Issues on Spec 2.1

Spec 2.1 is NOT edited sufficiently clear and consistently.

21 October 2019

XBRL Japan Advisor
Representative of JIS X 7206:2020 editorial working group
SAMBUICHI Nobuyuki
nobuyuki@sambuichi.jp
https://www.sambuichi.jp/?lang=en

Errata

74. Changed from inferring precision to inferring decimals

07 March 2011

Section 4.6.6, Section 4.10, Section 5.2.5.2

75.Changed rounding descriptions to IEEE standard recommendation

31 October 2011

Section 4.6.7.1, Section 4.6.7.2

Ambiguous definition for “round” and “truncate”

PREVIOUS

RECOMMENDATION - 2003-12-31 + Corrected Errata - 2008-07-02

It is often necessary to round significant figures following a calculation. This is known as rounding. To round a number to n significant figures, discard all digits to the right of the nth place. This step is known as truncation. Then, if the left-most discarded digit is less than 5, leave the nth digit unchanged; if the left-most discarded digit is greater than or equal to 5, add 1 to the nth digit (propagating any carries to digits further to the left according to the normal rules of arithmetic and removing the final 0 if necessary).

CURRENT

Recommendation 31 December 2003 with errata corrections to 20 February 2013

It is often necessary to round significant figures following a calculation. This is known as rounding [IEEE] [IEEE 4.3.1 Rounding-direction attributes to nearest, roundTiesToEven]. To round a number to n significant figures, discard all digits to the right of the nth place. This step is known as truncation. Then, if the original number is equally near two truncated numbers, the one with an even nth digit is chosen.

PROBLEM 1
We can not get equally near <u>two truncated numbers</u> with this description.
Discarding all digits to the right of the nth place of original number results a single truncated number.
Requirement level is not clearly stated with statement is known as.

PROBLEM 2
IEEE Std. 754 is not general definition for rounding. It defines rounding for <u>representing floating-point in binary format or decimal format in microprocessor</u>.

XBRL THE BUSINESS REPORTING STANDARD

Clear definition “round” and “truncate” referencing to ISO 80000-1 Annex B

Rounding of numbers by JIS X 7206:2020

"Rounding" and "Truncation" MUST be performed as specified in [ISO] (ISO80000-1, Annex B Rounding of numbers). Rounding means replacing the magnitude of an original number by another number called the rounded number, selected from the sequence of integral multiples of a chosen rounding range. [IEEE] (IEEE 4.3 Rounding-direction) lists the rounding-direction attribute affects all computational operations that might be inexact.

The first rounding is “Rounding”, and its rounding-direction is “roundTiesToEven”.
If there is only one integral multiple nearest the original number, then that is accepted as the rounded number. If there are two successive integral multiples equally near the original number. the even multiple is selected as the rounded number. This rule is of special advantage when treating. for example. series of measurements in such a way that the rounding errors are minimized.

The second rounding is “Truncate”, and its rounding-direction is “roundTowardZzero”. An integral multiple closest to and no greater in multitude than infinity precise result the original number is accepted as the truncated number.

Rounding in more than one stage by the application of the rules given above may lead to errors. Therefore, the rounding MUST always be carried out in only one step.

EXAMPLE 12,254 should be rounded to 12,3 and not first to 12,25 and then to 12,2.

The rules given above should be used only if no special criteria for the selection of the rounded number MUST be taken into account. For instance, in case where safety requirements or other limits MUST be respected, it is advisable to round only in one direction.

Clear definition “round” and “truncate” referencing to ISO 80000-1 Annex B

Rounding of numbers by JIS X 7206:2020 (cont.)

<table>
<tr><th>original number</th><th>rounding range</th><th>integral multiple</th><th>Rounded number
roundTiesToEven</th><th>Truncated number
roundTowardZero</th></tr>
<tr><td>-12.251</td><td rowspan="5">0.1
(@decimals=“1”)</td><td rowspan="2">-12.4; -12.3, -12.2; -12.1; etc.</td><td>-12.3</td><td rowspan="2">-12.2</td></tr>
<tr><td>-12.233</td><td>-12.2</td></tr>
<tr><td>12.233</td><td rowspan="3">12.1; 12.2; 12.3; 12.4; etc.</td><td rowspan="2">12.2</td><td rowspan="3">12.2</td></tr>
<tr><td>12.250</td></tr>
<tr><td>12.275</td><td>12.3</td></tr>
<tr><td>-1 225.1</td><td rowspan="5">10
(@decimals=“-1”)</td><td rowspan="2">-1 240; -1 230; -1 220; -1210; etc.</td><td>-1 230</td><td rowspan="2">-1 220</td></tr>
<tr><td>-1 223.3</td><td>-1 220</td></tr>
<tr><td>1 223.3</td><td rowspan="3">1 210; 1 220; 1 230; 1 240; etc.</td><td rowspan="2">1 220</td><td rowspan="3">1 220</td></tr>
<tr><td>1 225.0</td></tr>
<tr><td>1 227.5</td><td>1 230</td></tr>
</table>

V-Equal

JIS Z 7206:2010 V-Equal definition

		Recommendation 31 December 2003 with errata corrections to 20 February 2013	JIS X 7206:2020
Numeric Items not of type fractionItemType or a type derived from fractionItemType by restriction	V-Equal	**A** and **B** are V-Equal if and only if all the following conditions apply: 1. **A** and **B** are C-Equal and U-Equal 2. the numeric values $\mathbf{A}_N$ and $\mathbf{B}_N$ are X-Equal where $\mathbf{A}_N$ is obtained by rounding the content of **A** to **N** significant figures and $\mathbf{B}_N$ is obtained by rounding the content of **B** to **N** significant figures where **N** is the lower of: a. the specified or inferred decimals for **A** and b. the specified or inferred decimals for **B** (If either Numeric Item has a @precision attribute value 0 then the v-equality is false.)	**A** and **B** are V-Equal if and only if all the following conditions apply: **1.** **A** and **B** are C-equal and U-Equal. **2.** $\mathbf{A}_N$ and $\mathbf{B}_N$ are X-Equal where $\mathbf{A}_N$ is obtained by rounding the content of **A** correct to **N** decimal places and $\mathbf{B}_N$ is obtained by rounding the content of **B** correct to **N** decimal places where **N** is the lower of: a. the specified or inferred decimals for **A** and b. the specified or inferred decimals for **B**

Round content to N significant figures
-> Round content correct o N decimal places

JIS Z 7206:2010 V-Equal definition (cont.)

		Recommendation 31 December 2003 with errata corrections to 20 February 2013	JIS X 7206:2020
Numeric Items of type fractionItem Type or a type derived from fractionItemType by restriction	V-Equal	**A** and **B** are V-Equal if and only if all the following conditions apply: 1. **A** and **B** are C-Equal and U-Equal **2.** $\mathbf{A_N}$ is X-Equal to $\mathbf{B_N}$ and $\mathbf{A_D}$ is X-Equal to $\mathbf{B_D}$ where: a. $\mathbf{A_N}$ is the numerator and $\mathbf{A_D}$ is the denominator of the normal form (defined below) of **A** and b. $\mathbf{B_N}$ is the numerator and $\mathbf{B_D}$ is the denominator of the normal form of **B.** For any item F of type fractionItemType or a type derived from fractionItemType by restriction, the normal form has numerator F_N and denominator F_D such that F_N and F_D are integers and have no integer common factor and there exists a number H such that multiplying F_N by H gives the numerator of F and multiplying F_D by H gives the denominator of F.	Same as left

Inconsistency between example 11 and example 52

4.6.4 The @precision attribute (optional)

TRUNCATE

Example 52. Calculations involving decimals and precision

ROUND

```
<a contextRef="c1" unitRef="u1" precision="2">1559</a>
<b contextRef="c1" unitRef="u1" precision="3">984.8</b>
<c contextRef="c1" unitRef="u1" decimals="1">582.334973</c>
```

Consuming application are required to handle inconsistent ways.
The rounding shall always be carried out in only one step.

Inconsistency between example 11 and example 52

4.6.4 The @precision attribute (optional)

TRUNCATE

If a numeric fact has a @precision attribute that has the value "n" then it is correct to "n" significant figures (see **Section 4.6.1** for the normative definition of 'correct to "n" significant figures').

An application **SHOULD** ignore (i.e. replace with zeroes) any digits after the first "*n*" decimal digits, counting from the left, starting at the first non-zero digit in the lexical representation of any number for which the value of precision is specified or inferred to be *n*.

```
<a contextRef="c1" unitRef="u1" precision="2">1559</a>
<b contextRef="c1" unitRef="u1" precision="3">984.8</b>
<c contextRef="c1" unitRef="u1" decimals="1">582.334973</c>
```

Example 52. Calculations involving decimals and precision

ROUND

Suppose that the numeric item a is a summation for numeric items b and c (with weight 1.0) and there exists a context with id 'c1' and unit with id 'u1' in the instance so that the summation binds.

To perform the calculation, **first round** 984.8 **to precision** 3 to give 985 and **then round** 582.334973 **to the inferred precision 4** to give 582.3 resulting in a total of 1567.3. The total is then equal to the summation item after **rounding to precision 2** (the precision of the summation item a) at 1600, so that this calculation is consistent.

Consuming application are required to handle inconsistent ways.
The rounding shall always be carried out in only one step.

Remove truncate example in example 11

4.6.4 The @precision attribute (optional)

TRUNCATE

If a numeric fact has a @precision attribute that has the value "n" then it is correct to "n" significant figures (see **Section 4.6.1** for the normative definition of 'correct to "n" significant figures').

An application **SHOULD** ignore (i.e. replace with zeroes) any digits after the first "*n*" decimal digits, counting from the left, starting at the first non-zero digit in the lexical representation of any number for which the value of precision is specified or inferred to be *n*.

JIS Z 7206:2020

~~TRUNCATE~~

~~If a numeric fact has a @precision attribute that has the value "n" then it is c~~ [illegible] ~~" significant figures (see S~~ [illegible] ~~on 4.6.~~ [illegible] ~~the normative~~ [illegible] ~~n of 'co~~ [illegible] ~~to "n" significant~~ [illegible] ~~ure~~ [illegible]

~~An applica~~ [illegible] ~~n SHO~~ [illegible] ~~D ig~~ [illegible] ~~e (i.e. replace wi~~ [illegible] ~~eroes) a~~ [illegible] ~~after the first "*n*" dec~~ [illegible] ~~l digits,~~ [illegible] ~~ng from the left, starting a~~ [illegible] ~~zero digit in the lexical represe~~ [illegible] ~~on of any number for which the value of precision is specified or inferred to be *n*.~~

Rounding in more than one stage by the application may lead to errors. Therefore, <u>the rounding MUST always be carried out in only one step</u>.

XBRL
THE BUSINESS REPORTING STANDARD

Modify example 52 to carry out rounding in single step

```
<a contextRef="c1" unitRef="u1" precision="2">1559</a>
<b contextRef="c1" unitRef="u1" precision="3">984.8</b>
<c contextRef="c1" unitRef="u1" decimals="1">582.334973</c>
```

Example 52. Calculations involving decimals and precision

Suppose that the numeric item a is a summation for numeric items b and c (with weight 1.0) and there exists a context with id 'c1' and unit with id 'u1' in the instance so that the summation binds.

To perform the calculation, first **round** 984.8 to precision 3 to give 985 and then **round** 582.334973 to the inferred precision 4 to give 582.3 resulting in a total of 1567.3. The total is then equal to the summation item after **rounding** to precision 2 (the precision of the summation item a) at 1600, so that this calculation is consistent.

JIS Z 7206:2020

Suppose that the numeric item a is a summation for numeric items b and c (with weight 1.0) and there exists a context with id 'c1' and unit with id 'u1' in the instance so that the summation binds.

To perform the calculation, first add 984.8 and 582.334 973 to get 1567.13497, then check if the total is v-equal. The inferred decimals of the summation item a is -2 and 1559 correct to -2 decimal places is1600. The total 1567.13497 correct to -2 decimal places is also1600. These two values are not only both C-Equal and U-equal but also V-Equal, so that this calculation is consistent.

The rounding shall always be carried out in only one step.

The final V-Equal test step.

Calculation based on 4.6.4 vs JIS

Example 52 Round	4.6.4 way		JIS Z 7206:2020 Rounding in final V-Equal check only
	Truncate: Read as (after omitting or zeroing any spurious digits)	value range based on roundToEven	
first <u>round</u> 984.8 <u>to precision 3</u> to give **985**	first <u>truncate</u> 984.8 <u>to precision 3</u> to give **984**	[983.5, 984.5] Closed interval from 983.5 incl. to 984.5 incl.	**984.8**
then <u>round</u> 582.334973 <u>to the inferred precision 4</u> to give **582.3**	then <u>truncate</u> 582.334973 <u>to the inferred precision 4</u> to give **582.3**	(582.25, 582.35) Open interval from 582.25 excl. to 582.35 excl.	**582.334973**
resulting in a total of **1567.3**	resulting in a total of **1566.3**	?	resulting in a total of **1567.13497**
		V-Equal test	
total 1567.3 <u>rounding to precision 2</u> is **1600**	total 1566.3 <u>truncating to precision 3</u> is **1570**	**1570** <u>rounding to @precision 2</u> is **1600**	Total 1567.13497 <u>rounding to @decimals -2</u> is **1600**
summation 1559 <u>rounding to precision 2</u> is **1600**	summation 1559 <u>truncating to precision 2</u> is **1500**	**1500** <u>rounding to @precision 2</u> is **1500** ✗	1599 <u>rounding to @decimals -2</u> is **1600** ○

Rounding in more than one stage by the application may lead to errors.
Therefore, <u>the rounding MUST always be carried out in only one step for V-Equal test</u>.

XBRL THE BUSINESS REPORTING STANDARD

21 October 2019

XBRL Japan Advisor
Representative of JIS X 7206:2020 editorial working group
ISO TC 295 Audit data services
Head of delegate Japanese Industrial Standards Committee

SAMBUICHI Nobuyuki

nobuyuki@sambuichi.jp
https://www.sambuichi.jp/?lang=en
https://bit.ly/32xIS9W

IEEE

IEEE Standard for Floating-Point Arithmetic

IEEE Computer Society

Sponsored by the
Microprocessor Standards Committee

Contents

Table 3.1—Relationships between different specification levels for a particular format

Level 1	$\{-\infty \ldots \mathbf{0} \ldots +\infty\}$	Extended real numbers.
many-to-one ↓	*rounding*	↑ projection (except for NaN)
Level 2	$\{-\infty \ldots -0\} \cup \{+0 \ldots +\infty\} \cup \mathbf{NaN}$	Floating-point data—an algebraically closed system.
one-to-many ↓	*representation specification*	↑ many-to-one
Level 3	***(sign, exponent, significand)*** $\cup \{-\infty, +\infty\} \cup$ **qNaN** $\cup$ **sNaN**	Representations of floating-point data.
one-to-many ↓	*encoding for representations of floating-point data*	↑ many-to-one
Level 4	**0111000…**	Bit strings.

The mathematical structure underpinning the arithmetic in this standard is the extended reals, that is, the set of real numbers together with positive and negative infinity. For a given format, the process of *rounding* (see 4) maps an extended real number to a *floating-point number* included in that format. A *floating-point datum*, which can be a signed zero, finite non-zero number, signed infinity, or a NaN (not-a-number), can be mapped to one or more *representations of floating-point data* in a format.

Table 3.2—Parameters defining basic format floating-point numbers

	Binary format (b=2)			**Decimal format** (b=10)	
parameter	**binary32**	**binary64**	**binary128**	**decimal64**	**decimal128**
p, digits	24	53	113	16	34
emax	+127	+1023	+16383	+384	+6144

IEEE Standard for Floating-Point Arithmetic

4. Attributes and rounding

4.1 Attribute specification

An attribute is logically associated with a program block to modify its numerical and exception semantics. A user can specify a constant value for an attribute parameter. Some attributes have the effect of an implicit parameter to most individual operations of this standard; language standards shall specify — rounding-direction attributes (see 4.3) and should specify — alternate exception handling attributes (see Clause 8). Other attributes change the mapping of language expressions into operations of this standard; language standards that permit more than one such mapping should provide support for: — preferredWidth attributes (see 10.3) — value-changing optimization attributes (see 10.4) — reproducibility attributes (see Clause 11). For attribute specification, the implementation shall provide language-defined means, such as compiler directives, to specify a constant value for the attribute parameter for the operations in a block; the scope of the attribute value is the block with which it is associated. Language standards shall provide for constant specification of the default and each specific value of the attribute.

4.2 Dynamic modes for attributes

Attributes in this standard shall be supported with the constant specification of 4.1. Particularly to support debugging, language standards should also support dynamic-mode specification of attributes. With dynamic-mode specification, a user can specify that the attribute parameter assumes the value of a dynamic-mode variable whose value might not be known until program execution. This standard does not specify the underlying implementation mechanisms for constant attributes or dynamic modes. For dynamic-mode specification, the implementation shall provide language-defined means to specify that the attribute parameter assumes the value of a dynamic-mode variable for the operations within the scope of the dynamic-mode specification in a block. The implementation initializes a dynamic-mode variable to the default value for the dynamic mode. Within its language-defined (dynamic) scope, changes to the value of a dynamic-mode variable are under the control of the user via the operations in 9.3.1 and 9.3.2. The following aspects of dynamic-mode variables are language-defined; language standards may explicitly defer the definitions to implementations: — The precedence of static attribute specifications and dynamic-mode assignments. — The effect of changing the value of the dynamic-mode variable in an asynchronous event, such as in another thread or signal handler. — Whether the value of the dynamic-mode variable can be determined by non-programmatic means, such as a debugger. NOTE—A constant value for an attribute can be specified and meet the requirements of 4.1 by a dynamic mode specification with appropriate scope of that constant value.

IEEE Standard for Floating-Point Arithmetic

4.3 Rounding-direction attributes

Rounding takes a number regarded as infinitely precise and, if necessary, modifies it to fit in the destination's format while signaling the inexact exception, underflow, or overflow when appropriate (see Clause 7). Except where stated otherwise, every operation shall be performed as if it first produced an intermediate result correct to infinite precision and with unbounded range, and then rounded that result according to one of the attributes in this clause. Except where stated otherwise, the rounding-direction attribute affects all computational operations that might be inexact. Inexact numeric floating-point results always have the same sign as the unrounded result. The rounding-direction attribute affects the signs of exact zero sums (see 6.3), and also affects the thresholds beyond which overflow (see 7.4) and underflow (see 7.5) are signaled. Implementations supporting both decimal and binary formats shall provide separate rounding-direction attributes for binary and decimal, the binary rounding direction and the decimal rounding direction. Operations returning results in a floating-point format shall use the rounding-direction attribute associated with the radix of the results. Operations converting from an operand in a floating-point format to a result in integer format or to an external character sequence (see 5.8 and 5.12) shall use the rounding-direction attribute associated with the radix of the operand. NaNs are not rounded (but see 6.2.3).

IEEE Standard for Floating-Point Arithmetic

4.3.1 Rounding-direction attributes to nearest

In the following two rounding-direction attributes, an infinitely precise result with magnitude at least $b^{emax} \times (b - \frac{1}{2} b^{1-p})$ shall round to ∞ with no change in sign; here *emax* and *p* are determined by the destination format (see 3.3). With:

— **roundTiesToEven**, the floating-point number nearest to the infinitely precise result shall be delivered; if the two nearest floating-point numbers bracketing an unrepresentable infinitely precise result are equally near, the one with an even least significant digit shall be delivered; if that is not possible, the one larger in magnitude shall be delivered (this can happen for one-digit precision, possible with convertToDecimalCharacter for example, as when rounding 9.5 to one digit in which case both 9 and 1×10^1 have odd significands)

— **roundTiesToAway**, the floating-point number nearest to the infinitely precise result shall be delivered; if the two nearest floating-point numbers bracketing an unrepresentable infinitely precise result are equally near, the one with larger magnitude shall be delivered.

IEEE Standard for Floating-Point Arithmetic

4.3.2 Directed rounding attributes

Three other user-selectable rounding-direction attributes are defined, the directed rounding attributes roundTowardPositive, roundTowardNegative, and roundTowardZero. With:

— **roundTowardPositive**, the result shall be the format's floating-point number (possibly $+\infty$) closest to and no less than the infinitely precise result

— **roundTowardNegative**, the result shall be the format's floating-point number (possibly $-\infty$) closest to and no greater than the infinitely precise result

— **roundTowardZero**, the result shall be the format's floating-point number closest to and no greater in magnitude than the infinitely precise result.

4.3.3 Rounding attribute requirements

An implementation of this standard shall provide roundTiesToEven and the three directed rounding attributes. A decimal format implementation of this standard shall provide roundTiesToAway as a user selectable rounding-direction attribute. The rounding attribute roundTiesToAway is not required for a binary format implementation. The roundTiesToEven rounding-direction attribute shall be the default rounding-direction attribute for results in binary formats. The default rounding-direction attribute for results in decimal formats is language-defined but should be roundTiesToEven.

Standards All about ISO Taking part Store

ICS › 01 › 01.060

ISO 80000-1:2009

Quantities and units — Part 1: General

Annex B
(normative)

Rounding of numbers

Annex B (normative) Rounding of numbers

B.1 Rounding means replacing the magnitude of a given number by another number called the *rounded number*, selected from the sequence of integral multiples of a chosen rounding range.

EXAMPLE 1

rounding range: 0,1
integral multiple: 12,1; 12,2; 12,3; 12,4; etc.

EXAMPLE 2

rounding range: 10
integral multiple: 1 210; 1 220; 1 230; 1 240; etc.

B.2 If there is only one integral multiple nearest the given number, then that is accepted as the rounded number.

EXAMPLE 1

rounding range: 0,1

given number	rounded number
12,233	12,2
12,251	12,3
12,275	12,3

EXAMPLE 2

rounding range: 10

given number	rounded number
1 223,3	1 220
1 225,1	1 230
1 227,5	1 230

rounding range should always be indicated.

Annex B (normative) Rounding of numbers (cont.)

B.3 If there are two successive integral multiples equally near the given number, two different rules are in use.

Rule A: The even multiple is selected as the rounded number.

EXAMPLE 1

rounding range: 0,1

given number	rounded number
12,25	12,2
12,35	12,4

EXAMPLE 2

rounding range: 10

given number	rounded number
1 225,0	1 220
1 235,0	1 240

Rule B: The greater in multitude is selected as the rounded number.

EXAMPLE 1

rounding range: 0,1

given number	rounded number
12,25	12,3
12,35	12,4
-12,25	-12,3
-12,35	-12,4

EXAMPLE 2

rounding range: 10

given number	rounded number
1 225,0	1 230
1 235,0	1 240
-1 225,0	-1 230
-1 235,0	-1 240

Rule A is generally preferable and of special advantage when treating, for example, series of measurements in such a way that the rounding errors are minimized.

Rule B is sometimes used in computers.

Annex B (normative) Rounding of numbers (cont.)

B.4 Rounding in more than one stage by the application of the rules given above may lead to errors.

Therefore, the rounding shall always be carried out in only one step.

EXAMPLE 12,254 should be rounded to 12,3 and not first to 12,25 and then to 12,2.

B.5 The rules given above should be used only if no special criteria for the selection of the rounded number have to be taken into account. For instance, in case where safety requirements or other limits have to be respected, it is advisable to round only in one direction.

B.6 The rounding range should always be indicated.