

Defect Reports UK 002 to 025

BSI reference: IST/5/-/14 CP021.

WG14/N388
X3J11/94-073

These DRs were all accepted for forwarding to WG14 at the 1994-11-10 meeting of IST/5/-/14, who have not, however, examined the Suggested Technical Corrigenda.

--
Clive D.W. Feather | Santa Cruz Operation | If you lie to the compiler,
clive@sco.com | Croxley Centre | it will get its revenge.
Phone: +44 1923 813541 | Hatters Lane, Watford | - Henry Spencer
Fax: +44 1923 813811 | WD1 8YN, United Kingdom |

Each Defect Report should be treated as if preceeded by the following boilerplate:

** Submitted to BSI by Clive D.W. Feather <clive@sco.com>

** In this Defect Report, identifiers lexically identical to those declared
** in standard headers refer to the identifiers declared in those standard
** headers, whether or not the header is explicitly mentioned.

** This Defect Report has been prepared with considerable help from Mark
** Brader, Jutta Degener, Ronald Guilmette, and a person whose employment
** conditions require anonymity. However, opinions expressed or implied
** should not be assumed to be those of any person other than myself.

Defect Report UK 002: consistency of implementation-defined values

The restrictions that apply to "implementation-defined" entities are not clear.

What restrictions apply to implementation-defined entities ? If the value of an expression is implementation-defined, need the implementation always produce the same result ?

For example, the value of the expressions "7/-3" and "8/-3" must each be either -3 or -2. Can an implementation make them different (that is, use a different implementation-defined choice for each), or must it make the same choice for all integral divisions involving a negative quantity ?

As another example, can the number of significant characters and the significance of case in an identifier with external linkage depend on the identifier itself, or must it be the same for all possible identifiers ?

Defect Report UK 003: zero sized allocations

The use of the word "unique" in subclause 7.10.3 is ambiguous, and the handling of zero size allocations is incomplete.

Part 1

7.10.3 reads:

|| If the size of the space requested is zero, the behavior is
|| implementation-defined; the value returned shall be either a null pointer
|| or a unique pointer.

Does the term "unique" mean "different every time", or does it mean "there is a single pointer returned by all calls with size zero" (as might be presumed from the ordinary dictionary definition of "unique") ?

In other words, if "malloc(0)" does not return a null pointer, is the following expression:

malloc(0) == malloc(0)

003

always zero, always non-zero, or implementation-defined ?

Part 2

If "unique" means "there is a single pointer", what is the result of attempting to free that pointer ? How does the wording of 7.10.3 apply: || The value of a pointer that refers to freed space is indeterminate. Possibly nothing happens, because the pointer does not really point to a block of memory. In that case, is the following code strictly conforming ?

```
#include <stdlib.h>
/* ... */
void *p = malloc (0);
if (p != NULL)
{
    free (p); /* Line A */
    free (p); /* Line B */
}
```

What is the behavior if each of lines A and B are reached ?

Part 3

If "unique" means "different every time", then each such call still consumes address space, even though no storage actually needs to be allocated, and therefore the call can fail due to exhaustion of memory. Thus malloc (0) can return a null pointer, while the Standard seems to suggest that an implementation can return either null pointers or unique pointers, but not both. This is a defect in the existing wording.

Suggested Technical Corrigendum

If "unique" means "there is a single pointer", then change the penultimate sentence of 7.10.3 from:

If the size of the space requested is zero, the behavior is implementation-defined; the value returned shall be either a null pointer or a unique pointer.

to:

If the size of the space requested is zero, the behavior is implementation-defined; the value returned shall be either a null pointer or a unique pointer. The values returned by two zero-length allocations shall compare equal. Freeing the value returned by a zero-length allocation shall have no effect. If that value is used as an operand of the unary * operator, or of a + or - operator except one whose other operand has integral type and value zero, the behavior is undefined.

If "unique" means "different every time", then change it to:

If the size of the space requested is zero, the behavior is implementation-defined; either a null pointer is always returned, or the behavior is as if the size were some unspecified non-zero value. In the latter case, if the returned pointer is not a null pointer and is used as an operand of the unary * operator, or of a + or - operator except one whose other operand has integral type and value zero, the behavior is undefined.

[See also Defect Report UK 006.]

Defect Report UK 004: closed streams

Calls to fsetpos with positions in closed and reopened streams are permitted, but should be undefined.

The definition of fsetpos (subclause 7.9.9.3) requires the fpos_t argument to have a value generated by a successful call to fgetpos on the same stream. However, it does not require the stream to refer to the same file. If the stream does not so refer, the effect should be explicitly undefined.

Suggested Technical Corrigendum

Line

145 In 7.9.9.3, change:
146 ... an earlier call to the fgetpos function on the same stream.
147 to:
148 ... an earlier call to the fgetpos function on the same stream; there
149 shall not have been an intervening call to the fclose or freopen
150 function with that stream.

151
152
153 Defect Report UK 005: legitimacy of type synonyms
154 =====

155 The Standard does not clearly indicate when the spelling of a type name is
156 or is not significant; in other words, when a type name may be replaced by
157 another type name representing the same type.

158
159 Part 1
160 -----

161 Subclause 6.5.4.3 reads in part:
162 || The special case of void as the only item in the list specifies that
163 || the function has no parameters.

164
165 Subclause 6.7.1 reads in part:
166 || (except in the special case of a parameter list consisting of a single
167 || parameter of type void, in which there shall not be an identifier).

168
169 In both cases, the word "void" is set in the typeface used to indicate C code.

170
171 In the code:

```
172     typedef void Void;  
173  
174     extern int f (Void);  
175  
176     int f (Void) { return 0; }
```

177
178 is the declaration on line 2 strictly conforming, and is the external
179 definition on line 3 strictly conforming ?

180
181
182 Part 2
183 -----

184 Subclause 5.1.2.2.1 reads in part:
185 || It can be defined with no parameters:
186 || int main (void) { /* ... */ }

187
188 Is the following definition of main strictly conforming ?

```
189     typedef int word;  
190  
191     word main (void) { /* ... */ }
```

192
193 Part 3
194 -----

195
196 Are there any circumstances in which a typedef name is not permitted instead
197 of the type it is a synonym for ? If so, what are they ?

198
199 Defect Report UK 006: null pointer conversions
200 =====

201 The Standard does not define semantics for the explicit conversion of null
202 pointer constants and for the implicit conversion of null pointers.

203
204 Subclause 6.2.2.3 reads in part:
205 || If a null pointer constant is assigned to or compared for equality to
206 || a pointer, the constant is converted to a pointer of that type. Such a
207 || pointer, called a null pointer, is guaranteed to compare unequal to a
208 || pointer to any object or function.

209
210 ||
211 || Two null pointers, converted through possibly different sequences of
212 || casts to pointer types, shall compare equal.

213
214 Given the definitions:

```
215     void * p = 0;  
216
```

005

217 int * i = 0;
218
219 does the standard guarantee that the expression
220
221 p == i
222
223 always evaluates to 1 ? The last quoted sentence only covers casts, and
224 not the implicit conversions of that comparison. Conversely, do the
225 expressions:
226
227 (int *) 0
228 1 ? 0 : (int *) 0
229
230 yield null pointers of type (int *) ? The quoted text does not cover the
231 case of a null pointer constant being converted other than by assignment
232 or in a test for equality, yet expressions such as these are widely used.

233 Suggested Technical Corrigendum

234 -----
235 In subclause 6.2.2.3, change:

236 Two null pointers, converted through possibly different sequences of
237 casts to pointer types, shall compare equal.
238

239 to:

240 Conversion of a null pointer to another pointer type yields a null
241 pointer of that type. Any two null pointers shall compare equal.
242

243 Alternatively, a common term could be introduced to more conveniently
244 describe the various forms of pointer that cannot be dereferenced. In
245 this case, replace the last two paragraphs of subclause 6.2.2.3 with:

246 For each pointer type, there exist values which can participate in
247 assignment and equality operations, but which cause undefined behavior
248 if dereferenced. These are referred to as undereferenceable. An
249 undereferenceable pointer compares unequal to any other value of the
250 same pointer type. For each pointer type, one particular undereferenceable
251 pointer value is called the null pointer. [*]

252 [*] Since there is only one such value, all null pointers of the same
253 type compare equal.
254

255
256 An integral constant expression with the value 0, or such an expression
257 cast to type void *, is called a null pointer constant. If a null
258 pointer constant is assigned to or compared for equality with an object
259 of pointer type, or cast to pointer type, then it is converted to the
260 null pointer of that type. Conversion of a null pointer to another pointer
261 type produces the null pointer of that type.
262

263 If the answer to Defect Report UK 003 is that "unique" means "different
264 each time", then replace the last two sentences of subclause 7.10.3 with:

265 If the size of the space requested is zero, an undereferenceable pointer
266 is returned. It is implementation-defined whether this is always a null
267 pointer or whether the implementation attempts to produce a value distinct
268 from any other undereferenceable pointer. Any pointer value returned by
269 an allocation can be passed to the free function; if the value is not a
270 null pointer, it becomes indeterminate[*]. The value of a pointer that
271 refers to any part of a freed object is also indeterminate.
272

273 [*] A subsequent allocation may return a pointer value with the same
274 bit pattern, but a strictly conforming program can't detect this.
275

276 Defect Report UK 007: consistency of the Standard

277 =====
278 Defects exist in the way the Standard refers to itself.
279

280 Part 1

281 -----
282 The introduction to the Standard reads in part:

283 || The introduction, the examples, the footnotes, the references, and
284 || the annexes are not part of this International Standard.
285

286
287 While it is not, strictly speaking, an inconsistency for text that is not
288 part of the Standard to specify which text is part of the Standard, it is

006

Line

289 confusing for this to be the case when other text that *looks* like part
290 of the Standard isn't - the examples and footnotes.
291
292 In particular, placing this information - necessary for interpreting the
293 text of the Standard itself - outside that text causes a danger that, when
294 some other document is produced that purports to contain the full text of
295 the Standard, the Introduction will be omitted while the footnotes and
296 examples are retained. A reader of such a document who is not aware of the
297 text of the introduction will then be misled as to the Standard's contents.
298 Whilst this is not the responsibility of ISO, it is another reason for
299 regularising the situation.

300
301 Note that this has definitely happened in the case of "The Annotated ANSI
302 C Standard" by Herbert Schildt, and I have been informed (but have not
303 confirmed) that it has also happened with the version of the Standard
304 distributed by the Australian National Body.

305 Part 2

306 -----

307 The introduction to the Standard reads in part:

308 || The language clause (clause 7) ...

309 || The library clause (clause 8) ...

310
311

312 These references are wrong.

313 Suggested Technical Corrigendum

314 -----

315 In the introduction, change:

316 The introduction, the examples, the footnotes, the references, and
317 the annexes are not part of this International Standard.

318
319 The language clause (clause 7) ...

320
321 The library clause (clause 8) ...

322 to:

323 As specified in the definitions and conventions clause (clause 3),
324 this introduction, the examples, the footnotes, the references, and
325 the annexes are not part of this International Standard.

326
327 The language clause (clause 6) ...

328
329 The library clause (clause 7) ...

330
331 Insert at the start of clause 3:

332 The introduction, the examples, the footnotes, the references, and
333 the annexes are not part of this International Standard.

334 Defect Report UK 008: reservation of identifiers

335 =====
336 The Standard is unclear in its description of what applications can and
337 cannot do with identifiers that are reserved to the implementation for
338 certain uses.

339 Subclause 7.1.3 reads in part:

340 || Each identifier with file scope listed in any of the following subclauses
341 || (including the future library directions) is reserved for use as an
342 || identifier with file scope in the same name space if any of its
343 || associated headers is included.

344
345 Does this include reservation as macros ? In particular, is the following
346 code:

347
348 #include <stddef.h>
349 #define size_t 42

350
351 strictly conforming, or could it cause a redefinition of the macro "size_t" ?
352 Similarly, can another macro legitimately defined by <stddef.h> (such as
353 offsetof) include size_t in its replacement list, so that:

354
355 #include <stddef.h>
356 #undef size_t

007


```
Line
361 #define size_t 42
362 /* ... */
363 offsetof (struct_type, field)
364
```

fails to expand correctly ? It is not clear how the wording of footnote 91 applies, and this is in any case not part of the Standard (except in Australia :-).

Defect Report UK 009: details of reserved symbols

=====

The wording of subclause 7.13 is unclear.

Does the term "any combination" in 7.13 include the empty combination ? In other words, are names like E2, tom, LC_X, and memo reserved ?

Defect Report UK 010: gmtime and localtime

=====

The Standard's description of the static objects used by <time.h> functions is misleading.

Subclause 7.12.3 reads in part:

```
|| these functions return values in one of two static objects: a broken-down
|| time structure and an array of char. Execution of any of the functions
|| may overwrite the information returned in either of these objects by any
|| of the other functions.
```

Does this mean that, for example, localtime and gmtime must share a single broken-down time structure, and so the value returned from gmtime, if not a null pointer, must equal the value returned from localtime (and this value cannot change during execution of the program) ?

The wording "the other functions" also implies that a call to gmtime can overwrite a previous call to localtime, but not a previous call to gmtime. This is clearly ridiculous.

Suggested Technical Corrigendum

In subclause 7.12.3, change:

```
these functions return values in one of two static objects: a broken-down
time structure and an array of char. Execution of any of the functions
may overwrite the information returned in either of these objects by any
of the other functions.
```

to:

```
these functions each return a pointer to an object of static storage
duration after assigning a value to it. Execution of any of these
functions may overwrite the information returned in any of these objects
by a previous call to any of these functions.
```

Defect Report UK 011: undeclared identifiers

=====

The Standard is not clear on whether the use of an undeclared identifier as a primary expression requires a diagnostic message.

Subclause 6.3.1 states that:

```
|| An identifier is a primary expression, provided it has been declared as
|| designating an object (in which case it is an lvalue) or a function (in
|| which case it is a function designator).
```

It has been suggested that if no declaration of some identifier is visible in the current scope when that identifier appears in an expression, the identifier is not a primary expression, and therefore the syntax of 6.3.1 is violated (in other words, there is no valid parse for the expression). This would thus require a diagnostic for an undeclared identifier.

Is this interpretation correct ? If yes, then it needs to be made clear that this does not prevent a previously undeclared function from being called by a strictly conforming program (see 6.3.2.2).

If not, does an undeclared identifier require a diagnostic, and if so, why ?

If not, is this a deliberate policy, or is it a defect that needs correction ?

Line

Defect Report UK 012: bad declarations

The Standard contains no constraint to prevent declarations involving types not defined by subclause 6.1.2.5.

Subclause 6.5 states that:

|| A declaration shall declare at least a declarator, a tag, or the members
|| of an enumeration.

There seems to be no constraint that a declarator generate a well-formed type. Consider the following code:

```
{
    int a [][5];      /* Line A */
    int x, b [][5];   /* Line B */
}
```

Neither a nor b has a well formed type. Does line A nevertheless "declare a declarator", or does it violate the quoted constraint? If it violates the constraint, does line B?

Is it the intent of the Standard that an ill-formed (but syntactically correct) type generate a diagnostic? If so, then is there one, or does one need to be added?

Defect Report UK 013: tags and incomplete types

The wording of subclause 6.5.2.3 concerning tags is defective in a number of ways.

Part 1

The first paragraph states that:

If this declaration of the tag is visible, a subsequent declaration that uses the tag and that omits the bracketed list specifies the declared structure, union, or enumerated type.

This neither handles the case of a type name (for example, in the operand of the sizeof operator), nor does it make it clear whether or not the rule applies within the braces of the first declaration (the tag is in scope from the open brace).

In other words, it fails to address either occurrence of "struct tag" in the following code:

```
{
    struct tag { int i [sizeof (struct tag)]; };
    int j [sizeof (struct tag)];
    /* ... */
}
```

Part 2

The second paragraph does not adequately distinguish between type specifiers which refer to an incomplete type and those which refer to a type in an outer scope. For example, in the following code, it fails to indicate whether or not all the uses of the tag refer to the same type:

```
struct tag;
struct tag *p;
{
    struct tag *q;
    /* ... */
}
struct tag { int member; };
```

Part 3

The handling of enumerated types before their content is defined is also unclear; this was covered to some extent in DR013Q5 and the subsequent discussion on the WG14 mailing list.

009

For example, what is the status of the following code:

```
enum tag { e = sizeof (enum tag *****) };
```

or of:

```
enum tag { e0, e1, e2, e3 };
{
    enum tag2 { e4 = sizeof (enum tag); };
    enum tag { e5 = sizeof (enum tag); };
}
```

If an enumeration tag cannot be used before the end of the list defining its contents, a diagnostic ought to be required.

Part 4

If the same tag is used in a type specifier with a contents list twice in the same scope, it is unclear whether or not a diagnostic is required. It could be argued that, since this is forbidden by the semantics in 6.5.2.3, it is not excluded from the second constraint of 6.5, and so a diagnostic is required by that constraint. However, this may be viewed as clutching at straws. An explicit constraint should be added.

Suggested Technical Corrigendum

Rather than making piecemeal changes to address each issue separately, the whole subclause should be rewritten. Footnote numbers have been chosen to match the present footnotes.

Constraints

A specific type shall have its content defined at most once.

A type specifier of the form

enum identifier

without an enumerator list shall only appear when the type it specifies is complete.

Semantics

All declarations of structure, union, or enumerated types that have the same scope and use the same tag declare the same type. The type is incomplete [63] until the closing brace of the list defining the content, and complete thereafter.

[63] An incomplete type may only be used when the size of an object of that type is not needed. [Append the present wording, or see Defect Report CA-2-09 - submitted independently - for alternative wording.]

Two declarations of structure, union, or enumerated types which are in different scopes or use different tags declare distinct types. Each declaration of a structure, union, or enumerated type which does not include a tag declares a distinct type.

A type specifier of the form

struct-or-union identifier { struct-declaration-list }

opt

or

enum identifier { enumerator-list }

opt

declares a structure, union, or enumerated type. The list defines the *structure content*, *union content*, or *enumeration content*. If an identifier is provided[64], the type specifier also declares the identifier to be the tag of that type.

[64] If there is no identifier, the type can, within the translation unit, only be referred to by the declaration of which it is a part. Of course, when the declaration is of a typedef name, subsequent declarations can make use of that typedef name to declare objects having the specified structure, union, or enumerated type.

010

Line

577 A declaration of the form
 578 struct-or-union identifier ;
 579 specifies a structure or union type and declares the identifier as the tag
 580 of that type[62].
 581
 582 [62] A similar construction with enum does not exist.
 583
 584 If a type specifier of the form
 585 struct-or-union identifier
 586 occurs other than as part of one of the above constructions, and no
 587 other declaration of the identifier as a tag is visible, then it
 588 declares a structure or union type which is incomplete at this point,
 589 and declares the identifier as the tag of that type[62].
 590
 591 If a type specifier of the form
 592 struct-or-union identifier
 593 or
 594 enum identifier
 595 occurs other than as part of one of the above constructions, and a
 596 declaration of the identifier as a tag is visible, then it specifies
 597 the same type as that other declaration, and does not redeclare the tag.
 598
 599

Defect Report UK 014: meaning of lvalue

600 =====
 601 Constraints that require something to be an lvalue place an unacceptable
 602 burden on the implementation.
 603
 604

605 Subclause 6.2.2.1 states in part:

606 || An lvalue is an expression (with an object type or an incomplete type
 607 || other than void) that designates an object.
 608

609 Given the declaration "int a [10], i;", the expression "a [i]" designates
 610 an object, and is thus an lvalue, if and only if "i" has a value between 0
 611 and 9 inclusive (see Defect Report 076 for further details). Now consider
 612 the Constraint in subclause 6.3.3.2:

613 || The operand of the unary & operator shall be either a function
 614 || designator or an lvalue that designates an object ...
 615

616 This means that the expression "&a[i]" is a constraint violation whenever
 617 "i" has a value outside the range 0 to 9 inclusive, and that therefore a
 618 diagnostic is required, at run-time !
 619

620 The defect is that the operand of the unary & operator does not need to
 621 be an lvalue that designates an object, but rather an lvalue which, if
 622 evaluated with its operands having suitable values, could designate an object.
 623

624 There are probably other parts of the Standard with the same problem,
 625 such as 6.3.2.4, 6.3.3.1, and 6.3.16.
 626
 627

Defect Report UK 015: consistency of the Standard

628 =====
 629 The change to the n conversion specifier in subclause 7.9.6.2 made by TC1,
 630 DR014Q2, should also be applied to subclause 7.9.6.1. Change:

631 No argument is converted.

632 to:

633 No argument is converted, but one is consumed. If the conversion
 634 specification with this conversion specifier is not one of %n, %ln,
 635 or %hn, the behavior is undefined.
 636
 637

638 In addition, an entry something like:

639 A %n conversion specification for the fprintf or fscanf functions
 640 is not one of %n, %ln, or %hn (7.9.6.1, 7.9.6.2).
 641 should be added to Annex G.2.
 642
 643

Defect Report UK 016: consistency of the Standard

644 =====
 645 The change to subclause 6.3 made by TC1, DR053Q1, should also be applied
 646 in Annex G.2 (page 200).
 647
 648

649 Defect Report UK 017: Trigraphs

650 =====

651 The standard's description of the replacement of trigraphs is contradictory.

652 Subclause 5.2.1.1 reads in part:

653 || All occurrences in a source file of the following sequences of three
654 || characters (called trigraph sequences [7]) are replaced with the
655 || corresponding single character.
656 [...]

657 || Each ? that does not begin one of the trigraphs listed above is not
658 || changed.

659 Since the second character in each trigraph is a ? that does not begin
660 the trigraph, this is a direct contradiction.

661 Suggested Technical Corrigendum

662 -----

663 Change the last sentence of the cited text to:

664 Each ? that is not part of one of the trigraphs listed above is not
665 changed.

666 Defect Report UK 018: Operators and Punctuators

667 =====

668 The description of operators and punctuators is confusing, and the
669 constraints are contradictory.

670 Subclause 6.1.5 Constraints reads:

671 || The operators [], (), and ? : shall occur in pairs, possibly
672 || separated by expressions. The operators # and ## shall occur in
673 || macro-defining preprocessing directives only.

674 Subclause 6.1.6 Constraints reads:

675 || The punctuators [], (), and { } shall occur (after translation phase
676 || 4) in pairs, possibly separated by expressions, declarations, or
677 || statements. The punctuator # shall occur in preprocessing directives
678 || only.

679 Consider the code:

```
680 #define STR(x)  #x
681 STR ({)        /* Line A */
682 STR (:)        /* Line B */
683 STR ([)        /* Line C */
684 STR (#)        /* Line D */
```

685 Line A appears to be strictly conforming, since the first sentence of
686 the constraint of 6.1.6 does not apply during translation phase 4. Line B
687 violates the constraint of 6.1.5. The interpretation of line C depends
688 on whether the [is an operator or a punctuator !

689 Line D violates both constraints, but again which one depends on whether
690 it is an operator or a punctuator, something which is not made clear in
691 the Standard.

692 Assuming that the intent was for line B to be strictly conforming, and
693 that "(after translation phase 4)" was inadvertently omitted from 6.1.5,
694 the first sentence of each of these Constraints is nugatory, as any
695 program which violates these constraints also violates a syntax rule
696 elsewhere in clause 6. The remaining sentences would be better expressed
697 as part of subclause 6.8. It is also arguable that the concepts of operator
698 and punctuator are better merged at the syntactic level, and separated out
699 only at the semantic level.

700 Suggested Technical Corrigendum

701 -----

702 Delete the Constraints of subclauses 6.1.5 and 6.1.6. Add the following
703 constraint to 6.8:

704 A # preprocessing token shall only occur within a replacement-list
705 or when permitted by the syntax rules of this subclause. A ##
706 preprocessing token shall only occur within a replacement-list.

Line

721 Add to the end of the Constraints of subclause 6.1, just before the full
722 stop:

723 , and shall not be # or ##

725 Alternative Suggested Technical Corrigendum

726 -----
727 In subclause 6.1 syntax, delete both occurrences of "operator" and replace
728 the second occurrence of "punctuator" by "pp-punctuator".

729 Delete subclauses 6.1.5 and 6.1.6, and replace them by the following:

732 6.1.5 Punctuators

733 Syntax:

734 pp-punctuator:

735 punctuator

736 pp-only-punctuator

737 pp-only-punctuator: one of

738 # ## defined

739 punctuator:

740 [] () { } . ->

741 ++ -- & * + - ~ ! sizeof

742 / % << >> < > <= >= == != ^ | && ||

743 ? : , : ; ...

744 = *= /= %= += -= <<= >>= &= ^= |=

746 Semantics:

747 A punctuator is a symbol that has independent syntactic and semantic
748 significance. Depending on context, some punctuators may specify an
749 operation to be performed (an /evaluation/) that yields a value, or
750 yields a designator, or produces a side-effect, or a combination
751 thereof; in that context, the punctuator is known as an /operator/. An
752 /operand/ is an entity on which an operator acts.

753 Add the following constraint to 6.8:

754 A # preprocessing token shall only occur within a replacement-list
755 or when permitted by the syntax rules of this subclause. A ##
756 preprocessing token shall only occur within a replacement-list.

760 Defect Report UK 019: ranges of integral types

761 =====
762 It appears to be possible to create implementations with unreasonable
763 arrangements of integral types.

764 Subclause 6.1.2.5 states various rules which allow the following
765 deductions to be made:

768 SCHAR_MAX <= SHRT_MAX

769 SHRT_MAX <= INT_MAX

770 INT_MAX <= LONG_MAX

771 SCHAR_MIN >= SHRT_MIN

772 SHRT_MIN >= INT_MIN

773 INT_MIN >= LONG_MIN

774 SCHAR_MAX <= UCHAR_MAX

775 SHRT_MAX <= USHRT_MAX

776 INT_MAX <= UINT_MAX

777 LONG_MAX <= ULONG_MAX

779 and, depending on the interpretation of the term "the same amount of
780 storage":

782 sizeof (unsigned short) == sizeof (short)

783 sizeof (unsigned int) == sizeof (int)

784 sizeof (unsigned long) == sizeof (long)

786 However, (based on the preliminary discussions of DR 069, which allow padding
787 bits in integral types) there does not appear to be any requirement for the
788 following:

790 UCHAR_MAX <= USHRT_MAX

791 USHRT_MAX <= UINT_MAX

792 UINT_MAX <= ULONG_MAX

size of (short) <= sizeof (int)
sizeof (int) <= sizeof (long)
UCHAR_MAX <= INT_MAX

The first five of these are necessary to allow reasonable deductions to be made about the behavior of types in the presence of padding bits (for example, that unsigned long can hold any value representable in any integral type). The sixth is necessary to allow the <ctype.h> functions to behave sensibly (it is also assumed by example 2 of subclause 5.1.2.3).

Suggested Technical Corrigendum

In subclause 6.1.2.5, change in the fourth paragraph:

In the list of signed integer types above, the range of values of each type is a subrange of the values of the next type in the list.

to:

In the list of signed integer types above, the range of values of each type is a subrange of the values of the next type in the list, and the size of an object of each type is not greater than the size of an object of the next type in the list.

Add to the fifth paragraph:

The range of values of each unsigned integer type is a subrange of the next type (in the list unsigned char, unsigned short, unsigned int, unsigned long).

Add to the fifth or eighth paragraph:

The range of values of the type unsigned char is a subrange of the values of the type int.

Defect Report UK 020: Relational and Equality operators

The descriptions of these operators with pointer operands contain several defects.

Part 1

Consider the following code:

```
char *s = "a string";  
if (s >= NULL)  
    /* ... */
```

Subclause 6.3.8 Semantics reads in part:

|| If the objects pointed to are not members of the same aggregate or union
|| object, the result is undefined

This implies that the comparison causes undefined behavior.

Subclause 6.2.2.1 reads in part:

|| Such a pointer, called a null pointer, is guaranteed to compare unequal
|| to a pointer to any object or function.

This implies that the comparison is guaranteed to yield "false".

This is a direct contradiction.

Part 2

Subclause 6.3.9 Semantics reads in part:

|| Where the operands have types and values suitable for the relational
|| operators, the semantics detailed in 6.3.8 apply.

This can reasonably be read as meaning that, whenever the constraints of 6.3.8 apply, its definitions should be used, even if that would result in undefined behavior. [The phrase "and values" can reasonably be read as requiring only that the pointers both be to objects; it does not necessarily mean that the result of the comparison must be defined.]

It further reads:

|| If two pointers to object or incomplete types are both null pointers,

Line

```

865 || they compare equal. If two pointers to object or incomplete types compare
866 || equal, they both are null pointers, or both point to the same object, or
867 || both point one past the last element of the same array object.

```

```

869 This says nothing about the comparison of any other pointers. Now, -subclause
870 3.16 reads in part:
871 || Undefined behavior is otherwise indicated [...] by the omission of any
872 || explicit definition of behavior.

```

```

873
874 Thus, in:

```

```

875     int a, b;
876     &a == &b

```

```

877
878 the comparison causes undefined behavior !

```

```

880
881 Part 3

```

```

882 -----
883 The above citation does not allow for the case where one pointer is to an
884 object, and the other is one past the last element of an array object. If
885 an implementation places two independent objects in adjacent memory
886 locations, a pointer to one would equal a pointer to just past the other
887 on many common implementations.

```

```

888
889 If these pointers are not to be viewed as identical, then the wording is
890 defective.

```

```

891
892 Suggested Technical Corrigendum

```

```

893 -----
894 In subclause 6.2.2.1, replace the cited text by:
895     Such a pointer is called a null pointer.

```

```

896
897 In subclause 6.3.9, replace the first paragraph of the semantics by:
898     The operators == (equal to) and != (not equal to) shall yield 1 if the
899     specified relation is true and 0 if it is false. If the operands have
900     types suitable for those of a relational operator and values that would
901     not cause undefined behavior if used with a relational operator, then
902     the result of the comparison, either greater than or less than (both
903     implying not equal to) or equal to, is the same as with a relational
904     operator.

```

```

905 insert at the start of the second paragraph:
906     Otherwise the operands are pointers, and they shall compare either
907     equal or not equal.

```

```

908
909 If part 3 is viewed as an issue, then in the same paragraph change:
910     or both point one past the last element of the same array object.

```

```

911 to:
912     both point one past the last element of the same array object, or one
913     points one past the last element of some array object and the other
914     points to the first element of a different array object.

```

```

915
916 Defect Report UK 021: Line numbers

```

```

917 =====
918 The concept of "line number" is not clearly defined when a token is split
919 over more than one physical source line.

```

```

920
921 Subclause 6.8.4 reads in part:
922 || The line number of the current source line is one greater than the number
923 || of new-line characters read or introduced in translation phase 1 (5.1.1.2)
924 || while processing the source file to the current token.

```

```

925
926 Subclause 6.8.8 reads in part:
927 || __LINE__ The line number of the current source line (a decimal constant).

```

```

928
929 Consider the program:

```

```

930
931     #include <stdio.h>

```

```

932
933     #define LNER __LINE__

```

```

934
935     /* The next statement is on physical source lines 6 to 8 */
936

```



```

937 int east_coast = __\
938 LINE\
939 __;
940 /* The next statement is on physical source lines 10 to 13 */
941 int main_line = L\
942 N\
943 E\
944 R;
945
946 int main (void)
947 {
948     printf ("%d %d\n", east_coast, main_line);
949     return 0;
950 }

```

In each of the two substitutions, it is unclear whether the line number is the number of new-lines read to the *start of* the current token, or to the *end of* the current token, or to a specified point within the current token.

What is the output of this program ?

Defect Report UK 022: Implicit conversions

The wording dealing with the usual arithmetic conversions contains a number of errors; while the correct meaning is usually clear, a strict reading of the Standard shows some contradictions and/or unwanted side-effects.

Subclause 6.2.1.5 reads in part:

```

966 || Many binary operators that expect operands of arithmetic type cause
967 || conversions and yield result types in a similar way. The purpose is to
968 || yield a common type, which is also the type of the result.
969

```

Subclause 6.3.15 reads in part:

```

971 || The second operand is evaluated only if the first compares unequal to
972 || 0; the third operand is evaluated only if the first compares equal to 0;
973 || the value of the second or third operand (whichever is evaluated) is
974 || the result.
975 ||
976 || If both the second and third operands have arithmetic type, the usual
977 || arithmetic conversions are performed to bring them to a common type
978 || and the result has that type.
979 [...]
980 || in which case the other operand is converted to type pointer to void,
981 || and the result has that type.
982

```

These citations have several defects:

- * The relational and equality operators apply the usual arithmetic conversions, but not to yield the type of result.
- * The conditional operator ?: is not a binary operator, but is specified as performing the usual arithmetic conversions.
- * The concept of conversions applies only to a value; 6.3.15 is therefore contradicting itself when it calls for both the second and third operands to be subject to conversion when only one of them is evaluated.
- * The value of the result of the ?: is not necessarily that of the second or third operand, as the value may have been converted (possibly yielding a different value).

Suggested Technical Corrigendum

In 6.2.1.5, change the cited sentences to:

Many operators cause the same pattern of conversions to be applied to two operands of arithmetic type. The purpose is to yield a common type, which, unless explicitly stated otherwise, is also the type of the operator's result.

In 6.3.15, change the cited wording to:

The second operand is evaluated only if the first compares unequal to 0; the third operand is evaluated only if the first compares equal to 0; the result of the operator is the value of the second or third operand (whichever is evaluated), converted to the type described below.

Line

1009 If both the second and third operands have arithmetic type, the type
1010 that the usual arithmetic conversions would yield if applied to those
1011 two operands is the type of the result.
1012 [...]
1013 in which case the type of the result is pointer to void.

1016 Defect Report UK 023: Correction to Technical Corrigendum number 1
1017 =====

1018 An example added by TC1 is wrong.

1020 TC1 added the following example to subclause 7.9.6.2:
1021 || Add to subclause 7.9.6.2, page 138, another Example:
1022 || In:

```
1024 || #include <stdio.h>  
1025 || /* ... */  
1026 || int d1, d2, n1, n2, i;  
1027 || i = sscanf("123", "%d%n%n%d", &d1, &n1, &n2, &d2);
```

1029 || the value 123 is assigned to d1 and the value 3 to n1. Because %n can
1030 || never get an input failure the value of 3 is also assigned to n2. The
1031 || value of d2 is not affected. The value 3 is assigned to i.

1033 This should set i to 1, not 3, as %n does not affect the returned assignment
1034 count.

1036 Suggested Technical Corrigendum
1037 -----

1038 In the example, change:

1039 The value 3 is assigned to i.

1040 to:

1041 The value 1 is assigned to i.

1044 Defect Report UK 024: diagnostics for #error
1045 =====

1046 The rules concerning whether #error generates a diagnostic are contradictory.

1048 Subclause 5.1.1.3 reads:

1049 || A conforming implementation shall produce at least one diagnostic message
1050 || (identified in an implementation-defined manner) for every translation
1051 || unit that contains a violation of any syntax rule or constraint. Diagnostic
1052 || messages need not be produced in other circumstances.

1054 Subclause 6.8.5 reads:

```
1055 || Semantics  
1056 || A preprocessing directive of the form  
1057 || # error pp-tokens new-line  
1058 || opt  
1059 || causes the implementation to produce a diagnostic message that includes  
1060 || the specified sequence of preprocessing tokens.
```

1062 Since this is not in a Constraints section, these two statements directly
1063 contradict one another. Furthermore, the second statement can be read as
1064 applying to a #error directive that is excluded by a false #if condition.

1066 Suggested Technical Corrigendum
1067 -----

1068 In 6.8.5, replace the entire subclause with:

1069 Constraints

1070 A #error preprocessing directive shall not occur in a translation unit.

1071 Any diagnostic message generated because of the violation of this
1072 constraint [*] shall include the sequence of preprocessing tokens in the
1073 directive.

1074 [*] The intent of this subclause is that #error indicates that translation
1075 should fail. As stated in 5.1.1.3, a translation unit excludes lines
1076 within the "false" side of #if...#else...#endif groups.

1079 Defect Report UK 025: preprocessing directives
1080 =====

Preprocessing directives are not removed from the translation unit at any point during or after translation phase 4, and thus wreck the syntax analysis in translation phase 7.

Subclause 5.1.1.1 reads in part:

```
|| A source file together with all the headers and source files included
|| via the preprocessing directive #include, less any source lines skipped
|| by any of the conditional inclusion preprocessing directives, is called
|| a /translation unit/.
```

Nothing here, in the description of translation phase 4, or in subclause 6.8, states that any preprocessing directive is removed (except for #include, which is "replaced").

Consider the source file:

```
#define QUIT return 0
#if 0
This is some junk
#else
int main (void)
{
    puts ("Hello world\n");
}
#endif
QUIT;
```

The translation unit resulting at the end of translation phase 4 is thus:

```
#define QUIT return 0
#if 0
#else
int main (void)
{
    puts ("Hello world\n");
}
#endif
return 0;
```

and this clearly does not match the syntax of "translation-unit" in subclause 6.7.

Suggested Technical Corrigendum

In subclause 5.1.1.2, add at the end of the description of translation phase 4:

All preprocessing directives are then removed from the translation unit.