associated asynchronous state prior to the assignment it gives up this reference.~~ releases any associated asynchronous state ([futures.state]).
- move assigns the contents of `rhs` to `*this`.

### 30.6.7 Class template shared\_future [futures.shared\_future]

Edit paragraph 13, regarding the destructor, as follows.

*Effects:*

- ~~gives up the reference to its~~ releases any associated asynchronous state (`[futures.state]`).
- destroys `*this`.

Edit paragraph 14, regarding the move assignment operator, as follows.

*Effects:*

- ~~if `*this` refers to an associated asynchronous state it gives up this reference.~~ releases any associated asynchronous state (`[futures.state]`).
- assigns the contents of `rhs` to `*this`.

Edit paragraph 16, regarding the copy assignment operator, as follows.

*Effects:*

- ~~if `*this` refers to an associated asynchronous state it gives up this reference.~~ releases any associated asynchronous state (`[futures.state]`).
- assigns the contents of `rhs` to `*this`. [*Note:* as a result, `*this` refers to the same associated asynchronous state as `rhs` (if any). — *end note*]

### **30.6.8 Class template `atomic_future` [`futures.atomic_future`]**

Edit paragraph 9, regarding the destructor, as follows.

*Effects:*

- ~~gives up the reference to its~~ releases any associated asynchronous state (`[futures.state]`).
- destroys `*this`.

Edit paragraph 10, regarding the copy assignment operator, as follows.

*Effects:*

- releases any associated asynchronous state.
- assigns the contents of `rhs` to `*this`. [*Note:* as a result, `*this` refers to the same associated asynchronous state as `rhs` (if any). — *end note*]

### 30.6.10.1 packaged\_task member functions [futures.task.members]

Edit paragraph 9, regarding the move assignment operator, as follows.

*Effects:*

- releases any associated asynchronous state ([futures.state]).
- `packaged_task<R, ArgTypes...>(other).swap(*this).`

Edit paragraph 10, regarding the destructor, as follows.

*Effects:* ~~if the associated asynchronous state of \*this is not ready, stores an exception object of type future\_error with an error code of broken\_promise. Any threads blocked in a function waiting for the associated asynchronous state of \*this to become ready are unblocked. Destroys \*this and releases its reference to its associated asynchronous state (if any). If this is the last reference to that associated asynchronous state, destroys that state.~~ abandons any associated asynchronous state ([futures.state]).

Edit within paragraph 24, regarding `make_ready_at_thread_exit`, as follows.

.... this shall be done without ~~making the state ready~~ making ready that state ([futures.state]) immediately. ....

Edit paragraph 27, regarding `reset`, as follows.

*Effects:* ~~returns the object to a state as if a newly constructed instance had just been assigned to \*this by \*this = packaged\_task(std::move(f)), where f is the task stored in \*this.~~ [Note: this constructs a new associated asynchronous state for \*this. The old state is discarded, abandoned ([futures.state]), as described in the destructor for packaged\_task.get\_future may now be called again for \*this. —end note]

**FCD 14882**

**UNITED KINGDOM (GB)**

**ADDITIONAL DETAILS TO BALLOT COMMENTS**

**FCD 14882**

**GB 9**

B) The use of maximal in the definition of release sequence

-----  
The current wording of the standard suggests that release sequences are maximal with respect to sequence inclusion, i.e. that if there are two release operations in the modification order,

```
      mod      mod  
rel1----->rel2----->w
```

then [rel1;rel2;w] is the only release sequence, as the other candidate [rel2;w] is included in it. This interpretation precludes synchronizing with releases which have other releases sequenced-before them. We believe that the intention is actually to define the maximal release sequence from a particular release operation, which would admit both [rel1;rel2;w] and [rel2;w].

We suggest that 1.10:6 be changed to:


- 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, and every subsequent operation
- is performed by the same thread that performed the release, or
  - is an atomic read-modify-write operation.

**FCD 14882**  
**GB 10**

C) Inter-thread-happens-before is not acyclic

-----


Inter-thread-happens-before (1.10:10) is not acyclic at the moment. The following example execution is valid according to the current draft and has cycles in its inter-thread-happens-before relation, but we believe the intent is to forbid it:



The diagram consists of four memory actions, each represented by either an 'R' or a 'W' for read or write, followed by a letter identifying an atomic object and a memory ordering parameter. For example the action 'Rx\_consume' is a read of x that is ordered in memory as a consume.

The labelled arrows represent relationships between the operations. The vertical arrows labelled 'sb' are sequenced-before edges. The crossed diagonal arrows labelled 'rf' are reads-from edges that point from a write to the read that takes its value.

The reads-from edges make dependency-order edges with the same arrows. Transitively closing inter-thread-happens-before (ithb) gives the following relation:



Note that this is cyclic and that the execution is valid. We believe the intention is for inter thread happens before to be acyclic, and the standard will have to explicitly state this.

We suggest the following sentence be added to 1.10:

Candidate executions with a cyclic inter-thread-happens-before should not be considered [Note: the existence of such a candidate execution does not introduce undefined behaviour].

or

The inter-thread happens-before relation shall be acyclic.

with a similar note.

**FCD 14882**  
**GB 11**

E) Non-unique visible sequences of side effects and happens-before ordering

-----

In 1.10:12 the standard allows multiple visible sequences of side effects (vsse's) for a given read (despite the use of "The" at the start of 1.10:13). We will demonstrate this by constructing an execution with two vsse's. The following execution has five memory operations, four of which are read modify writes (RMW's). There are two threads, one with four operations each ordered by sequenced before (sb), the other with a single RMW release.

```
RMW1          +---RMW3_release
|
|sb          do/
v           /
R_consume<---+
|
|sb
v
RMW2
|
|sb
v
RMW4
```

The modification order in this example is as follows:

```
      mod      mod      mod
RMW1----->RMW2----->RMW3_release----->RMW4
```

There are two visible sequences of side effects here for the read consume:

[RMW1,RMW2] and [RMW3,RMW4]

The R\_consume here must read from the later vsse in modification order for the dependency\_ordered edge to exist. The existence of two vsse's relies on the lack of transitivity of happens before (which only occurs in the presence of consume operations).

Given that there is not a unique vsse for a given side effect, the standard should be changed from defining "the" vsse to defining "a" vsse. A note can explain that we can only ever read from one vsse.

In addition to non-uniqueness, if every element in a vsse happens-before a read, the read should not take the value of the visible side effect. We can prevent this by removing the word



"subsequent" from the definition of vsse in 1.10:12 (as suggested by Hans).

Collecting the two issues together, our suggestion is to change 1.10:12 to:

"A 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 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 a visible sequence of M with respect to B. [...] [ Note: For a given value computation B, there is only one visible sequence of side effects that B can read from, even if more than one sequence can be constructed. -end note ]"

FCD 14882  
GB 12

F) Alternative definition of the value read by an atomic operation  
-----


The standard introduces visible side effects, which are used first to define the values read by non-atomic operations. They are then re-used to constrain the value read by atomic operations:

1.10:13 says that an atomic operation must read from somewhere in "the" visible sequence of side effects, which must start from \*a\* visible side effect, ie a side effect that (a) happens before the read, and (b) is not happens-before-hidden.

We suspect that this re-use of the notion of visible side effect may be a drafting artifact, and tentatively suggest that it would be clearer to remove the requirement that there is a vse for atomics, replacing the first two sentences of 1.10:13 by

"An atomic operation must read from somewhere in the modification order that is not happens-before-hidden and does not follow (in modification order) any side effect that happens-after the read."

Note that the current text allows the following, whereas the revised text would forbid it. We believe that in a direct implementation on hardware, this would be forbidden by coherence.



**FCD 14882**

**GB 13**

G) Wording of the read-read coherence condition

-----  
In 1.10:13 a coherence condition is stated on the values of atomic reads:

"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."

The words "corresponds to" are not used elsewhere in the standard, as far as we can see, and it is unclear whether they have a special meaning here. In addition taking the value of the read B from the value read by A seems unnecessarily indirect. B could take its value from X instead. We suggest that the prose 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."

**FCD 14882**

**GB 14**

I) Implementation of dependencies on Power/ARM

-----

Carries-a-dependency includes via-memory dependencies (from a write to a program-order-later read from the same location on the same processor), but the Power and ARM architectures do not guarantee that such dependencies are respected by all processors. Hence, we wonder how implementations are expected to guarantee dependency-ordered before.

**FCD 14882**  
**GB 15**

J) Control dependencies for atomics

-----

Given the examples of compilers interchanging data and control dependencies, and that control dependencies are 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.

**FCD 14882**  
**GB 16**

L) Minor editorial changes

-----

1.10:12 last note: "...as defined here..." should be "...as defined below...".

Date: 2010-06-28

## Problems with bitmask types in the library

The library defines the *bitmask types* (17.5.2.1.3) `ios_base::fmtflags`, `ios_base::iostate`, `ios_base::openmode`, `regex_constants::syntax_option_type` and `regex_constants::match_flag_type`. Each is defined as an unscoped enumeration without a fixed underlying type. Each has `operator&`, `operator|` and `operator~` overloaded.

The library also defines the *enumerated types* (17.5.2.1.2) `ios_base::seekdir` and `regex_constants::error_type`.

There are several problems related to these types:

- Bitmask type operators declared as members of `ios_base`,
- bitmask type operators defined for enumerated types,
- `operator^` is missing,
- `operator~` is defined incorrectly,
- the underlying type should be fixed,
- value of `goodbit` is not unspecified,
- bitmask types are inconsistent with the requirements and over-specified.

### Bitmask operators declared as members of `ios_base`

The bitwise operators for the `ios_base` bitmask types are declared in the body of `ios_base`, making them members of the class, but they should be non-member functions. This is largely an editorial issue, but is addressed by the proposed changes below.

### Bitmask type operators defined for enumerated types

As specified in 27.5.2p1 and 27.5.2.1.5p1, `ios_base::seekdir` is an enumerated type (17.5.2.1.2) not a bitmask type (17.5.2.1.3) and so should not have bitwise operators defined. The operators were added by the `constexpr` proposal, n2349, but should be removed.

The same applies to `regex_constants::error_type` (28.5.3) although 28.5.3p1 incorrectly uses the term "enumeration type" instead of "enumerated type".

### `operator^` is missing

n2349 proposed the following additional overload which is missing from the FCD:

```
constexpr fmtflags operator^(fmtflags lhs, fmtflags rhs)
{
```

```
    return fmtflags(int(lhs) ^ int(rhs));  
}
```

This operator would be needed for all the bitmask types if they were defined as enumeration types.

### **operator~ is defined incorrectly**

The FCD specifies:

```
constexpr fmtflags operator~(fmtflags f)  
{  
    return fmtflags(f);  
}
```

This seems to be an editorial error as n2349 proposed:

```
constexpr fmtflags operator~(fmtflags f)  
{  
    return fmtflags(~(f));  
}
```

This isn't right either, since it will lead to infinite recursion. I think the intended definition is:

```
constexpr fmtflags operator~(fmtflags f)  
{  
    return fmtflags(~int(f));  
}
```

### **The underlying type should be fixed**

Bitmask types need `operator~` in order to clear a value from the bitmask type, as show in 17.5.2.1.3 paragraph 4:

— To *clear* a value  $Y$  in an object  $X$  is to evaluate the expression  $X \&= \sim Y$ .

However the definition for `fmtflags` above does not have well-specified behaviour if the underlying type is smaller than `int`, because `~int(f)` is likely to produce a value outside the range of `fmtflags`.

The same problem is present in the example implementation in 17.5.2.1.3 because `int_type` is only stated to be capable of holding all the values of *bitmask* so there is no guarantee that `sizeof(int_type) == sizeof(bitmask)` and therefore no guarantee that the code given for `operator~` works as intended.

This could be solved by giving the enumeration a fixed underlying type which is the same type as used for the conversions used by the bitwise operators. An alternative solution would be to do all conversions to and from `std::underlying_type<bitmask>`.



**Value of `goodbit` is not unspecified.**

The synopsis in 27.5.2 includes:

```
goodbit = unspecified,
```

but 27.5.2.1.3 specifies that `goodbit` has the value zero.

**Bitmask types are inconsistent with the requirements and over-specified.**

In C++03 implementations were allowed to use any of the options in 17.5.2.1.3 for the bitmask types defined by the library and e.g. `std::ios::in` is an object of bitmask type with an implementation-defined value. The FCD specifies that bitmask types are implemented as enumerations and `std::ios::in` is an enumerator not an object. Not only is this inconsistent with the stated requirements in 17.5.2.1.3 but it means that existing implementations which define bitmask types as an integer type must change, and it breaks the following valid C++03 code:

```
#include <ios>
std::ios::openmode* p = &std::ios::in;
```

If all the library's bitmask types are required to be enumeration types then there is no reason for 17.5.2.1.3 to allow integer types or a `bitset`.

Existing C++03 implementations choose to implement bitmask types as integer types for efficiency, because enumerations with overloaded operators are less efficient without `constexpr`. Since the FCD defines the bitmask types as enumerations with overloaded `constexpr` operators, implementations cannot provide a common definition for C++03 and C++0x. This is of no benefit to users or implementations.

**Proposed Solution**

The over-specification of the library's bitmask types should be reverted. The example bitmask in clause 17 can be updated to use `constexpr` and an enumeration type with a fixed underlying type. `constexpr` can also be used for declaring the objects of bitmask type. Proposed wording for this change is given below.

An alternative solution would be to fix all the issues listed above except the over-specification, but this would still require existing implementations to change, has the chance of breaking valid C++03 programs, and leaves the definitions of bitmask types inconsistent with the requirements in Clause 17.

**Proposed Wording**

## Modifications to Bitmask types

Make the following changes to 17.5.2.1.3:

The bitmask type *bitmask* can be written:

```

// For exposition only.
// int_type is an integral type capable of
// representing all values of bitmask
enum bitmask : int_type {
    V0 = 1 << 0, V1 = 1 << 1, V2 = 1 << 2, V3 = 1 << 3, .....
};

static constexpr bitmask C0 (V0);
static constexpr bitmask C1 (V1);
static constexpr bitmask C2 (V2);
static constexpr bitmask C3 (V3);
.....

// For exposition only.
// int_type is an integral type capable of
// representing all values of bitmask
constexpr bitmask operator& (bitmask X , bitmask Y ) {
    return static_cast<bitmask >(
        static_cast<int_type>(X ) &
        static_cast<int_type>(Y ));
}
constexpr bitmask operator| (bitmask X , bitmask Y ) {
    return static_cast<bitmask >(
        static_cast<int_type>(X ) |
        static_cast<int_type>(Y ));
}
constexpr bitmask operator^ (bitmask X , bitmask Y ){
    return static_cast<bitmask >(
        static_cast<int_type>(X ) ^
        static_cast<int_type>(Y ));
}
constexpr bitmask operator~ (bitmask X ){
    return static_cast<bitmask >(~static_cast<int_type>(X ));
}
constexpr bitmask & operator&=(bitmask & X , bitmask Y ){
    X = X &Y ; return X ;
}
constexpr bitmask & operator|=(bitmask & X , bitmask Y ) {
    X = X |Y ; return X ;
}
constexpr bitmask & operator^=(bitmask & X , bitmask Y ) {
    X = X ^Y ; return X ;
}

```

## Modifications to Class ios\_base

Make the following changes to the synopsis in 27.5.2:

```
class ios_base {  
public:  
    class failure;
```

```
    typedef T1 fmtflags;  
    constexpr fmtflags boolalpha = unspecified;  
    constexpr fmtflags dec = unspecified;  
    constexpr fmtflags fixed = unspecified;  
    constexpr fmtflags hex = unspecified;  
    constexpr fmtflags internal = unspecified;  
    constexpr fmtflags left = unspecified;  
    constexpr fmtflags oct = unspecified;  
    constexpr fmtflags right = unspecified;  
    constexpr fmtflags scientific = unspecified;  
    constexpr fmtflags showbase = unspecified;  
    constexpr fmtflags showpoint = unspecified;  
    constexpr fmtflags showpos = unspecified;  
    constexpr fmtflags skipws = unspecified;  
    constexpr fmtflags unitbuf = unspecified;  
    constexpr fmtflags uppercase = unspecified;  
    constexpr fmtflags adjustfield = unspecified;  
    constexpr fmtflags basefield = unspecified;  
    constexpr fmtflags floatfield = unspecified;
```

```
    // 27.5.2.1.2 fmtflags  
    enum fmtflags {  
        boolalpha = unspecified,  
        dec = unspecified,  
        fixed = unspecified,  
        hex = unspecified,  
        internal = unspecified,  
        left = unspecified,  
        oct = unspecified,  
        right = unspecified,  
        scientific = unspecified,  
        showbase = unspecified,  
        showpoint = unspecified,  
        showpos = unspecified,  
        skipws = unspecified,  
        unitbuf = unspecified,  
        uppercase = unspecified,  
        adjustfield = unspecified,  
        basefield = unspecified,  
        floatfield = unspecified,  
    };
```

```
    constexpr fmtflags operator~(fmtflags f);  
    constexpr fmtflags operator&(fmtflags lhs, fmtflags rhs);  
    constexpr fmtflags operator|(fmtflags lhs, fmtflags rhs);
```

```
    typedef T2 iostate;  
    constexpr iostate badbit = unspecified;
```

```
constexpr iostate eofbit = unspecified;
constexpr iostate failbit = unspecified;
constexpr iostate goodbit{ 0 };

// 27.5.2.1.3 iostate
enum iostate {
    badbit = unspecified,
    eofbit = unspecified,
    failbit = unspecified,
    goodbit = unspecified,
};

constexpr iostate operator~(iostate f);
constexpr iostate operator&(iostate lhs, iostate rhs);
constexpr iostate operator|(iostate lhs, iostate rhs);

typedef T3 openmode;
constexpr openmode app = unspecified;
constexpr openmode ate = unspecified;
constexpr openmode binary = unspecified;
constexpr openmode in = unspecified;
constexpr openmode out = unspecified;
constexpr openmode trunc = unspecified;

// 27.5.2.1.4 openmode
enum openmode {
    app = unspecified,
    ate = unspecified,
    binary = unspecified,
    in = unspecified,
    out = unspecified,
    trunc = unspecified,
};

constexpr openmode operator~(openmode f);
constexpr openmode operator&(openmode lhs, openmode rhs);
constexpr openmode operator|(openmode lhs, openmode rhs);

typedef T4 seekdir;
constexpr seekdir beg = unspecified;
constexpr seekdir cur = unspecified;
constexpr seekdir end = unspecified;

// 27.5.2.1.5 seekdir
enum seekdir {
    beg = unspecified,
    cur = unspecified,
    end = unspecified,
};

constexpr seekdir operator~(seekdir f);
constexpr seekdir operator&(seekdir lhs, seekdir rhs);
```

```
constexpr seekdir operator|(seekdir lhs, seekdir rhs);
```

Revert the changes from n2349, making the following changes to 27.5.2.1.2:

```
typedef T1 fmtflags;
enum fmtflags;
```

1 The type `fmtflags` is a bitmask type (17.5.2.1.3). Setting its elements has the effects indicated in Table 112.

2 Type `fmtflags` also defines the constants indicated in Table 113.

```
constexpr fmtflags ios_base::operator~(fmtflags f);
3 Returns: fmtflags(~f);
constexpr fmtflags ios_base::operator&(fmtflags lhs, fmtflags rhs);
4 Returns: fmtflags(int(lhs) & int(rhs));
constexpr fmtflags ios_base::operator|(fmtflags lhs, fmtflags rhs);
5 Returns: fmtflags(int(lhs) | int(rhs));
```

Make the following changes to 27.5.2.1.3:

```
typedef T2 iostate;
enum iostate;
```

1 The type `iostate` is a bitmask type (17.5.2.1.3) that contains the elements indicated in Table 114.

```
constexpr iostate ios_base::operator~(iostate f);
2 Returns: iostate(f);
constexpr iostate ios_base::operator&(iostate lhs, iostate rhs);
3 Returns: iostate(int(lhs) & int(rhs));
constexpr iostate ios_base::operator|(iostate lhs, iostate rhs);
4 Returns: iostate(int(lhs) | int(rhs));
```

Make the following changes to 27.5.2.1.4:

```
typedef T3 openmode;
enum openmode;
```

1 The type `openmode` is a bitmask type (17.5.2.1.3). It contains the elements indicated in Table 115.

```
constexpr openmode ios_base::operator~(openmode f);

2 Returns: openmode(f);

constexpr openmode ios_base::operator&(openmode lhs, openmode rhs);

3 Returns: openmode(int(lhs) & int(rhs));

constexpr openmode ios_base::operator|(openmode lhs, openmode rhs);

4 Returns: openmode(int(lhs) | int(rhs));
```

Make the following changes to 27.5.2.1.5:

```
typedef T4 seekdir;
enum seekdir;
```

1 The type `seekdir` is an enumerated type (17.5.2.1.2) that contains the elements indicated in Table 116.

```
constexpr seekdir ios_base::operator~(seekdir f);

2 Returns: seekdir(f);

constexpr seekdir ios_base::operator&(seekdir lhs, seekdir rhs);

3 Returns: seekdir(int(lhs) & int(rhs));

constexpr seekdir ios_base::operator|(seekdir lhs, seekdir rhs);

4 Returns: seekdir(int(lhs) | int(rhs));
```

## Modifications to Bitmask Type `regex_constants::syntax_option_type`

Make the following changes to 28.5.1:

```
namespace std {
  namespace regex_constants {
    typedef bitmask_type syntax_option_type;
```

```

constexpr syntax_option_type icase = unspecified;
constexpr syntax_option_type nosubs = unspecified;
constexpr syntax_option_type optimize = unspecified;
constexpr syntax_option_type collate = unspecified;
constexpr syntax_option_type ECMAScript = unspecified;
constexpr syntax_option_type basic = unspecified;
constexpr syntax_option_type extended = unspecified;
constexpr syntax_option_type awk = unspecified;
constexpr syntax_option_type grep = unspecified;
constexpr syntax_option_type egrep = unspecified;

enum syntax_option_type {
    icase = implementation defined,
    nosubs = implementation defined,
    optimize = implementation defined,
    collate = implementation defined,
    ECMAScript = implementation defined,
    basic = implementation defined,
    extended = implementation defined,
    awk = implementation defined,
    grep = implementation defined,
    egrep = implementation defined,
};
constexpr syntax_option_type operator~(syntax_option_type f);
constexpr syntax_option_type operator&(syntax_option_type lhs,
syntax_option_type rhs);
constexpr syntax_option_type operator|(syntax_option_type lhs,
syntax_option_type rhs);
}
}

```

1 The type `syntax_option_type` is an implementation-defined bitmask type (17.5.2.1.3). Setting its elements has the effects listed in table 128. A valid value of type `syntax_option_type` shall have exactly one of the elements `ECMAScript`, `basic`, `extended`, `awk`, `grep`, `egrep`, `set`.

```

constexpr syntax_option_type operator~(syntax_option_type f);
2 Returns: syntax_option_type(f);

constexpr syntax_option_type operator&(syntax_option_type lhs,
syntax_option_type rhs);
3 Returns: syntax_option_type(int(lhs) & int(rhs));

constexpr syntax_option_type operator|(syntax_option_type lhs,
syntax_option_type rhs);
4 Returns: syntax_option_type(int(lhs) | int(rhs));

```

## Modifications to Bitmask Type `regex_constants::syntax_option_type`

Make the following changes to 28.5.2:

```

namespace std {
  namespace regex_constants{
    typedef bitmask_type match_flag_type;
    constexpr match_flag_type match_default{ 0 };
    constexpr match_flag_type match_not_bol = unspecified;
    constexpr match_flag_type match_not_eol = unspecified;
    constexpr match_flag_type match_not_bow = unspecified;
    constexpr match_flag_type match_not_eow = unspecified;
    constexpr match_flag_type match_any = unspecified;
    constexpr match_flag_type match_not_null = unspecified;
    constexpr match_flag_type match_continuous = unspecified;
    constexpr match_flag_type match_prev_avail = unspecified;
    constexpr match_flag_type format_default{ 0 };
    constexpr match_flag_type format_sed = unspecified;
    constexpr match_flag_type format_no_copy = unspecified;
    constexpr match_flag_type format_first_only = unspecified;

    enum match_flag_type {
      match_default = 0,
      match_not_bol = implementation_defined,
      match_not_eol = implementation_defined,
      match_not_bow = implementation_defined,
      match_not_eow = implementation_defined,
      match_any = implementation_defined,
      match_not_null = implementation_defined,
      match_continuous = implementation_defined,
      match_prev_avail = implementation_defined,
      format_default = 0,
      format_sed = implementation_defined,
      format_no_copy = implementation_defined,
      format_first_only = implementation_defined,
    };
    constexpr match_flag_type operator~(match_flag_type f);
    constexpr match_flag_type operator&(match_flag_type lhs,
match_flag_type rhs);
    constexpr match_flag_type operator|(match_flag_type lhs,
match_flag_type rhs);
  }
}

```

1 The type `regex_constants::match_flag_type` is an implementation-defined bitmask type (17.5.2.1.3). Matching a regular expression against a sequence of characters `[first, last)` proceeds according to the rules of the grammar specified for the regular expression object, modified according to the effects listed in table 129 for any bitmask elements set.

```

constexpr match_flag_type operator~(match_flag_type f);

```

2 **Returns:** `match_flag_type(f)`.



```
constexpr match_flag_type operator&(match_flag_type lhs,
match_flag_type rhs);
```

~~3 Returns: match\_flag\_type(int(lhs) & int(rhs)).~~

```
constexpr match_flag_type operator|(match_flag_type lhs,
match_flag_type rhs);
```

~~4 Returns: match\_flag\_type(int(lhs) | int(rhs)).~~

## Modifications to Bitmask Type `regex_constants::error_type`

Make the following changes to 28.5.3:

```
namespace std {
  namespace regex_constants {
    typedef implementation-defined error_type;
    constexpr error_type error_collate = unspecified;
    constexpr error_type error_ctype = unspecified;
    constexpr error_type error_escape = unspecified;
    constexpr error_type error_backref = unspecified;
    constexpr error_type error_brack = unspecified;
    constexpr error_type error_paren = unspecified;
    constexpr error_type error_brace = unspecified;
    constexpr error_type error_badbrace = unspecified;
    constexpr error_type error_range = unspecified;
    constexpr error_type error_space = unspecified;
    constexpr error_type error_badrepeat = unspecified;
    constexpr error_type error_complexity = unspecified;
    constexpr error_type error_stack = unspecified;

    enum error_type {
      error_collate = implementation defined,
      error_ctype = implementation defined,
      error_escape = implementation defined,
      error_backref = implementation defined,
      error_brack = implementation defined,
      error_paren = implementation defined,
      error_brace = implementation defined,
      error_badbrace = implementation defined,
      error_range = implementation defined,
      error_space = implementation defined,
      error_badrepeat = implementation defined,
      error_complexity = implementation defined,
      error_stack = implementation defined,
    };
    constexpr error_type operator~(error_type f);
    constexpr error_type operator&(error_type lhs, error_type rhs);
    constexpr error_type operator|(error_type lhs, error_type rhs);
  }
}
```

1 The type `error_type` is an implementation-defined enumeration type (17.5.2.1.2). Values of type `error_type` represent the error conditions described in table 130:

```
constexpr error_type operator~(error_type f);
```

*2 Returns:* `error_type(f)`.

```
constexpr error_type operator&(error_type lhs, error_type rhs);
```

*3 Returns:* `error_type(int(lhs) & int(rhs))`.

```
constexpr error_type operator|(error_type lhs, error_type rhs);
```

*4 Returns:* `error_type(int(lhs) | int(rhs))`.

**FCD 14882**  
**GB 80**

A) Reads of indeterminate value result in undefined behaviour  
-----

In 20.2.3, NullablePointer requirements [nullablepointer.requirements], the standard specifies the behaviour of programs that read indeterminate values:

... A default-initialized object of type P may have an indeterminate value. [ Note: Operations involving indeterminate values may cause undefined behaviour. end note ]

The note uses the word "may", but we believe the intention is that such reads will cause undefined behaviour, but implementations are not required to produce an error. We suggest changing the note to:

[ Note: Operations involving indeterminate values cause undefined behaviour. end note ]

**FCD 14882**

**GB 122**

D) Imposed happens-before edges are not made transitive

-----

At various points in the standard new edges are added to happens-before, for example 27.2.3:2 adds happens-before edges between writes and reads from a stream:

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 happens before b.

Happens-before is defined in 1.10:11 in a deliberate way that makes it not explicitly transitively closed. Adding edges to happens-before directly, as in 27.2.3:2, does not provide transitivity with sequenced-before or any other existing happens-before edge. This lack of transitivity seems to be unintentional. In order to achieve transitivity we suggest each edge be added to inter-thread-happens-before as a synchronises-with edge (as per conversation with Hans Boehm). In the standard, each use of the words "happens-before" should be replaced with the words "synchronizes-with" in the following sentences:

27.2.3:2  
30.3.1.2:6  
30.3.1.5:7  
30.6.4:7  
30.6.9:5  
30.6.10.1:23

**FCD 14882**  
**GB 131**

H) Overlapping evaluations are allowed  
-----

29.3:8 states:

"An atomic store shall only store a value that has been computed from constants and program input values by a finite sequence of program evaluations, such that each evaluation observes the values of variables as computed by the last prior assignment in the sequence."

... but 1.9.13 states:

"If A is not sequenced before B and B is not sequenced before A, then A and B are unsequenced. [ Note: The execution of unsequenced evaluations can overlap. -end note ]"

Overlapping executions can make it impossible to construct the sequence described in 29.3:8. We are not sure of the intention here and do not offer a suggestion for change, but note that 29.3:8 is the condition that prevents out-of-thin-air reads.

**FCD 14882**  
**GB 136**

K) Initialisation of atomics

-----

We believe the intent is that for any atomics there is a distinguished initialisation write, but that this need not happens-before all the other operations on that atomic - specifically so that the initialisation write might be non-atomic and hence give rise to a data race, and hence undefined behaviour, in examples such as this (from Hans):

```
atomic< atomic<int> * > p
f()
{ atomic<int>x;           |
  p.store(&x,mo_rlx);    |   W_na  x
                        |   W_rlx p=&x
}
```

(where na is nonatomic and rlx is relaxed). We suspect also that no other mixed atomic/nonatomic access to the same location is intended to be permitted. Either way, a note would probably help.

FCD 14882: GB 138

Date: 2010-04-20

Author: [Anthony Williams](#)  
Just Software Solutions Ltd

## Lockable requirements for C++0x

This paper provides a proposed resolution for [LWG issue 1268](#). The basic premise of that issue is that the "Mutex requirements" from the current working draft are worded as if they are requirements on all lockable types, including user-defined mutexes and instantiations of `unique_lock`. However, the requirements really only need apply to the standard mutex types such as `std::mutex`, and are too strong when applied to user-defined mutex types.

This paper therefore proposes to separate the existing requirements on the standard mutex types from the general requirements on all lockable types.

### Proposed wording

Add a new section to 30.2 [thread.req] after 30.2.4 [thread.req.timing] as follows:

#### 30.2.5 Requirements for Lockable types

The standard library templates `unique_lock` (30.4.3.2 [thread.lock.unique]), `lock_guard` (30.4.3.1 [thread.lock.guard]), `lock`, `try_lock` (30.4.4 [thread.lock.algorithm]) and `condition_variable_any` (30.5.2 [thread.condition.condvarany]) all operate on user-supplied lockable objects. Such an object must support the member functions specified for either the `BasicLockable` requirements, the `Lockable` requirements or the `TimedLockable` requirements as appropriate to acquire or release ownership of a lock by a given thread. [Note: the nature of any lock ownership and any synchronization it may entail are not part of these requirements. — end note]

##### 30.2.5.1 BasicLockable Requirements

In order for a type `L` to qualify as a `BasicLockable` type, the following expressions must be supported, with the specified semantics, where `m` denotes a value of type `L`:

The expression `m.lock()` shall be well-formed and have the following semantics:

Effects:

Block until a lock can be acquired for the current thread. If an exception is thrown then a lock shall not have been acquired for the current thread.

Return type:

void

The expression `m.unlock()` shall be well-formed and have the following semantics:

Effects:

Release a lock on `m` held by the current thread.

Return type:

void

Throws:

Nothing if the current thread holds a lock on `m`.

### 30.2.5.2 Lockable Requirements

In order for a type `L` to qualify as a `Lockable` type, it must meet the `BasicLockable` requirements. In addition, the following expressions must be supported, with the specified semantics, where `m` denotes a value of type `L`:

The expression `m.try_lock()` shall be well-formed and have the following semantics:

Effects:

Attempt to acquire a lock for the current thread without blocking. If an exception is thrown then a lock shall not have been acquired for the current thread.

Return type:

bool

Returns:

true if the lock was acquired, false otherwise.

### 30.2.5.3 TimedLockable Requirements

For a type `TL` to qualify as `TimedLockable` it must meet the `Lockable` requirements, and additionally the following expressions must be well-formed, with the specified semantics, where `m` is an instance of a type `TL`, `rel_time` denotes instantiation of `duration` (20.10.3 [time.duration]) and `abs_time` denotes an instantiation of `time_point` (20.10.4 [time.point])

The expression `m.try_lock_for(rel_time)` shall be well-formed and have the following semantics:

Effects:



Attempt to acquire a lock for the current thread within the specified time period. If an exception is thrown then a lock shall not have been acquired for the current thread.

Return type:

bool

Returns:

true if the lock was acquired, false otherwise.

The expression `m.try_lock_until(abs_time)` shall be well-formed and have the following semantics:

Effects:

Attempt to acquire a lock for the current thread before the specified point in time. If an exception is thrown then a lock shall not have been acquired for the current thread.

Return type:

bool

Returns:

true if the lock was acquired, false otherwise.

Replace 30.4.1 [thread.mutex.requirements] paragraph 2 with the following:

2 This section describes requirements on ~~template argument types used to instantiate templates defined in~~ the mutex types supplied by the C++ standard library. ~~The template definitions in the C++ standard library refer~~ These types shall conform to the named `Mutex` requirements whose details are set out below. In this description, `m` is an object of ~~a `Mutex` type~~ one of the standard library mutex types `std::mutex`, `std::recursive_mutex`, `std::timed_mutex` or `std::recursive_timed_mutex`.

Add the following paragraph after 30.4.1 [thread.mutex.requirements] paragraph 2:

A `Mutex` type shall conform to the `Lockable` requirements (30.2.5.2).

Replace 30.4.2 [thread.timedmutex.requirements] paragraph 1 with the following:

The C++ standard library `TimedMutex` types `std::timed_mutex` and `std::recursive_timed_mutex` ~~A `TimedMutex` type~~ shall meet the requirements for a `Mutex` type. In addition, ~~it they~~ shall meet the requirements set out ~~in this~~ Clause 30.4.2 below, where `rel_time` denotes an instantiation of `duration` (20.10.3 [time.duration]) and `abs_time` denotes an instantiation of `time_point` (20.10.4 [time.point]).

Add the following paragraph after 30.4.2 [thread.timedmutex.requirements] paragraph 1:

A `TimedMutex` type shall conform to the `TimedLockable` requirements (30.2.5.3).

Add the following paragraph following 30.4.3.1 [thread.lock.guard] paragraph 1:

The supplied `Mutex` type shall meet the `Lockable` requirements (30.2.5.2).

Add the following paragraph following 30.4.3.2 [thread.lock.unique] paragraph 1:

The supplied `Mutex` type shall meet the `Lockable` requirements (30.2.5.2). `unique_lock<Mutex>` meets the `Lockable` requirements. If `Mutex` meets the `TimedLockable` requirements (30.2.5.3) then `unique_lock<Mutex>` also meets the `TimedLockable` requirements.

Replace the use of "mutex" or "mutex object" with "lockable object" throughout clause 30.4.3.30.4.3 [thread.mutex.locks] paragraph 1:

1 A lock is an object that holds a reference to a `mutexlockable object` and may unlock the `mutexlockable object` during the lock's destruction (such as when leaving block scope). A thread of execution may use a lock to aid in managing `mutex` ownership of a `lockable object` in an exception safe manner. A lock is said to own a `mutexlockable object` if it is currently managing the ownership of that `mutexlockable object` for a thread of execution. A lock does not manage the lifetime of the `mutexlockable object` it references. [ Note: Locks are intended to ease the burden of unlocking the `mutexlockable object` under both normal and exceptional circumstances. — end note ]

30.4.3 [thread.lock] paragraph 2:

2 Some lock constructors take tag types which describe what should be done with the `mutexlockable` object during the lock's construction.

30.4.3.1 [thread.lock.guard] paragraph 1:

1 An object of type `lock_guard` controls the ownership of a `mutexlockable` object within a scope. A `lock_guard` object maintains ownership of a `mutexlockable` object throughout the `lock_guard` object's lifetime. The behavior of a program is undefined if the `mutexlockable object` referenced by `pm` does not exist for the entire lifetime (3.8) of the `lock_guard` object. `Mutex` shall meet the `Lockable` requirements (30.2.5.2).

30.4.3.2 [thread.lock.unique] paragraph 1:

1 An object of type `unique_lock` controls the ownership of a `mutexlockable object` within a scope. `Mutex` Ownership of the `lockable object` may be acquired at construction or after construction, and may be transferred, after acquisition, to another `unique_lock` object. Objects of type `unique_lock` are not copyable but are movable. The behavior of a program is undefined if the contained pointer `pm` is not null and the `mutex` pointed to by `pm` does not exist for the entire remaining

lifetime (3.8) of the `unique_lock` object. Mutex shall meet the Lockable requirements (30.2.5.2).

Add the following to the precondition of `unique_lock(mutex_type& m, const chrono::time_point<Clock, Duration>& abs_time)` in 30.4.3.2.1 [thread.lock.unique.cons] paragraph 18:

```
template <class Clock, class Duration>
    unique_lock(mutex_type& m, const chrono::time_point<Clock,
Duration>& abs_time);
```

*18 Requires:* If `mutex_type` is not a recursive mutex the calling thread does not own the mutex. The supplied `mutex_type` type shall meet the TimedLockable requirements (30.2.5.3).

Add the following to the precondition of `unique_lock(mutex_type& m, const chrono::duration<Rep, Period>& rel_time)` in 30.4.3.2.1 [thread.lock.unique.cons] paragraph 22

*22 Requires:* If `mutex_type` is not a recursive mutex the calling thread does not own the mutex. The supplied `mutex_type` type shall meet the TimedLockable requirements (30.2.5.3).

Add the following as a precondition of `bool try_lock_until(const chrono::time_point<Clock, Duration>& abs_time)` before 30.4.3.2.2 [thread.lock.unique.locking] paragraph 10

```
template <class Clock, class Duration>
    bool try_lock_until(const chrono::time_point<Clock, Duration>&
abs_time);
```

*Requires:* The supplied `mutex_type` type shall meet the TimedLockable requirements (30.2.5.3).

Add the following as a precondition of `bool try_lock_for(const chrono::duration<Rep, Period>& rel_time)` before 30.4.3.2.2 [thread.lock.unique.locking] paragraph 15

```
template <class Rep, class Period>
    bool try_lock_for(const chrono::duration<Rep, Period>&
rel_time);
```

*Requires:* The supplied `mutex_type` type shall meet the TimedLockable requirements (30.2.5.3).

Replace 30.4.4 [thread.lock.algorithm] p1 with the following:

```
template <class L1, class L2, class... L3> int try_lock(L1&, L2&,
L3&...);
```

1 *Requires:* Each template parameter type shall meet the ~~Mutex~~ Lockable requirements (30.2.5.2), ~~except that a call to try\_lock() may throw an exception.~~ [Note: The unique\_lock class template meets these requirements when suitably instantiated. — end note]

Replace 30.4.4 [thread.lock.algorithm] p4 with the following:

```
template <class L1, class L2, class... L3> void lock(L1&, L2&, L3&...);
```

4 *Requires:* Each template parameter type shall meet the ~~Mutex~~ Lockable requirements (30.2.5.2), ~~except that a call to try\_lock() may throw an exception.~~ [Note: The unique\_lock class template meets these requirements when suitably instantiated. — end note]

Replace 30.5.2 [thread.condition.condvarany] paragraph 1 with:

1 A Lock type shall meet the ~~requirements for a Mutex type, except that try\_lock is not required~~ BasicLockable requirements (30.2.5.1). [Note: All of the standard mutex types meet this requirement. — end note]

**FCD 14882**

**FINLAND (FI)**

**ADDITIONAL DETAILS TO BALLOT COMMENTS**

# Defaulting non-public special member functions on first declaration

**Date:** 2010-06-13

**Version:** FCD14882: FI-1, FI-2, FI-3

**Authors:** Ville Voutilainen <[ville.voutilainen@gmail.com](mailto:ville.voutilainen@gmail.com)>

## Abstract

The resolution of Core Issue 906 forbids defaulting a non-public special member function on its first declaration. I believe this resolution to be incorrect, and this document explains why. This document proposes that defaulting a non-public special member function on its first declaration is an important part of design vocabulary, and is an important facility for writing readable code. This document is to be considered as supplemental explanatory document for FCD NB comments FI-1, FI-2 and FI-3. Triviality concerns are left out of this document, focusing just on the avoidance of having to specify defaulting on the definition outside the class body. The examples illustrate cases where the special member functions are non-public, but the same arguments (readability, code that's easy to write) also apply to being able to default virtual or explicit special member functions on their first declaration.

I wish to thank Daniel Krüger for performing a sanity check for an earlier version of this paper, and Lawrence Crowl for reviewing those papers and providing feedback and suggestions for improvement.

## The problem explained

I think it's problematic that non-public special member functions can't be defaulted on their first declaration. The defaulting is something that people do when they want to mark a special

member function defaulted regardless of the access of said member function, and people want to do it in a simple and concise manner, without having to write boiler-plate code. Forbidding defaulting on first declaration seems to restrict the design vocabulary considerably.

I shall repeat the use cases provided on the mailing list. Let's first consider a case where a copy constructor is protected and defaulted:

```
struct B
{
protected:
    B(const B&) = default;
};

struct D : B
{
};
```

The use case for this defaulting would be that the user is attempting to forbid slicing copies from D to B. Having to write the defaulting outside the class is tedious:

```
struct B
{
protected:
    B(const B&);
};
B::B(const B&) = default;
```

When the design evolves, the user decides that instances of the base are harmful, so she modifies the base thus:

```
struct B
{
protected:
    B() = default;
    B(const B&) = default;
};
```

Having to write the defaulting outside is once again tedious:

```
struct B
{
protected:
    B();
    B(const B&);
};
B::B(const B&) = default;
B::B() = default;
```

For other subobjects besides base classes, let's consider the following case where private constructor and copy constructor are defaulted:

```
struct part
```

```

{
friend class aggregate;
private:
    part() = default;
    part(const part&) = default;
};

struct aggregate
{
    part x;
};

```

The defaulting is equally tedious as in previous examples. Templates make the situation worse:

```

template<class T> struct TX
{
    template<class U> struct TY
    {
        protected:
            TY();
    };
protected:
    TX();
};
template<class T> TX<T>::TX() = default;
template<class T>
template<class U>
TX<T>::TY<U>::TY() = default;

```

In comparison, if the defaulting is allowed on the first declaration, the example is arguably much more readable:

```

template<class T> struct TX
{
    template<class U> struct TY
    {
        protected:
            TY() = default;
    };
protected:
    TX() = default;
};

```

According to my surveys, users understand that defaulting possibly affects e.g. triviality. Regardless of that, they want to do the defaulting concisely, and they don't want to be forced to write it outside the class. It seems to be a relatively common case to mark a special member function protected or private, default it, and move on. Forcing the defaulting to be done outside the class definition is cumbersome.

I consider forbidding defaulting non-public special member functions on first declaration to be overkill for what it's seemingly trying to achieve. Forcing users to write code outside of class definitions for defaulting, when they are never forced to do that for any other reason, is difficult to explain, difficult to teach, and is going to lead to people having to write error-prone out-of-line



declarations. For these reasons, Finnish NB comments FI-1/FI-2/FI-3 propose removing the restrictions for defaulting special member functions on their first declaration.

# Triviality of non-public special member functions defaulted on first declaration

**Date:** 2010-06-17

**Version:** FCD14882: FI-4, FI-5

**Authors:** Ville Voutilainen <[ville.voutilainen@gmail.com](mailto:ville.voutilainen@gmail.com)>

## Abstract

The resolution of Core Issue 906 forbids defaulting a non-public special member function on its first declaration. I believe this resolution to be incorrect, and this document explains why. This document proposes that defaulting a non-public special member function on its first declaration should not be considered as something that affects triviality. Defaulting should retain triviality, or retain the lack of triviality, and people shouldn't expect defaulting to bring forth triviality. Furthermore, non-public special member functions defaulted on first declaration should not be considered user-provided. This document is to be considered as supplemental explanatory comment for FCD NB comments FI-4 and FI-5.

I wish to thank Daniel Krüger for performing a sanity check for this paper, and Lawrence Crowl for reviewing the paper and providing feedback and suggestions for improvement.

## Background

N2346 states that "A special member function is user-provided if it is user-declared and not explicitly defaulted on its first declaration." Further, N3000 states in 12.8 Copying class objects [class.copy] p6 that "A copy constructor for class X is trivial if it is not user-provided". Therefore, I would assume that if it's allowed to default a non-public copy constructor on its first declaration, a non-public copy constructor could be trivial. For constructors, 12.1 Constructors

[class.ctor] p5 states that "A default constructor is trivial if it is not user-provided". Therefore, I'd assume that if it's allowed to default a non-public constructor on its first declaration, a non-public default constructor could be trivial.

The current status quo is that in order to be defaulted on the first declaration, a special member function must be public. During the discussion of Core Issue 906 in Santa Cruz it was suggested by John Spicer that triviality should be made more explicit by forcing the user to do the defaulting after the first declaration, assuming that defaulting a non-public member function will remove triviality because such a function would be considered user-provided.

Jason Merrill and I were of the opinion that access should not be the deciding factor when allowing defaulting on the first declaration, but it should be decided on whether the defaulted special member function is considered user-provided. It seems to me that N2346 strongly suggests that a defaulted special member function should not be considered user-provided, and thus non-public member functions should not affect triviality if they are defaulted.

Jens Maurer thought in message 15490 that a base class with a protected copy constructor makes a derived class not trivially copyable, because it has a subobject that is not trivially copyable. Mike Miller asked for a use case, which was provided later in reflector discussions, in message 15492 on the core reflector.

## **The problem, cases that should be allowed and trivial**

As a simple rule, if an implicitly declared, or explicitly defaulted special member function is (indirectly or directly) accessible, I'd expect triviality to follow. All cases where triviality would be lost would be cases where access is lost as well and offending code becomes ill-formed, so I don't see cases where triviality would be subtly lost. Cases that make classes non-trivial today (adding a non-trivial subobject) would still be non-trivial even with defaulted non-public special member functions.

I think it's problematic that non-public special member functions can't be defaulted on their first declaration. The defaulting is something that people do when they want to retain triviality if possible. The access control is orthogonal to that, because access control in this case is a design tool. Forbidding defaulting on first declaration seems to restrict the design vocabulary considerably.

Furthermore, regarding the concern by Jens Maurer, a base class with a protected copy constructor is not CopyConstructible to begin with; a derived class can be, however, if the derived class has a public copy constructor. My take on that is that a derived class can be trivially copyable, even if the base class alone would not be, as long as the copy constructor of the base is accessible to the derived class, and as long as neither copy constructor is user-provided.

This applies to all subobjects, a class aggregating a subobject having a private copy constructor can still be trivially copyable if the aggregating class is the friend of the aggregated class.

I shall repeat the use cases provided on the mailing list. Let's first consider a case where a copy constructor is protected and defaulted:

```
struct B
{
protected:
    B(const B&) = default;
};

struct D : B
{
};
```

The use case for this defaulting would be that the user is attempting to forbid slicing copies from D to B, but still retaining triviality.

When the design evolves, the user decides that instances of the base are harmful, so she modifies the base thus:

```
struct B
{
protected:
    B() = default;
    B(const B&) = default;
};
```

There's still no reason for the user to assume that D wouldn't be trivially copyable, because there's no user-provided function in sight. D should also be trivially constructible, for the same reason.

For other subobjects besides base classes, let's consider the following:

```
struct part
{
friend class aggregate;
private:
    part() = default;
    part(const part&) = default;
};

struct aggregate
{
    part x;
};
```

I'd expect it to be perfectly reasonable for the user to assume aggregate to be trivially copyable (and trivially constructible).

If we apply the previous rule, even the following should be trivial:

```
struct B
{
protected:
    B() = default;
    B(const B&) = default;
};

struct D : B
{
protected:
    D() = default;
    D(const D&) = default;
};

struct DD : D
{
};
```

The subobjects of DD have accessible (to DD) special member functions, which are defaulted. Further, the subobjects of D have accessible (to D) special member functions, which are defaulted. This should result in DD being both trivially constructible and trivially copyable. The rule seems generic, and implementations need to walk the subobject chains anyway to check for access violations and triviality, so I'd think it's feasible to take the defaulting into account when computing whether a given class is trivial. B and D aren't trivial in isolation, and I don't propose changing that. The crux of the matter is allowing examples like DD to be trivial. Previously this was practically impossible because the only trivial things were implicitly declared special member functions. When defaulting is added, it would be good to allow triviality for cases where user-declared (but not user-provided) special member functions can retain triviality.

## Some cases that should be allowed but aren't trivial

Let's consider the following example:

```
struct base
{
    std::shared_ptr<int> member;
};
```

The implicitly declared copy constructor will not be trivial, since `shared_ptr`'s copy constructor isn't. Now, if this is a base class that people don't want to have instances of, we may write

```
struct base
{
    std::shared_ptr<int> member;
protected:
    base() = default;
```

```
    base(const base&) = default;
};
```

This shouldn't affect the triviality in any way. Yet users expect it to be allowed. While destructors work just fine with an empty definition, copy constructors don't. Thus the brevity of being able to default on first declaration seems superior to the alternatives. Consider the following:

```
struct base
{
    std::shared_ptr<int> member;
protected:
    base(const base&) {} // doesn't work at all!
};
```

Another, correct attempt would be this:

```
struct base
{
    std::shared_ptr<int> member;
protected:
    base(const base&);
};

base::base(const base&) = default; // tedious to write outside the
class definition
```

## Conclusion

As suggested by this paper, I don't think a protected or even private special member function should result in loss of triviality, and most importantly, such a function should definitely not be considered user-provided just because it has different access than the implicitly declared function would have.

I'm inclined to think it's a quality-of-implementation issue to diagnose the lack of triviality, whether in the case of implicitly declared special member functions or in the case of special member functions defaulted on first declaration. I consider forbidding defaulting non-public special member functions on first declaration to be overkill for what it's seemingly trying to achieve, and I consider it grossly incorrect from the design vocabulary point of view.