

The `xintexpr` and allied packages

JEAN-FRANÇOIS BURNOL

jfbu (at) free (dot) fr

Package version: 1.4m (2022/06/10); documentation date: 2022/06/11.

From source file xint.dtx. Time-stamp: <11-06-2022 at 19:36:36 CEST>.

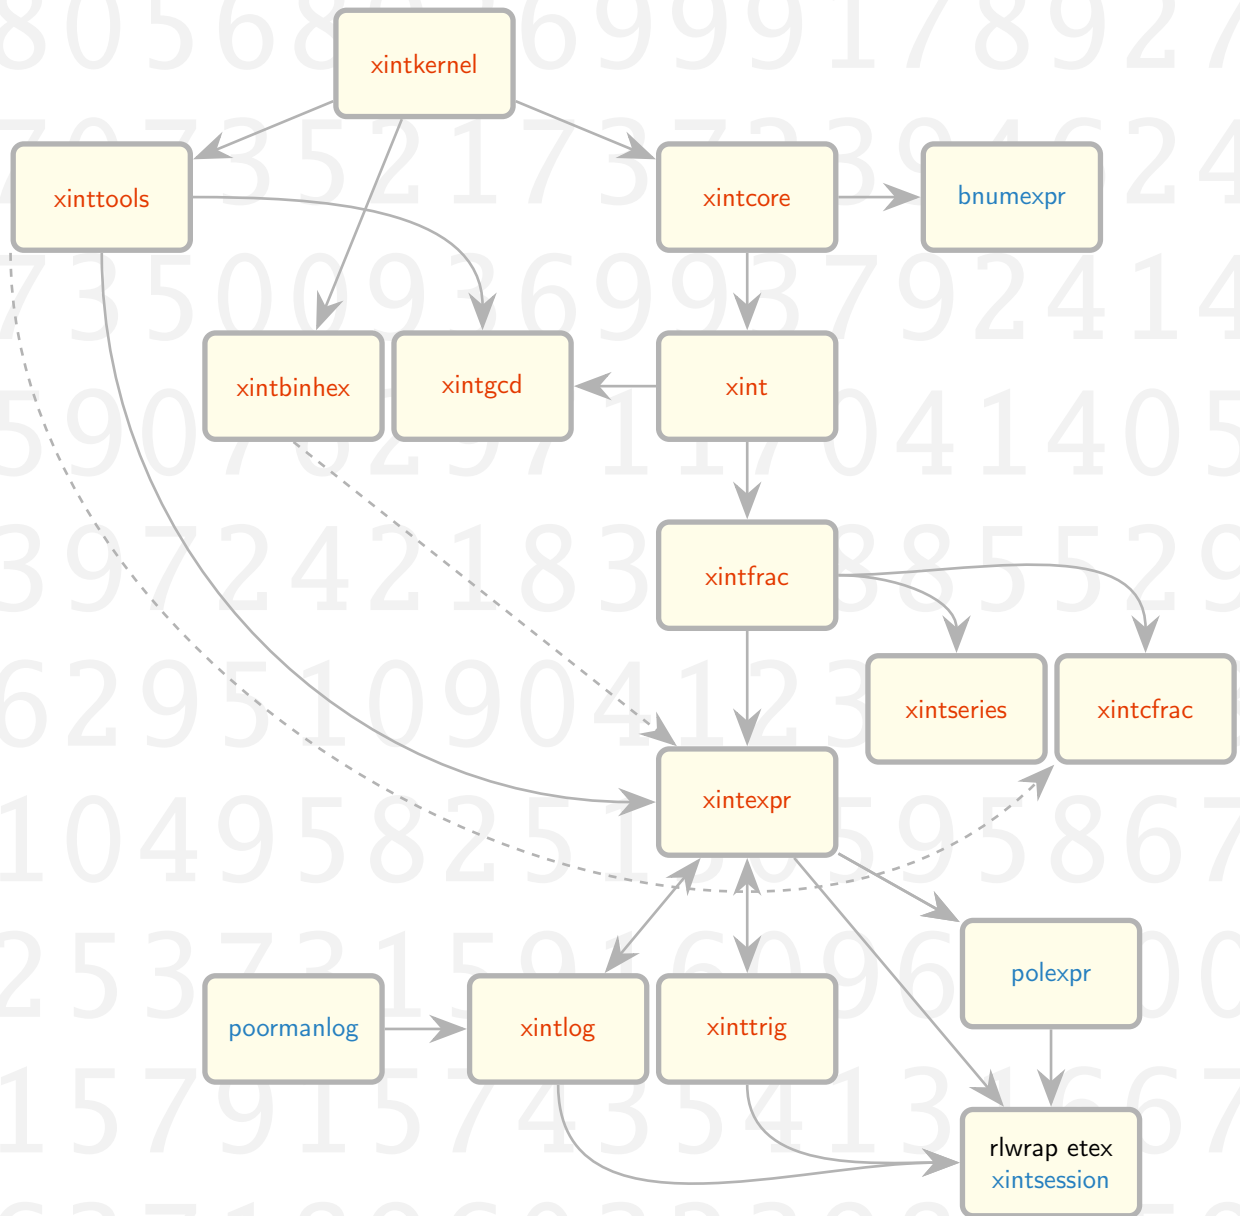
Part I. The `xintexpr` package

1	Introduction	3
1.1	Basic usage and purpose	3
1.2	<code>xintsession</code>	4
1.3	<code>polexpr</code>	5
1.4	<code>bnumexpr</code>	5
1.5	License and installation instructions	6
1.6	Printing big numbers on the page	6
1.7	The package documentation	6
2	Syntax reference and user guide	7
2.1	The three parsers	7
2.2	Output customization	9
2.3	Built-in operators and their precedences	13
	<i>Table of precedence levels of operators</i>	15
2.4	Built-in functions	18
	<i>Table of functions in expressions</i>	19
2.5	Generators of arithmetic progressions	35
2.6	Python slicing and indexing of one-dimensional sequences	36
2.7	NumPy like nested slicing and indexing for arbitrary oples and nutples	37
2.8	Tacit multiplication	38
2.9	User defined variables	39
2.10	User defined functions	42
2.11	Examples of user defined functions	49
2.12	Links to some (old) examples within this document	53
2.13	Oples and nutples: the 1.4 terminology	53
2.14	Expansion (for geeks only)	58
2.15	Known bugs/features (last updated at 1.4m)	59
3	The macros of <code>xintexpr</code> (ancient documentation, mostly)	61
4	The <code>xinttrig</code> package	78
5	The <code>xintlog</code> package	82
6	Macros of the <code>xinttools</code> package	85
7	Additional (old) examples with <code>xinttools</code> or <code>xintexpr</code> or both	106

Part II. The macro layer for expandable computations: `xintcore`, `xint`, `xintfrac`, and some extras

8	The <code>xint</code> bundle	125
9	Macros of the <code>xintkernel</code> package	140
10	Macros of the <code>xintcore</code> package	146
11	Macros of the <code>xint</code> package	151
12	Macros of the <code>xintfrac</code> package	163
13	Macros of the <code>xintbinhex</code> package	193
14	Macros of the <code>xintgcd</code> package	196
15	Macros of the <code>xintseries</code> package	198
16	Macros of the <code>xintcfac</code> package	214

Part III. The `xintexpr` and allied packages source code (not included, see `sourcexint.pdf`)



Dependency graph for the `xint` bundle components: modules pointed to by arrows **automatically** import the modules originating the continuous line ended by an arrow. Dashed lines indicate a partial dependency, and to enable the corresponding functionalities of the lower module it is thus necessary to use a suitable `\usepackage` (L^AT_EX) or `\input` (Plain T_EX.)

`bnumexpr` is a separate (L^AT_EX only) package by the author which uses (by default) `xintcore` as its mathematical engine.

`polexpr` handles definitions and algebraic operations on one-variable polynomials, as well as root localization to arbitrary precision. It works both with Plain T_EX and with L^AT_EX.

`xintrig` and `xintlog` are loaded automatically by `xintexpr`; they should not be loaded directly via a separate `\usepackage` (in L^AT_EX).

`poormanlog` is a T_EX and L^AT_EX package by the author which is loaded automatically by `xintlog`.

`xintsession` is invoked on the command line as `etex xintsession` (or, if available, `rnowrap etex xintsession`). It loads `xintbinhex` automatically (but this is not indicated above graphically).

Part I.

The **xintexpr** package

1	Introduction	3
2	Syntax reference and user guide	7
3	The macros of xintexpr (ancient documentation, mostly)	61
4	The xinttrig package	78
5	The xintlog package	82
6	Macros of the xinttools package	85
7	Additional (old) examples with xinttools or xintexpr or both	106

1. Introduction

.1	Basic usage and purpose	3	.5	License and installation instructions	6
.2	xintsession	4	.6	Printing big numbers on the page	6
.3	polexpr	5	.7	The package documentation	6
.4	bnumexpr	5			

JÜRGEN GILG's interest into what he called "XINT" was instrumental in keeping the author motivated over the years. We exchanged on many topics extending beyond \TeX and often reacted similarly to private and public events. I knew he was a very kind and devoted person, who took care of the needs of others prior to his own, although he never mentioned it. Jürgen suffered a sudden, unexpected, and deadly stroke in May 2022. I will miss his friendship profoundly.

1.1. Basic usage and purpose

- To use with **etex**, **pdftex**, ..., i.e. with \TeX engines activating the $e\TeX$ extensions:

```
\input xintexpr.sty
```

- To use with the \LaTeX macro layer (**latex**, **pdflatex**, ...):

```
\usepackage{xintexpr}
```

Some random examples:

```


$$\sqrt{13}, \cos(1), \exp(13.3) = \text{\xintfloateval{\sqrt{13}, \cos(1), \exp(13.3)}}\par
\sqrt{13}, \cos(1), \exp(13.3) = 3.605551275463989, 0.5403023058681397, 597195.6137928163$$


```

\begin{center}
 \xintDigits*:=32;% this sets the precision to 32 digits for the math functions
 \begin{minipage}{34\fontcharwd\font`0}
 \xintfloateval{\sqrt{13}, \cos(1), \exp(13.3)}\par
 \end{minipage}%

```


```

```
\end{center}
```

```
3.6055512754639892931192212674705,
0.54030230586813971740093660744298,
597195.61379281625101872789722711
```

```
$2^{1000}=\printnumber{\xinteval{2^1000}}$\par% exact computations
```

```
21000 = 1071508607186267320948425049060001810561404811705533607443750388370351051124936122493
19837881569585812759467291755314682518714528569231404359845775746985748039345677748242309854
21074605062371141877954182153046474983581941267398767559165543946077062914571196477686542167
660429831652624386837205668069376
```

```
% dummy variables
```

```
\[\sum_{i=1}^{50}\frac{i}{i^2}=\xintTeXFrac{\xinteval{reduce(add(1/i, i=1..50))}}\]
```

$$\sum_{i=1}^{50} \frac{1}{i} = \frac{13943237577224054960759}{3099044504245996706400}$$

```
\[\sum_{i=1}^{50}\frac{1}{i^2}=\xintTeXFrac{\xinteval{reduce(add(1/i^2, i=1..50))}}\]
```

$$\sum_{i=1}^{50} \frac{1}{i^2} = \frac{3121579929551692678469635660835626209661709}{1920815367859463099600511526151929560192000}$$

Computations are done expandably. This means that for example

```
\typeout{\xinteval{sqrt(13, 70)}}
```

will compute on the fly while writing out to the log file (`\typeout` is the \TeX idiom), here the square root of 13 correctly rounded to 70 significant digits.

Indeed, `xintexpr` is a package to do expandable computations, either exactly with arbitrarily big inputs (fractions, arbitrarily long decimal expansions, ...), or in the sense of floating point numbers (logarithm, exponential, sine, cosine, ...). The math functions are implemented up to 62 digits of precision. The square root (as well of course as the four operations) achieve correct rounding in arbitrary precision (i.e. even with more than 62 digits, and may even handle thousands of digits but reasonable range is not beyond a few hundreds).

The syntax to modify the precision used for floating point evaluations is

```
\xintDigits*:= <Number>;
```

For reasons of the history of the package, one usually will want to use indeed the `*` (see `\xintDigits`), as the more lightweight variant without it does let the four operations and the square root obey the new setting but it does not recalibrate the `xinttrig` and `xintlog` libraries and their built-in constants.

The current precision is available as `\xinttheDigits`, and defaults to 16, but in this documentation I might be using simply `Digits` to refer to it.

The table of [built-in operators](#) and the one of [built-in functions](#) will give a quick overview of the available syntax.

1.2. xintsession

The simplest way¹ to test the syntax is to work interactively on the command line (this feature is available since April 2021, the version of `xintsession` used here is 1.3a):

```
rlwrap etex xintsession
[...welcome banner...]
Magic words: '&pause' (or `;'), '&help', '&bye',
             '&exact', '&fp', '&int', '&pol'.

\jobname is xintsession
Transcript will go to log and to xintsession-210609_12h00.tex
Starting in exact mode (floating point evaluations use 16 digits)
```

¹ I am assuming here in the displayed example a Unixen system, i.e. Mac OS or Linux, adapt to your environment.

```

>>> 2^100;
@_1      1267650600228229401496703205376
>>> cos(1);
@_2      0.5403023058681397
>>> &fp=32
(/usr/local/texlive/2021/texmf-dist/tex/generic/xint/xintlog.sty)
(/usr/local/texlive/2021/texmf-dist/tex/generic/xint/xinttrig.sty)
fp mode (log and trig reloaded at Digits=32)
>>> cos(1);
@_3      0.54030230586813971740093660744298
>>> 3^1000;
@_4      1.3220708194808066368904552597521e477
>>> &exact
exact mode (floating point evaluations use 32 digits)
>>> 3^1000;

@_5      13220708194808066368904552597521443659654220327521481676649203682268285
9734670489954077831385060806196390977769687258235595095458210061891186534272525
7953674027620225198320803878014774228964841274390400117588618041128947815623094
4380615661730540866744905061781254803444055470543970388958174653682549161362208
3026856377858229022841639830788789691855640408489893760937324217184635993869551
6765018940588109060426089671438864102814350385648747165832010614366132173102768
902855220001
>>> &bye
Did I say something wrong?
Session transcript written on xintsession-210609_12h00.tex
)
No pages of output.
Transcript written on xintsession.log.

```

1.3. `polexpr`

The package `polexpr` enriches the `\xinteval` syntax (but not the one of `\xintfloateval`) with a polynomial type with associated constructor `pol([c0,c1,...])`, and polynomial specific functions such as `polgcd(pol1, pol2, ...)`.

Full usage of polynomials (and algebraic notations $c_0 + c_1 x + c_2 x^2 + \dots$ for input and also output) goes through a dedicated `\poldef` parser which is based upon `\xintdefvar/\xintdeffunc` and is a necessary step to then access via a dedicated macro interface operations such as identifying all rational roots and isolating all real roots to arbitrary precision.

The simplest manner to experiment with `polexpr` is via the `&pol` mode of `xintsession`.

1.4. `bnumexpr`

This small package loads `xintcore` and provides the functionality of the integer-only `\xintiexpr/\xintiieval` parser, but dropping support for nested structures, functions, variables, boolean branching, etc... It includes by default support for hexadecimal input (with `"` prefix) and output based upon `xintbinhex`. Compared to `\the\numexpr` it adds support for arbitrarily large integers, powers with `**` and `^`, factorials via `!` postfix operator, floored division and associated modulo, and comma separated multi-item expressions.

It also offers a user-level interface to add extra infix or postfix operators.

1.5. License and installation instructions

`xint` is made available under the [LaTeX Project Public License 1.3c](#) and is included in the major \TeX distributions, thus there is probably no need for a custom install: just use the package manager to update if necessary `xint` to the latest version available.

Else, [CTAN](#) access provides `xint.tds.zip` which has all source code and documentation in a TDS-compliant archive, only waiting to be `unzip -d <DIR>` into some suitable hierarchical structure.

The [README.md](#) at [CTAN](#) lists some alternatives.

1.6. Printing big numbers on the page

When producing very long numbers there is the question of printing them on the page, without going beyond the page limits. In this document, I have most of the time made use of a `\printnumber` macro, not provided by the package, and which is coded like this (with some extra decorations):

```
\def\allowsplits #1{\ifx #1\relax \else #1\hskip 0pt plus 1pt\relax
\expandafter\allowsplits\fi}%
\def\printnumber #1{\expandafter\allowsplits \romannumeral-`0#1\relax }%
```

This macro triggers `f-expansion` of its argument (and indeed `\xinteval` and friends expand completely under such triggering), then it goes through the computation result character by character inserting \TeX potential break points in-between them.

```
\printnumber{\xintiieval{100!^3}}\newline
81285103704665697929058034741394527800954175275203119077085794747670888482337305968567201883
75050478138776220712647125923141159206411609199354037545836490698436012619000519089702481351
07234498895796609463150334493880799668742586291763030205250590988746228607583652771623341365
91629009247695685942955467213561895127511100771737329147330105403484204308951158469957099274
14697054763835474153299936479805440000000000000000000000000000000000000000000000000000000000000
00000000000000
```

1.7. The package documentation

With \TeX distributions providing a `texdoc` or similar utility,

```
texdoc --list xint
```

will offer the choice to visualize one of those files:

1. `xint.pdf` (this file)
2. `sourcexint.pdf` (commented (hmm...) source code)
3. `CHANGES.html`
4. `README.md`

Warning: I don't have the time to maintain perfectly such large documentation. It currently combines old documentation which never really got updated and may be locally obsolete with more recent stuff mostly written on occasion of the 1.4 release of January 2020 and the 1.4e one of May 2021, and the intervening changes might also have made some of it not completely accurate, despite my best efforts. Bug reports and feature requests are most welcome.

The section formerly included with this documentation and documenting the changes since the 1.4 release has been removed as it was too time-consuming to maintain it here. Please refer to [CHANGE S.html](#) for the detailed listing:

```
texdoc --list xint then select CHANGES.html
```

On the internet:

<http://mirrors.ctan.org/macros/generic/xint/CHANGES.html>

A one-page condensed `timeline` is to be found at the start of `sourcexint.pdf`.

2. Syntax reference and user guide

.1	The three parsers.....	7	arbitrary oples and nutples	37	
.2	Output customization	9	.8	Tacit multiplication.....	38
.3	Built-in operators and their precedences ..	13	.9	User defined variables.....	39
	<i>Table of precedence levels of operators</i>	15	.10	User defined functions	42
.4	Built-in functions.....	18	.11	Examples of user defined functions	49
	<i>Table of functions in expressions</i>	19	.12	Links to some (old) examples within this document.....	53
.5	Generators of arithmetic progressions	35	.13	Oples and nutples: the 1.4 terminology ...	53
.6	Python slicing and indexing of one-dimensional sequences.....	36	.14	Expansion (for geeks only)	58
.7	NumPy like nested slicing and indexing for		.15	Known bugs/features (last updated at 1.4m)	59

2.1. The three parsers

`xintexpr` provides three numerical expression parsers corresponding to these three respective tasks:

- `\xinteval`: exact evaluations with fractions, decimal fixed point numbers, numbers in scientific notation, with no size limitation,
- `\xintiieval`: evaluations allowing only integers with no size limitation,
- `\xintfloateval`: evaluations with floating points numbers according to the prevailing precision (see `\xintDigits*`),

and two secondary ones which act like the exact evaluator then round the output to a given number of fractional digits, or convert them to `false` or `true` according to whether they vanish or do not vanish.

Changed at 1.4m!

Please note the following:

- Although `\xinteval` manipulates arbitrarily long integers or fractions it also accepts scientific notation on input, as well as all the mathematical functions (evaluated using the prevailing digits precision), and (depending on customization) can thus produce also scientific notation on output.
- So far, individual operations and the printing routine of `\xinteval` do not automatically reduce fractions to their lowest terms.

The interface is:

- `\xinteval{<expression>}` handles integers, decimal numbers, numbers in scientific notation and fractions. The algebraic computations are done *exactly*, and in particular `/` simply constructs fractions. Use `//` for floored division.

```
\xinteval{add(x/(x+1), x = 1000..1014)}\par
4648482709767835886400149017599415343/310206597612274815392155150733157360
```

In this example, the fraction obtained by addition is already irreducible, but this is not always the case, as pointed above

Arbitrarily long numbers are allowed in the input. The space character (contrarily to the situation inside `\numexpr`) and also the underscore character (as allowed in Python too) can serve to separate groups of digits for better readability. But the package currently provides no macros to let the output be formatted with such separators.

```
\xinteval{123_456_789_012^5}
28679718616935524442942783005582105858543331562763768832
```

- `\xintieval[⟨D⟩]{⟨expression⟩}` is the same parser as `\xinteval`, i.e. accepts the same inputs and does all computations exactly in the same manner, but it then rounds its final result to the nearest integer, or, in case there is an optional argument `[D]`, to:
 - if `D>0`: the nearest fixed point number with `D` digits after the decimal mark,
 - if `D=0`: the nearest integer (as for `\xintieval` with no optional argument),
 - if `D<0`: the rounded quotient by $10^{(-D)}$.

Prior to 1.4k the optional argument `[⟨D⟩]` had to be located *within* the braces at the start of the expression. The legacy syntax is and will keep being allowed.

- `\xintiieval{⟨expression⟩}` executes computations on (big) integers only. It is (only slightly) faster than `\xinteval` for the same expression.

Attention: the forward slash `/` does the rounded integer division to match behaviour of `\num_2 expr`. The `//` operator does floored division as in `\xinteval`. The `/:` is the associated modulo operator (we could easily let the catcode `12 %` character be an alias, but using such an unusual percent character would be a bit cumbersome in a \TeX workflow, if only for matters of syntax highlighting in \TeX -aware text editors).

```
% add the i^5 only if i is a multiple of 7
\xintiieval{add((i/:7)?{omit}{i^5}, i=1000..1020)}
3122939154402144
```

- `\xintfloateval[⟨Q⟩]{⟨expression⟩}` does floating point computations with a given precision `P`, as specified via a prior assignment `\xintDigits:=P\relax` (the value `P` can be recovered via `\xinttheDigits`).

Its optional argument `[Q]`, if present, means to do a *final* float rounding to a mantissa of `Q` digits (this thus makes sense only if `Q<P`).

A negative `Q` is allowed and means to round to `P+Q` digits only.

Prior to 1.4k the optional argument `[⟨Q⟩]` had to be located *within* the braces at the start of the expression. The legacy syntax is and will keep being allowed.

The infix operator `/` will compute the correct rounding of the exact fraction. The operator `//` is floored division and `/:` is its associated modulo (see also `divmod()`).

```
\begingroup
\xintDigits:=64\relax
\xintfloateval{sqrt(3)}
\endgroup
1.732050807568877293527446341505872366942805253810380628055806979
```

The four basic operations and the square root achieve *correct rounding*.²

On output, `\xintfloateval` uses `\xintPFloat` for each numeric leaf. This can be modified (cf. `\xintfloatexprPrintOne`).

There is a core syntax:

- `\xintexpr⟨expression⟩\relax`,
- `\xintiexpr⟨expression⟩\relax`,
- `\xintiieval⟨expression⟩\relax`,
- `\xintfloatexpr⟨expression⟩\relax`,
- `\xintboolexpr⟨expression⟩\relax`.

² when the inputs are already floating point numbers with at most `P`-digits mantissas.

Changed at 1.4m!

`\xintboolexpr<expression>\relax` does all computations like `\xintexpr` then converts all (non-empty) leaves³ to `true` or `false` (cf. `\xintboolexprPrintOne`). There is no `\xintboolval`.⁴

Formerly the `\xintexpr... \relax` legacy syntax had to be prefixed by `\xintthe` in order to appear make it to \TeX typesetting engine, else would raise (deliberately) an error. The `\xintthe` prefix is optional since 1.4.

In an `\edef` these constructs expand to some braced nested data, all computations having been completely done, which is prefixed with some `\protected` “typesetter” macros.

In an `\edef`, `\xinteval` (in contrast to `\xintexpr`), or `\xintfloateval` (in contrast to `\xintfloatexpr`) expand the “typesetting macros” and the final complete expansion consists of explicit digits and other characters such as those of scientific notation or square brackets.⁵

In \TeX it is possible to use the core syntax `\xintexpr<expression>\relax` also in so-called moving arguments, because when written out to a file the final expansion outcome uses only standard catcodes and thus will get retokenized and expand as expected if it has been written to an external file which is then reloaded.

One needs `\xinteval` et al. only if one really wants the final digits (and other characters), for example in a context where \TeX expects a number or a dimension.

As alternative to `\xinteval{<expression>}`, an equivalent is `\xintthe\xintexpr<expression>\relax`. Similarly `\xintthe` can prefix all other core parsers. And one can also use `\xinttheexpr` as shortcut for `\xintthe\xintexpr`.

Doing exact computations with fractions leads very quickly to very big results (and furthermore one needs to use explicitly the `reduce()` function to convert the fractions into smallest terms). Thus most probably what you want is `\xintfloateval` and `\xintfloatexpr`.

2.2. Output customization

2.2.1. `\xintfloatexprPrintOne` et al. for numerical values

The package provides only minimal facilities for formatting the output from `\xinteval` or `\xintfloateval` or ... And this output may well consist of comma separated values, even nested ones with, by default, square brackets. First we explain how to influence the handling of individual “leaves”.

Here are the default definitions to this effect:

```
\def\xintexprEmptyItem{[]}\def\xintexprPrintOne #1{\xintFracToSci{#1}}\def\xintiexprPrintOne #1{\xintDecToString{#1}}\def\xintiiexprPrintOne#1{#1}\def\xintfloatevalPrintOne [#1]#2{\xintPFloat[#1]{#2}}\def\xintboolexprPrintOne#1{\xintiiifNotZero{#1}{True}{False}}
```

They can be re-defined to one's wishes. \TeX users will want to use `\renewcommand` for this.

\TeX -hackers note:

- Actually, the defaults are more done in the style

```
\let\xintexprPrintOne\xintFracToSci
```

thus sparing grabbing the argument #1. And one can do

```
\def\xintexprPrintOne{\xintFracToSci}
```

too.

- `\xintexprPrintOne` defaults in truth to some private variant of `\xintFracToSci` with exactly the same output but able to understand only certain limited types of inputs as used internally. This private variant is not *f-expandable*.

³ Currently, empty leaves are output using `\xintexprEmptyItem`, i.e. default to `[]`. This may change. ⁴ This was `True` and `False` prior to 1.4m. ⁵ `\xinteval` and `\xintexpr` both expand completely in exactly two steps. And `\xintexpr` expands fully under *f-expansion* (of the `\romannumeral 0` or `-`0` type). As per `\xinteval` attention that it may expand to nothing, then naturally *f-expansion* propagates to tokens following up in the input stream.

- `\xintfloatexprPrintOne` defaults in fact to a private variant of `\xintPFloat` which assumes the optional argument `[P]` is present as it will be the case always in this context. This optional argument `[P]` is the optional argument `[Q]` of `\xintfloateval` (or `Digis+Q` if `Q<0`).
- The typesetter for `\xintiexpr` simply prints ``as is'`, but this may change in future, if some internal format is used requiring a conversion step.

Here is a possibly not up-to-date list of macros of interest, whose documentations you might consider reading (the first two require math mode):

- `\xintTeXFromSci`,
- `\xintTeXFrac`,
- `\xintDecToString`,
- `\xintPRaw`,
- `\xintFracToSci`,
- `\xintFracToDecimal`,
- `\xintPFloat`,
- and `\xintFloatToDecimal`.

Naming scheme, as one can see, has been pretty much incoherent, apologies.

Among packages providing macros formatting numeric values, there are `numprint` and its macro `\np` (or `\numprint` without the option `np`), and `siunitx` and its `\num`, and possibly more packages not known to the author.⁶ These macros are suitable in combination with `\xintFloat` as in the example below to customize the `\xintfloateval` output. Numerical output from `\xinteval` is more challenging as individual values may naturally contain the `/` character for fractions which the above mentioned packages will not know how to handle, as far as I know.

Here an example, with \TeX and `\num` from `siunitx`:

```
\def\xintfloatexprPrintOne[#1]#2{\num{\xintFloat[#1]{#2}}}
```

We used `\def` with delimited parameters because the optional argument will always be present at time of use, and further the macro will be submitted to `\expanded`, but \TeX 2e commands with optional arguments defined via `\newcommand` are not compatible with being submitted to `\expanded`.

\TeX -hackers note: `\num` is a `\protected` macro, hence it will remain intact in the `\expanded` phase. The argument `\xintFloat[#1]{#2}` will thus expand before `\num` itself. In the case at hand `\num` is a macro which would anyhow have triggered expansion of its argument before further processing.

With the `\numprint` macro from `numprint`, one can do this:

```
\protected\def\xintfloatexprPrintOne[#1]#2{\numprint{\xintPFloat[#1]{#2}}}
```

For the `\protected` see the explanations in fine print. The other difference is usage of `\xintPFloat` rather than `\xintFloat`. This is because (in my limited testing) `\numprint` will not silently remove a zero scientific exponent but it will typeset it, for example as $1.5 \cdot 10^0$. So we use our own `\xintPFloat` poor man “prettifier”.

The current behaviour of `\xintfloateval` corresponds to this set-up:

```
\def\xintfloatexprPrintOne [#1]#2{\xintPFloat[#1]{#2}}
```

and to this default configuration of `\xintPFloat`:

```
\def\xintPFloatE{e}
\def\xintPFloatZero{0}
\def\xintPFloatIntSuffix{}
\def\xintPFloatLengthOneSuffix{}
\def\xintPFloatNoSciEmax{5}
\def\xintPFloatNoSciEmin{-4}
\def\xintPFloatMinTrimmed{4}
```

With the custom replacement

```
\def\xintfloatexprPrintOne{\xintFloatToDecimal}
```

⁶ There does not seem to be yet a \LaTeX user level interface to the `l3str-format` package, part of `l3experimental`, which provides an implementation of the Python `format` function.

the `\xintfloateval` output will use decimal fixed point notation, i.e. no scientific exponents, and as many zeros as are needed (but no more, as trailing zeros will be removed from the significant digits). Here is an example comparing outputs from the default configuration and custom ones:

```
\xintfloateval{exp(-32.456)/2000} (default, i.e. PFloat)\newline
\def\xintfloatexprPrintOne{\xintFloatToDecimal}%
\xintfloateval{exp(-32.456)/2000} (FloatToDecimal)\newline
\def\xintfloatexprPrintOne[#1]#2{\xintTeXFromSci{\xintFloat[#1]{#2}}}%
$\xintfloateval{exp(-32.456)/2000}$ (TeXFromSci on Float)\par % math mode required
4.013361680161317e-18 (default, i.e. PFloat)
0.00000000000000000004013361680161317 (FloatToDecimal)
4.013361680161317 · 10-18 (TeXFromSci on Float)
```

Some examples showing now the effect of sensible customizations on `\xinteval`:

```
\xinteval{exp(-32.456)/2000} (default, i.e. FracToSci)\newline
\def\xintexprPrintOne{\xintFracToDecimal}%
\xinteval{exp(-32.456)/2000} (FracToDecimal)\newline
\def\xintexprPrintOne#1{\xintTeXFromSci{\xintFracToSci{#1}}}%
$\xinteval{exp(-32.456)/2000}$ (TeXFromSci on FracToSci)\par % math mode required
8.026723360322633e-15/2000 (default, i.e. FracToSci)
0.00000000000000008026723360322633/2000 (FracToDecimal)
8.026723360322633 · 10-15 · 2000-1 (TeXFromSci on FracToSci)
```

Notice that the `/2000` denominator remains ``as is'' in the output, in conformity with the documented behaviour of `\xintFracToSci` in the first example and of `\xintFracToDecimal` for the second example. This has not changed since 1.4 (the handling of the numerator part has changed at 1.4e and again slightly at 1.4k, the zero value being now always printed as 0 and not 0 or 0.0 depending on the input) but is to be considered unstable and undecided so far.

A slightly more costly typesetter could be for example:

```
\def\xintexprPrintOne#1{\xintDecToStringREZ{\xintIrr{#1}}}
```

Then

- the fraction (inclusive of its power of ten part) will be reduced to lowest terms (see `\xintIrr`),
- next the trailing zeros will be moved as an exponent (positive or negative) to the numerator,
- this numerator with a power of ten part will be printed in decimal fixed point notation, with as few zeros as are needed,
- and finally the denominator `B`, which has been trimmed of trailing zeros, will be printed as `/B` or not at all if `B=1`.

With the use case above:

```
\def\xintexprPrintOne#1{\xintDecToStringREZ{\xintIrr{#1}}}%
\xinteval{exp(-32.456)/2000}\par
0.00000000000000008026723360322633/2
```

This trailing `/2` is somewhat of a pain, but as documented and mentioned already `\xintDecToStringREZ` currently has not been educated to identify its presence and handle it. Slightly faster (see `\xintPIrr`) is

```
\def\xintexprPrintOne#1{\xintDecToStringREZ{\xintPIrr{#1}}}
```

which with the used example produces the same output.

One can also consider this for math mode:

```
\def\xintexprPrintOne#1{\xintTeXFromSci{\xintDecToStringREZ{\xintIrr{#1}}}%
$\xinteval{exp(-32.456)/2000}$\par

$$\frac{0.00000000000000008026723360322633}{2}$$

```

See our hesitations about what `\xintTeXFromSci` should do with denominators.

TeX-hackers note: One can hope that in future `\xintDecToString` will identify denominators being products of only two's and five's, but even then of course `\xintTeXFromSci` will have to decide how to handle other denominators.

T_EX-hackers note: The macro used as customization of `\xintexprPrintOne` (whose default is a private variant of `\xintFracToSci` with exactly same output) must understand the internal `xintfrac` format `A/B[N]`, but with the `/B` and `[N]` parts being only optional. This is not a problem when using for this task (nested) macros of `xintfrac`, as they of course accept such inputs as argument and in fact much more general ones.

In particular one can benefit from `\xintRaw`, or `\xintRawBraced`, to convert the argument into a well defined shape (`A/B[N]` for the former and `{N}{A}{B}` for the latter) and then work from there.

The macro used by `\xintfloatexprPrintOne` has the guarantee that the `[P]` will be always present at expansion time.

The customization should be compatible with being exposed to `\expanded` (which is like expansion in an `\edef`), either from being completely expandable or at the opposite from being `\protected`. \TeX commands defined via `\newcommand` as macros with one optional parameter are not compatible with this requirement.

Attention! The interface requirements described above for the macros customizing the behaviours of `\xintexprPrintOne` and `\xintfloatexprPrintOne` may change at any release... as they depend on some internal structures and it is not certain backwards compatibility will be maintained systematically in case of evolution.

2.2.2. `\xintthealign` for output of general oples

With `\xintthealign` one can get nested data use a T_EX alignment in the output. Here is an example :

```
\xintthealign\xintexpr ndseq(1/(i+j), i = 1..10; j=1..10)\relax
[[ 1/2, 1/3, 1/4, 1/5, 1/6, 1/7, 1/8, 1/9, 1/10, 1/11 ],
 [ 1/3, 1/4, 1/5, 1/6, 1/7, 1/8, 1/9, 1/10, 1/11, 1/12 ],
 [ 1/4, 1/5, 1/6, 1/7, 1/8, 1/9, 1/10, 1/11, 1/12, 1/13 ],
 [ 1/5, 1/6, 1/7, 1/8, 1/9, 1/10, 1/11, 1/12, 1/13, 1/14 ],
 [ 1/6, 1/7, 1/8, 1/9, 1/10, 1/11, 1/12, 1/13, 1/14, 1/15 ],
 [ 1/7, 1/8, 1/9, 1/10, 1/11, 1/12, 1/13, 1/14, 1/15, 1/16 ],
 [ 1/8, 1/9, 1/10, 1/11, 1/12, 1/13, 1/14, 1/15, 1/16, 1/17 ],
 [ 1/9, 1/10, 1/11, 1/12, 1/13, 1/14, 1/15, 1/16, 1/17, 1/18 ],
 [ 1/10, 1/11, 1/12, 1/13, 1/14, 1/15, 1/16, 1/17, 1/18, 1/19 ],
 [ 1/11, 1/12, 1/13, 1/14, 1/15, 1/16, 1/17, 1/18, 1/19, 1/20 ]]
```

Attention, this `\xintthealign` must be a prefix to `\xintexpr`, or `\xintfloatexpr` etc..., but there will be low-level T_EX errors if it is used to prefix `\xinteval` et al. or `\xinttheexpr` et al.

It is possible to customize the behaviour of `\xintthealign`. For example:

```
\protected\def\xintexpralignbegin    {\halign\bgroup\tabskip2ex\hfil##&&##\hfil\cr}%
\def\xintexpralignend                {\crcr\egroup}% removed \protected at 1.4c
\protected\def\xintexpralignlinesep  {\,\cr}% separates "lines"
\protected\def\xintexpralignleftsep  {\&}% at left of first item in a "line"
                                        % (after "left bracket")
\protected\def\xintexpraligninnersep  {\,&}% at the left of non-first items
\protected\def\xintexpralignrightsep  {\&}% at right of last item in a "line"
                                        % (before "right bracket")

\protected\def\xintexpralignleftbracket  {\[%
\protected\def\xintexpralignrightbracket}{\%]
```

The above definitions use `\protected` with no strong reason, as the replacement tokens are not expanding anyhow, but the idea is that this allows to execute a computation via an `\edef` and later one can change the meaning of the auxiliary macros depending on what one wants to do with the expansion result.

T_EX-hackers note: `\xintexpralignend` is expanded once, after the body has been submitted to exhaustive expansion (`\expanded` induced), and prior to the expansion of `\xintexpralignbegin`.

Although we will try to keep stable the way "regular arrays" as in the above example are rendered by default, the `\xintthealign` macro (and its associated customizability) is to be considered work-in-progress and may experience breaking changes.

Use for example this for outputting to a file or a terminal:⁷

```
% Better here without \protected.
% We assume here \newlinechar has the LaTeX setting.
\def\xintexpralignbegin    {}%
\def\xintexpralignend      {}%
\def\xintexpralignlinesep  {,^^}% separates "lines"
\def\xintexpralignleftsep  { }% at left of first item in a "line" (after brackets)
\def\xintexpraligninnersep {, }% at the left of non-first items
\def\xintexpralignrightsep { }% at right of last item in a "line" (before brackets)
\def\xintexpralignleftbracket {[}%
\def\xintexpralignrightbracket {]}%
```

In the \TeX example next using a `pmatrix` environment, `\noexpand` rather than `\protected` is used. This environment will not break across pages, contrarily to the display produced by the default `\xintthealign` configuration which uses \TeX 's `\halign`.

```
\[
\def\xintexpralignbegin    {\begin{pmatrix}}%
\def\xintexpralignend      {\end{pmatrix}}%
\def\xintexpralignlinesep  {\noexpand\\}% needed to counteract an internal \expanded
\def\xintexpraligninnersep  {&}%
\let\xintexpralignleftbracket\empty \let\xintexpralignleftsep\empty
\let\xintexpralignrightbracket\empty \let\xintexpralignrightsep\empty
% by default amsmath matrices can have 10 columns at most
% (cf amsmath documentation for what to do to allow more)
l.c.m.=\xintthealign\xintiexpr ndmap(lcm, 1..12; 1..10)\relax
\]
```

$$\text{l.c.m.} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ 2 & 2 & 6 & 4 & 10 & 6 & 14 & 8 & 18 & 10 \\ 3 & 6 & 3 & 12 & 15 & 6 & 21 & 24 & 9 & 30 \\ 4 & 4 & 12 & 4 & 20 & 12 & 28 & 8 & 36 & 20 \\ 5 & 10 & 15 & 20 & 5 & 30 & 35 & 40 & 45 & 10 \\ 6 & 6 & 6 & 12 & 30 & 6 & 42 & 24 & 18 & 30 \\ 7 & 14 & 21 & 28 & 35 & 42 & 7 & 56 & 63 & 70 \\ 8 & 8 & 24 & 8 & 40 & 24 & 56 & 8 & 72 & 40 \\ 9 & 18 & 9 & 36 & 45 & 18 & 63 & 72 & 9 & 90 \\ 10 & 10 & 30 & 20 & 10 & 30 & 70 & 40 & 90 & 10 \\ 11 & 22 & 33 & 44 & 55 & 66 & 77 & 88 & 99 & 110 \\ 12 & 12 & 12 & 12 & 60 & 12 & 84 & 24 & 36 & 60 \end{pmatrix}$$

2.3. Built-in operators and their precedences

The parser implements precedence rules based on concepts which are summarized below (only for binary infix operators):

- an infix operator has two associated precedence levels, say **L** for left and **R** for right,
- the parser proceeds from left to right, pausing each time it has found a new number and an operator following it,
- the parser compares the left-precedence **L** of the new found operator to the right-precedence **R_{last}** of the last delayed operation (which already has one argument and would like to know if it can use the new found one): if **L** is at most equal to it, the delayed operation is now executed, else the new-found operation is kept around to be executed first, once it will have gathered its arguments, of which only one is known at this stage.

⁷ With the `xetex` engine this will need its `-8bit` option else the `^^` in `\xintexpralignlinesep` will be printed literally instead of being converted into a line separator in the file or terminal output.

This means for example in the case of the multiplication `*` and the division operators `/`, `//`, `/:` that they are parsed in a left-associative way because they all share the same (left and right) precedence level. This is the case with the analogous operators from the Python language, as well.

At 1.4g the power operators were changed to act in a right associative way. Again, this matches the behaviour of e.g. Python:

```
\xinteval{2^3^4}
1/2417851639229258349412352
```

The entries of Table 1 are hyperlinked to the more detailed discussion at each level. In these entries the number within parentheses indicates the right-precedence, if it differs from the left.

∞ At this highest level of precedence, one finds:

functions and variables Functions (even the logic functions `!()` and `?()` whose names consist of a single non-letter character) must be used with parentheses. These parentheses may arise from expansion after the function name is parsed (there are exceptions which are documented at the relevant locations.)

* Python-like “unpacking” prefix operator. Sometimes one needs to use it as function `*()` (but I can't find an example right now) but most of the time parentheses are unneeded.

. is decimal mark; the number scanner treats it as an inherent, optional and unique component of a being formed number. `\xintexpr 0.^2+2^.0\relax` is interpreted as `0^2+2^0` and thus produces 1.

Since release 1.2 an isolated decimal mark is illegal input in the `xintexpr` parsers (it remains legal as argument to the macros of `xintfrac`).

e scientific notation.

E scientific notation. For output, see `\xintPFloatE`.

" prefix for hexadecimal input. Only uppercase letters, and one optional . separating integer and fractional hexadecimal parts. This functionality

requires to load explicitly package `xintbinhex`.

```
\xintexpr "FEDCBA9876543210\relax\newline
\xintexpr ".FEDCBA9876543210\relax\newline
\xintexpr 16^5-("F75DE.0A8B9+"8A21.F5746+16^-5)\relax
18364758544493064720
0.9955555555555555555555559410496613281793543137609958648681640625
0
```

It is possible that in future the " prefix could be dropped in favour of `0x` prefix. This would free " to be used for input of “string”-like entities.

20 The postfix operators `!` and the branching conditionals `?`, `??`.

`!` computes the factorial of an integer.

`?` is used as `(stuff)?{yes}{no}`. It evaluates `stuff` and chooses the `yes` branch if the result is non-zero, else it executes `no`. After evaluation of `stuff` it acts as a macro with two mandatory arguments within braces, chooses the correct branch *without evaluating the wrong one*. Once the braces are removed, the parser scans and expands the uncovered material.

`??` is used as `(stuff)??{<0}{=0}{>0}`, where `stuff` is anything, its sign is evaluated and depending on the sign the correct branch is un-braced, the two others are discarded with no evaluation of their contents.

- As unary operator, the minus sign inherits as precedence the minimum of 12 (which is the precedence for addition and subtraction) and of the (right-) precedence of the operators preceding it (if any).

<p>∞: at this top level the syntax elements whose execution is done prior to operators preceding them:</p> <ul style="list-style-type: none"> • built-in or user-defined functions, • variables, • the * unpacking operator, • and intrinsic constituents of numbers: decimal mark ., e and E of scientific notation, hexadecimal prefix ". 	
Precedence	``Operators'' at this level
20	postfix ! and branching ? , ?? operators
-	minus sign as unary operator inherits the right-precedence of the infix operator it follows, if that precedence is higher than the one of binary + and - , else it inherits the latter
18 (17)	^ and ** are synonymous; they act in a right-associative way (Changed at 1.4g!)
16 (14)	Tacit multiplication has an elevated (left) precedence
14	* , / , // (floored division), and /: (associated modulo, alias 'mod')
12	+ , -
10	< , > , == , <= , >= , != (they can be chained)
8	Boolean conjunction && and its alias 'and'
6	Boolean disjunction and its alias 'or' . Also 'xor' and .. , ..[,].. , and : have this precedence
4	the brackets for slicers and extractors [,]
3	the comma ,
2	the bracketers [,] construct nestable "arrays"
1	the parentheses (,) , and the semi-colon ; in iter() , rseq() , and further structures
<ul style="list-style-type: none"> • Binary operators have a left and a right precedence, which for most coincide. The right precedence is indicated within parentheses. • Tacit multiplication has an elevated left precedence level: $(1+2)/(3+4)5$ is computed as $(1+2)/((3+4)*5)$ and $x/2y$ is interpreted as $x/(2*y)$ when using variables. 	

Table 1: Precedence levels

```
\xintexpr -3-4*-5^7, (-3)-(4*(-(5^(-7))))\relax\newline
\xintexpr -3^4*-5-7, (-((3^(-4))*(-5)))-7\relax\newline
|2^-10| gives \xintexpr 2^-10\relax\space
-234371/78125, -234371/78125
-562/81, -562/81
```

2^{-10} gives $1/1024$ and is thus perfectly legal, no need for parentheses.

The + character as prefix unary operator is simply ignored during input parsing.

18

^

** Both compute powers. They act in a right associative way.

```
\xintiexpr 2^3^4\relax
2417851639229258349412352
```

16 see [Tacit multiplication](#).

14

* multiplication

/ division:

- in `\xinteval`: exact division in the field of rational numbers (not automatically reduced to lowest terms),
- in `\xintfloateval`: correct rounding of the exact division; the two operands are, if necessary, float-rounded before the fraction is evaluated and rounded (to obtain the correctly rounded A/B without prior rounding of A and B see `qfloat()`),
- in `\xintiieval`: for compatibility with the legacy behaviour of `/` in `\numexpr`, it rounds the exact fraction *with half-integers going towards the infinity of the same sign*.

The division is left-associative. Example:

```
\xintexpr reduce(100/50/2)\relax
1
```

// floored division (and thus produces an integer, see `divmod()` for details)

/: the associated modulo (see `divmod()` and `mod()`)

Left-associativity applies to the division operators:

```
\xintexpr 100000/:13, 100000 'mod' 13\relax, \xintexpr 100000/:13/13\relax
4, 4, 4/13
```

Nothing special needs to be done in contexts such as $\text{\TeX 3 \ExplSyntaxOn}$ where `:` is of cat-code letter, but if `:` is an active character (for example in \TeX with `babel+french`) one needs to use input such as `\string :` (or replace it with usage of the function `mod()`).

'mod' is same as `/:`.



Attention: with `polexpr` loaded, which allows ' in variable and function names, 'mod' can not follow a variable name. Add parentheses around the variable, or use `/:`.

12

+ addition

- subtraction. According to the general left-associativity rule in case of equal precedence, it is left associative:

```
\xintiexpr 100-50-2\relax
48
```


10 Comparison operators are (as in Python) all at the same level of precedence, use parentheses for disambiguation.

< `a<b` evaluates to 1 if the strict inequality holds to 0 if not.

> `a>b` evaluates to 1 if the strict inequality holds to 0 if not.

== `a==b` evaluates to 1 if equality holds to 0 if not.

<= `a<=b` evaluates to 1 if left hand side is at most equal to right hand side, to 0 if not.

>= `a>=b` evaluates to 1 if left hand side is at least equal to right hand side, to 0 if not.

!= `a!=b` evaluates to 1 if they differ, to 0 if not.


Comparisons can be chained arbitrarily, e.g., `x < y <= z != t` is equivalent to `x < y 'and' y <= z 'and' z != t` (and also to `all(x<y, y<=z, z!=t)`), except that if `y` and `z` involve computations, they are evaluated only once. Currently there is no short-circuit here, i.e. even if some intermediate comparison turns out false (in fact 0), all the remaining conditionals will still be evaluated.

```
\xintifboolexpr{1<=2!=3<4>1}{true}{\error},
\xintifboolexpr{1<=2>=3<4>1}{\error}{false},
\xintifboolexpr{3 != 3! == 6 != 4! == 24}{true}{\error}
true, false, true
```

8

&& logical conjunction. Evaluates to 1 if both sides are non-zero, to 0 if not.


'and' same as **&&**. See also the [all\(\)](#) multi-arguments function.

 { **Attention:** with [polexpr](#) loaded, which allows ' in variable and function names, 'and' can not follow a variable name. Add parentheses around the variable, or use **&&**.


6

|| logical (inclusive) disjunction. Evaluates to 1 if one or both sides are non-zero, to 0 if not.

'or' same as **||**. See also the [any\(\)](#) multi-arguments function.

 { **Attention:** with [polexpr](#) loaded, which allows ' in variable and function names, 'or' can not follow a variable name. Add parentheses around the variable, or use **||**.

'xor' logical (exclusive) disjunction.

 { **Attention:** with [polexpr](#) loaded, which allows ' in variable and function names, 'xor' can not follow a variable name. Add parentheses around the variable, or use the [xor\(\)](#) function syntax.

..

..[

].. Syntax for arithmetic progressions. See [subsection 2.5](#).

: This is a separator involved in `[a:b]` Python-like slicing syntax.

4

[

] Involved in Python-like slicing `[a:b]` and extracting `[N]` syntax. And its extension à la NumPy `[a:b,N,c:d,...,:]`. Ellipsis `...` is not yet implemented. The "step" parameter as in `[a:b:step]` is not yet implemented.

3

, The comma separates expressions (or function arguments).⁸

```
\xintiexpr 2^3,3^4,5^6\relax
8, 81, 15625
```

2

[

] The bracketers construct nestable “array-like” structures. Arbitrary (heterogeneous) nesting is allowed. For output related matters see [\xintthealign](#) (its usage is optional, without it rendering is “one-dimensional”). Output shape of non-homogeneous arrays is to be considered unstable at this time.

1

(

) The parentheses serve as mandatory part of the syntax for functions, and to disambiguate precedences.⁹ They do not construct any nested structure.

; The semi-colon as involved as part of the syntax of [iter\(\)](#), [rseq\(\)](#), [ndseq\(\)](#), [ndmap\(\)](#) has the same precedence as a closing parenthesis.

\relax This is the expression terminator for [\xintexpr](#) et al. It may arise from expansion during the parsing itself. As alternative to [\xintexpr](#) (et al.) use [\xinteval](#) (et al.) which have the usual macro interface (with one mandatory argument).

The **;** also serves as syntax terminator for [\xintdefvar](#) and [\xintdeffunc](#). It can in this rôle not arise from expansion as the expression body up to it is fetched by a delimited macro. But this is done in a way which does not require any specific hiding for inner semi-colons as involved in the syntax of [iter\(\)](#), etc...

2.4. Built-in functions

See [Table 2](#) whose elements are hyperlinked to the corresponding definitions.

Functions are at the same top level of priority. All functions even [?\(\)](#) and [!\(\)](#) require parentheses around their arguments.

Miscellaneous notes:

- since release [1.3d](#) [gcd\(\)](#) and [lcm\(\)](#) are extended to apply to fractions too, and do NOT require the loading of [xintgcd](#),
- The randomness related functions [random\(\)](#), [qrand\(\)](#) and [randrange\(\)](#) require that the \TeX engine provides the [\uniformdeviate](#) or [\pdfuniformdeviate](#) primitive. This is currently the case for [pdftex](#), [\(u\)ptex](#), [luatex](#), and also for [xetex](#) since \TeX Live 2019.
- [togl\(\)](#) is provided for the case [etoolbox](#) package is loaded,
- [bool\(\)](#), [togl\(\)](#) use delimited macros to fetch their argument and the closing parenthesis must be explicit, it can not arise from on the spot expansion. The same holds for [qint\(\)](#), [qfrac\(\)](#), [qfloat\(\)](#), [qraw\(\)](#), [random\(\)](#) and [qrand\(\)](#).
- Also [functions with dummy variables](#) use delimited macros for some tasks. See the relevant explanations there.
- Functions may be called with *oples* as arguments as long as the total length is the number of arguments the function expects.



⁸ The comma is really like a binary operator, which may be called “join”. It has lowest precedence of all (apart the parentheses) because when it is encountered all postponed operations are executed in order to finalize its *first* operand; only a new comma or a closing parenthesis or the end of the expression will finalize its *second* operand. ⁹ It is not apt to describe the opening parenthesis as an operator, but the closing parenthesis is analogous to a postfix unary operator. It has lowest precedence which means that when it is encountered all postponed operations are executed to finalize its operand. The start of this operand was decided by the opening parenthesis.

!()	atan2()	first()	iter()	num()	rbit()	subs()
?()	atan2d()	flat()	iterr()	nuple()	reduce()	subsm()
`*`()	binomial()	float()	inv()	odd()	reversed()	subsn()
`+`()	bool()	float_dgt()	last()	pArg()	round()	tan()
abs()	ceil()	floor()	lcm()	pArgd()	rrseq()	tand()
add()	cos()	frac()	len()	pfactorial()	rseq()	tg()
all()	cosd()	gcd()	log()	pow()	sec()	togl()
any()	cot()	if()	log10()	pow10()	secd()	trunc()
acos()	cotd()	ifint()	max()	preduce()	seq()	unpack()
acosd()	cotg()	ifone()	min()	qfloat()	sgn()	xor()
Arg()	csc()	ifsgn()	mod()	qfrac()	sin()	zip()
Argd()	cscd()	ilog10()	mul()	qint()	sinc()	
asin()	divmod()	iquo()	ndmap()	qrand()	sind()	
asind()	even()	irem()	ndseq()	qraw()	sqr()	
atan()	exp()	isint()	ndfillraw()	random()	sqrt()	
atand()	factorial()	isone()	not()	randrange()	sqrtr()	

Table 2: Functions (click on names)

.4.1	Functions with no argument	19
.4.2	Functions with one argument	20
.4.3	Functions with an alphanumeric argument	23
.4.4	Functions with one mandatory and a second but optional argument	24
.4.5	Functions with two arguments	25
.4.6	Functions with 3 or 4 arguments	28
.4.7	Functions with an arbitrary number of arguments	28
.4.8	Functions requiring dummy variables	30

2.4.1. Functions with no argument

random() returns a random float x verifying $0 \leq x < 1$. It obeys the prevailing precision as set by `\xintDigits`: i.e. with P being the precision the random float multiplied by 10^P is an integer, uniformly distributed in the $0..10^P-1$ range.

This description implies that if x turns out to be <0.1 then its (normalized) mantissa has $P-1$ digits and a trailing zero, if $x < 0.01$ it has $P-2$ digits and two trailing zeros, etc... This is what is observed also with Python's `random()`, of course with `10` replaced there by radix `2`.

```
\pdfsetrandomseed 12345
\xintDigits:=37\relax
\xintthefloatexpr random()\relax\newline
\xintthefloatexpr random()\relax\par
0.2415544817596207455547929850209500042
0.2584863529993996627285461554203021352
```

qrand() returns a random float $0 \leq x < 1$ using 16 digits of precision (i.e. $10^{16}x$ is an integer). This is provided when speed is a at premium as it is optimized for precision being precisely 16.

```
% still with 37 digits as prevailing float precision
\xintthefloatexpr qrand(), random()\relax\newline
\xintDigits:=16\relax
\xintthefloatexpr qrand(), random()\relax\par
0.4883568991327765, 0.09165461826072383107532471669335645230
0.9069127435402274, 0.9106687541716861
```

One can use both `qrand()` and `random()` inside the `\xintexpr` parser too. But inside the integer only `\xintiexpr` parser they will cause some low-level error as soon as they get involved in any kind of computation as they use an internal format not recognized by the integer-only parser.

See further `randrange()`, which generates random integers.

Currently there is no `uniform()` function¹⁰ but it can be created by user:

```
\xintdeffloatfunc uniform(a, b):= a + (b-a)*random();
\romannumeral\xintreplicate{10}%
{%
\xintthefloatexpr uniform(123.45678, 123.45679)\relax\newline
}%
```

```
123.4567849497100
123.4567812033226
123.4567863308250
123.4567896366777
123.4567849656655
123.4567849908270
123.4567889123433
123.4567896262979
123.4567846543719
123.4567832664043
```

`rbit()` returns a random 0 or 1.

2.4.2. Functions with one argument

`num(x)` truncates to the nearest integer (truncation towards zero). It has the same sign as `x`, except of course with `-1<x<1` as then `num(x)` is zero.

```
\xinttheexpr num(3.1415^20), num(1e20)\relax
```

```
8764785276, 10000000000000000000000
```

The output is an explicit integer with as many zeros as necessary. Even in float expressions, there will be an intermediate stage where all needed digits are there, but then the integer is immediately reparsed as a float to the target precision, either because some operation applies to it, or from the output routine of `\xintfloatexpr` if it stood there alone. Hence, inserting something like `num(1e10000)` is costly as it really creates ten thousand zeros, even though later the whole thing becomes a float again. On the other hand naturally `1e10000` without `num()` would be simply parsed as a floating point number and would cause no specific overhead.

`frac(x)` fractional part. For all numbers `x=num(x)+frac(x)`, and `frac(x)` has the same sign as `x` except when `x` is an integer, as then `frac(x)` vanishes.

```
\xintthefloatexpr frac(-355/113), frac(-1129.218921791279)\relax
```

```
-0.1415929203539820, -0.218921791279
```

`reduce(x)` reduces a fraction to smallest terms

```
\xinttheexpr reduce(50!/20!/20!/10!)\relax
```

```
1415997888807961859400
```

Recall that this is NOT done automatically, for example when adding fractions.

¹⁰ Because I am not sure how to handle rounding issues: should the computation proceed exactly and a rounding be done only at very end?

preduce(x) internally, fractions may have some power of ten part (for example when they got input in scientific notation). This function ignores the decimal part when doing the reduction. See [\xintPIrr](#).

```
\xinttheexpr preduce(10e7/2), reduce(10e7/2)\relax
5e7, 50000000
```

abs(x) absolute value

sgn(x) sign. See also [\xintifsgnexpr](#).

inv(x) inverse.

floor(x) floor function.

ceil(x) ceil function.

sqr(x) square.

ilog10(x) in [\xintiexpr](#) the integer exponent a such that $10^a \leq \text{abs}(x) < 10^{a+1}$; returns (this may evolve in future) `-2147450880` if x vanishes (i.e. `0x7fff8000`).

```
\xintiieval{ilog10(1), ilog10(-1234567), ilog10(-123456789123456789), ilog10(2**31)}\par
0, 6, 17, 9
```

See [ilog10\(\)](#) for the behaviour in [\xintexpr](#)-essions.

sqrt(x) in [\xintiexpr](#), truncated square root; in [\xintexpr](#) or [\xintfloatexpr](#) this is the floating point square root, and there is an optional second argument for the precision. See [sqrt\(\)](#).

sqrtr(x) available *only* in [\xintiexpr](#), rounded square root.

factorial(x) factorial function (like the post-fix `!` operator.) When used in [\xintexpr](#) or [\xintfloatexpr](#) there is an optional second argument. See [factorial\(\)](#).

?(x) is the truth value, 1 if non zero, 0 if zero. Must use parentheses.

!(x) is logical not, 0 if non zero, 1 if zero. Must use parentheses.

not(x) logical not.

even(x) is the evenness of the truncation [num\(x\)](#).

```
\xintfloatexpr [3] seq((x,even(x)), x=-5/2..[1/3]..+5/2)\relax
-2.50, 1, -2.17, 1, -1.83, 0, -1.50, 0, -1.17, 0, -0.833, 1, -0.500, 1, -0.167, 1, 0.167, 1,
0.500, 1, 0.833, 1, 1.17, 0, 1.50, 0, 1.83, 0, 2.17, 1, 2.50, 1
```

odd(x) is the oddness of the truncation [num\(x\)](#).

```
\xintfloatexpr [3] seq((x,odd(x)), x=-5/2..[1/3]..+5/2)\relax
-2.50, 0, -2.17, 0, -1.83, 1, -1.50, 1, -1.17, 1, -0.833, 0, -0.500, 0, -0.167, 0, 0.167, 0,
0.500, 0, 0.833, 0, 1.17, 1, 1.50, 1, 1.83, 1, 2.17, 0, 2.50, 0
```

isint(x) evaluates to 1 if x is an integer, to 0 if not. See [ifint\(\)](#).

```
\xinttheexpr -5/3..[1/3]..+5/3\relax
\rightarrow \xinttheexpr seq(isint(x), x=-5/3..[1/3]..+5/3)\relax$
-5/3, -4/3, -3/3, -2/3, -1/3, 0, 1/3, 2/3, 3/3, 4/3, 5/3 → 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0
```

isone(x) evaluates to 1 if x is 1, to 0 if not. See [ifone\(\)](#).

```
\xintfloatexpr subs(((x-1)/x, x/x, (x+1)/x), x=2**30)\relax
```

```
\rightarrow
\xintthefloatexpr seq(isone(y), y=subs(((x-1)/x, x/x, (x+1)/x), x=2**30))\relax$
0.9999999990686774, 1, 1.000000000931323 → 0, 1, 0
```

qint(x) belongs with `qfrac()`, `qfloat()`, `qraw()` to a special category:

1. They require the closing parenthesis of their argument to be immediately visible, it can not arise from expansion.
2. They grab the argument and store it directly; the format must be compatible with what is expected at macro level.
3. And in particular the argument can not be a variable, it has to be numerical.

`qint()` achieves the same result as `num`, but the argument is grabbed as a whole without expansion and handed over to the `\xintiNum` macro. The `q` stands for ``quick'', and `qint` is thought out for use in `\xintiexpr...\relax` with integers having dozens of digits.

Testing showed that using `qint()` starts getting advantageous for inputs having more (or *f-expanding* to more) than circa 20 explicit digits. But for hundreds of digits the input gain becomes a negligible proportion of (for example) the cost of a multiplication.

Leading signs and then zeroes will be handled appropriately but spaces will not be systematically stripped. They should cause no harm and will be removed as soon as the number is used with one of the basic operators. This input mode *does not accept decimal part or scientific part*.

```
\def\x{...many many many ... digits}\def\y{...also many many many digits...}
\xinttheiexpr qint(\x)*qint(\y)+qint(\y)^2\relax\par
```

qfrac(x) does the same as `qint` except that it accepts fractions, decimal numbers, scientific numbers as they are understood by the macros of package `xintfrac`. Thus, it is for use in `\xintexpr...\relax`. It is not usable within an `\xintiexpr`-ession, except if hidden inside functions such as `round` or `trunc` which then produce integers acceptable to the integer-only parser. It has nothing to do with `frac` (sigh...).

qfloat(x) does the same as `qfrac` and then converts to a float with the precision given by the setting of `\xintDigits`. This can be used in `\xintexpr` to round a fraction as a float with the same result as with the `float()` function (whereas using `\xintfloatexpr A/B\relax` inside `\xintexpr...\relax` would first round `A` and `B` to the target precision); or it can be used inside `\xintfloatexpr...\relax` as a faster alternative to wrapping the fraction in a sub-`\xintexpr`-ession. For example, the next two computations done with 16 digits of precision do not give the same result:

```
\xintthefloatexpr qfloat(12345678123456785001/12345678123456784999)-0.5\relax\newline
\xintthefloatexpr 12345678123456785001/12345678123456784999-0.5\relax\newline
\xintthefloatexpr 1234567812345679/1234567812345678-0.5\relax\newline
\xintthefloatexpr \xintexpr12345678123456785001/12345678123456784999\relax-0.5\newline
0.5
0.5000000000000010
0.5000000000000010
0.5
```

because the second is equivalent to the third, whereas the first one is equivalent to the fourth one. Equivalently one can use `qfrac` to the same effect (the subtraction provoking the rounding of its two arguments before further processing.)

Note that if the input needs no special rounding, the internal form of the output keeps a short mantissa (it does not add padding zeros to make it of length equal to the float precision). For example `qfloat(2[20])` would keep internally the input format.

TOC, Start here, [xintexpr](#), [xintexpr \(old doc\)](#), [xinttrig](#), [xintlog](#), [xinttools](#), [Examples](#), [xint bundle](#)

float_dgt(x) is like `float()` and avoids `float()`'s check whether it used with its second optional argument. This is useful in the context of converting function definitions done via `\xint-deffunc` (see explanations there) to functions usable in `\xintfloateval`.

Do not use! (1.4)

nuple(x) is currently same as `[...]`. Reserved for possible alternative meaning in future.

```
\xinteval{nuple(1,2,3)}  
[1, 2, 3]
```

unpack(x) is alternative for `*` unpacking operator.

```
\xinteval{unpack([1,2,3])}  
1, 2, 3
```

flat(ople) removes all nesting to produce a (non-bracketed) ople having the same leaves (some possibly empty) but located at depth 1.

```
\xinteval{flat([[[[1, [], 3], [4, [[5, 6, []], [8, 9], [], 11]], 12], [13, 14]], [], [16]]], [])}  
1, [], 3, 4, 5, 6, [], 8, 9, [], 11, 12, 13, 14, [], 16, []
```

unstable?

I almost delayed indefinitely release because I was hesitating on the name: perhaps better with `flattened()`, but long names add (negligible, but still) overhead compared to short names. For this reason, consider that name may change.

2.4.3. Functions with an alphanumeric argument

bool(name) returns 1 if the \TeX conditional `\ifname` would act as `\iftrue` and 0 otherwise. This works with conditionals defined by `\newif` (in \TeX or \LaTeX) or with primitive conditionals such as `\ifmmode`. For example:

```
\xintifboolexpr{25*4-if(mmode),100,75}{YES}{NO}
```

will return NO if executed in math mode (the computation is then $100-100 = 0$) and YES if not (the `if()` conditional is described below; the `\xintifboolexpr` test automatically encapsulates its first argument in an `\xintexpr` and follows the first branch if the result is non-zero (see subsection 3.14)).

The alternative syntax `25*4-\ifmmode100\else75\fi` could have been used here, the usefulness of `bool(name)` lies in the availability in the `\xintexpr` syntax of the logic operators of conjunction `&&`, inclusive disjunction `||`, negation `!` (or `not`), of the multi-operands functions `all`, `any`, `xor`, of the two branching operators `if` and `ifsgn` (see also `?` and `??`), which allow arbitrarily complicated combinations of various `bool(name)`.

togl(name) returns 1 if the \LaTeX package `etoolbox`¹¹ has been used to define a toggle named `name`, and this toggle is currently set to `true`. Using `togl` in an `\xintexpr..relax` without having loaded `etoolbox` will result in an error from `\iftoggle` being a non-defined macro. If `etoolbox` is loaded but `togl` is used on a name not recognized by `etoolbox` the error message will be of the type `ERROR: Missing \endcsname inserted.'`, with further information saying that `\protect` should have not been encountered (this `\protect` comes from the expansion of the non-expandable `etoolbox` error message).

When `bool` or `togl` is encountered by the `\xintexpr` parser, the argument enclosed in a parenthesis pair is expanded as usual from left to right, token by token, until the closing parenthesis is found, but everything is taken literally, no computations are performed. For example `togl(2+3)` will test the value of a toggle declared to `etoolbox` with name `2+3`, and not `5`. Spaces are gobbled in this process. It is impossible to use `togl` on such names containing spaces, but `\iftoggle{name with spaces}{1}{0}` will work, naturally, as its expansion will pre-empt the `\xintexpr` scanner.

¹¹ <https://ctan.org/pkg/etoolbox>

There isn't in `\xintexpr...` a `test` function available analogous to the `test{\ifsometest}` construct from the `etoolbox` package; but any expandable `\ifsometest` can be inserted directly in an `\xintexpr`-ession as `\ifsometest10` (or `\ifsometest{1}{0}`), for example `if(\ifsometest{1}{0},YES,NO)` (see the `if` operator below) works.

A straight `\ifsometest{YES}{NO}` would do the same more efficiently, the point of `\ifsometest 10` is to allow arbitrary boolean combinations using the (described later) `&&` and `||` logic operators: `\ifsometest10 && \ifsomeothertest10 || \ifsomehirdtest10`, etc... `YES` or `NO` above stand for material compatible with the `\xintexpr` parser syntax.

See also `\xintifboolexpr`, in this context.

2.4.4. Functions with one mandatory and a second but optional argument

round(x[, n]) Rounds its first argument to an integer multiple of 10^{-n} (i.e. it *quantizes*). The case of negative `n` is new with 1.4a. Positive `n` corresponds to conversion to a fixed point number with `n` digits after decimal mark.

```
\xinteval{round(-2^30/3^5,12), round(-2^30/3^5,-3)}
-4418690.633744855967, -4419e3
```

trunc(x[, n]) Truncates its first argument to an integer multiple of 10^{-n} . The case of negative `n` is new with 1.4a.

```
\xinteval{trunc(-2^30/3^5,12), trunc(-2^30/3^5,-3)}
-4418690.633744855967, -4418e3
```

float(x[, n]) Rounds its first argument to a floating point number, with a precision given by the second argument, which must be positive.

```
\xinteval{float(-2^30/3^5,12), float(-2^30/3^5, 1)}
-4.41869063374e6, -4e6
```

For this example and earlier ones if the parser had been `\xintfloateval`, not `\xinteval`, the first argument (here $2^{30}/3^5$) would already have been computed as floating point number with numerator and denominator rounded separately first to the prevailing precision. To avoid that, use `\xintexpr...\relax` wrapper. Then the rounding or truncation will be applied to an exact fraction.

sfloat(x[, n]) It is the same as `float()`, but in case of a short (non-fractional) input it gets stored internally without adding zeros to make the mantissa have the `\xinttheDigits` length. One may wonder then what is the utility of `sfloat()`? See for an example of use the documentation of `\xintdefunc`. Notice however that this is a bit experimental and may evolve in future when `xint` gets a proper internal data structure for floating point numbers. The non-normalized format is useful for multiplication or division, but float additions and subtractions usually convert their arguments to a normalized mantissa.

ilog10(x[, n]) If there is an optional argument `n`, returns the (relative) integer `a` such that $10^a \leq \text{abs}(\text{float}(x, n)) < 10^{a+1}$. In absence of the optional argument:

- in `\xintexpr`, it returns the exponent `a` such that $10^a \leq \text{abs}(x) < 10^{a+1}$.
- in `\xintfloatexpr`, the input is first rounded to `\xinttheDigits` float precision, then the exponent `a` is evaluated.

```
\xintfloateval{ilog10(99999999/10000000, 8), ilog10(-99999999/10000000, 8),
ilog10(-99999999/10000000, 9)}\newline
\xinteval{ilog10(-99999999/100000000), ilog10(-99999999/100000000, 8)}
```


TOC, Start here, [xintexpr](#), [xintexpr \(old doc\)](#), [xinttrig](#), [xintlog](#), [xinttools](#), [Examples](#), [xint bundle](#)

0, 1, 0
0, 1

If the input vanishes the function outputs `-2147450880` (i.e. `-0x7fff8000` which is near the minimal TeX number `-0x7fffffff`). This is also subject to change.

The **integer-only** variant for `\xintiexpr` admits no optional argument.

sqrt(x[, n]) in `\xintexpr...\relax` and `\xintfloatexpr...\relax` it achieves the precision given by the optional second argument. For legacy reasons the `sqrt` function in `\xintiexpr` truncates (to an integer), whereas `sqrt` in `\xintfloatexpr...\relax` (and in `\xintexpr...\relax` which borrows it) rounds (in the sense of floating numbers). There is `sqrtr` in `\xintiexpr` for rounding to nearest integer.

```
\xinttheexpr sqrt(2,31)\relax\ and \xinttheiexpr sqrt(num(2e60))\relax
1.41421356237309504880168872421 and 1414213562373095048801688724209
```

There is an **integer only** variant for `\xintiexpr`.

factorial(x[, n]) when the second optional argument is made use of inside `\xintexpr...\relax`, this switches to the use of the float version, rather than the exact one.

```
\xinttheexpr factorial (100,32)\relax, {\xintDigits:=32\relax \xintthefloatexpr
factorial (100)\relax}\newline
\xinttheexpr factorial (50)\relax\newline
\xinttheexpr factorial (50, 32)\relax
9.3326215443944152681699238856267e157, 9.3326215443944152681699238856267e157
30414093201713378043612608166064768844377641568960512000000000000
3.0414093201713378043612608166065e64
```

The **integer only variant** of course has no optional second argument.

randrange(A[, B]) when used with a single argument `A` returns a random integer $0 \leq x < A$, and when used with two arguments `A` and `B` returns a random integer $A \leq x < B$. As in Python it is an “empty range” error in first case if `A` is zero or negative and in second case if `B` $\leq A$.

Attention that the arguments are first converted to integers using `\xintNum` (i.e. truncated towards zero).

The function can be used in all three parsers. Of course the size is not limited (but in the float parser, the integer will be rounded if involved in any operation).

```
\pdfsetrandomseed 12345
\xinttheiexpr randrange(10**20)\relax\newline
\xinttheiexpr randrange(1234*10**16, 1235*10**16)\relax\newline
\printnumber{\xinttheiexpr randrange(10**199,10**200)\relax}\par
12545314555479298502
12341249468233524155
3872427149656655225094489636677708166243633082496887337312033225820004454949709978664331
9106687541716861906912743540227448009165461826072383107532471669335645234883568991327765
395258486352999399662728
```

For the support macros see `\xintRandomDigits`, `\xintiiRandRange`, `\xintiiRandRangeAtoB`. For some details regarding how `xint` uses the engine provided generator of pseudo-random numbers, see `\xintUniformDeviate`.

2.4.5. Functions with two arguments

iquo(m, n) Only available in `\xintiexpr/\xintiieval` context. Computes the Euclidean quotient. Matches with the remainder defined in next item. See `\xintiiQuo`.

irem(m, n) Only available in `\xintiexpr/\xintieval` context. Computes the Euclidean remainder. Attention that, following mathematical definition, it is always non-negative. See `\xintiRem`.

mod(f, g) computes $f - g \cdot \text{floor}(f/g)$. Hence its output is a general fraction or floating point number or integer depending on the used parser. If non-zero, it has the same sign as g .

Prior to 1.2p it computed $f - g \cdot \text{trunc}(f/g)$.

The `/:` and `'mod'` infix operators are both mapped to the same underlying macro as this `mod(f, g)` function. At 1.3 this macro produces smaller denominators when handling fractions than formerly.

```
\xinttheexpr mod(11/7,1/13), reduce(((11/7)/(1/13))*1/13+mod(11/7,1/13)),
mod(11/7,1/13)- (11/7)/(1/13), (11/7)/(1/13)\relax\newline
\xintthefloatexpr mod(11/7,1/13)\relax\par
3/91, 11/7, 0, 20
0.03296703296703260
```

Attention: the precedence rules mean that $29/5 /: 3/5$ is handled like $((29/5)/:3)/5$. This is coherent with behaviour of Python language for example:

```
>>> 29/5 % 3/5, 11/3 % 17/19, 11/57
(0.5599999999999999, 0.19298245614035087, 0.19298245614035087)
>>> (29/5) % (3/5), (11/3) % (17/19), 5/57
(0.4, 0.08771929824561386, 0.08771929824561403)
```

For comparison (observe on the last lines how `\xintfloatexpr` is more accurate than Python!):

```
\noindent\xinttheexpr 29/5 /: 3/5, 11/3 /: 17/19\relax\newline
\xinttheexpr (29/5) /: (3/5), (11/3) /: (17/19)\relax\newline
\xintthefloatexpr 29/5 /: 3/5, 11/3 /: 17/19, 11/57\relax\newline
\xintthefloatexpr (29/5) /: (3/5), (11/3) /: (17/19), 5/57\relax\newline
5/57 = \xinttheexpr trunc(5/57, 20)\relax\dots\newline
14/25, 11/57
2/5, 5/57
0.56, 0.1929824561403509, 0.1929824561403509
0.4, 0.08771929824561420, 0.08771929824561404
5/57 = 0.08771929824561403508...
```

Regarding some details of behaviour in `\xintfloatexpr`, see discussion of `divmod` function next.

divmod(f, g) computes the two mathematical values `floor(f/g)` and `mod(f,g)=f - g*floor(f/g)` and produces them as a bracketed pair in other terms it is analogous to the Python `divmod` function. Its output is equivalent to using `f//g`, `f/:g` but its implementation avoids doing twice the needed division.

In `\xintfloatexpr... \relax` the modulo is rounded to the prevailing precision. The quotient is like in the other parsers an exact integer. It will be rounded as soon as it is used in further operations, or via the global output routine of `\xintfloatexpr`. *Those examples behave as in 1.3f because assignments to multiple variables tacitly unpack if this is necessary.*

```
\xintdefvar Q, R := divmod(3.7, 1.2);%
\xinttheexpr Q, R, 1.2Q + R\relax\newline
\xintdefiivar Q, R := divmod(100, 17);%
\xinttheiexpr Q, R, 17Q + R\relax\newline
\xintdeffloatvar Q, R := divmod(100, 17e-20);%
\xintthefloatexpr Q, R, 17e-20 * Q + R\relax\newline
```

```
% show Q exactly, although defined as float it can be used in iiexpr:
\xinttheiiexpr Q\relax\ (we see it has more than 16 digits)\par
\xintunassignvar{Q}\xintunassignvar{R}%
```

```
3, 0.1, 3.7
5, 15, 100
5.882352941176471e20, 9e-20, 100
5882352941176471[5] (we see it has more than 16 digits)
```



Again: f/g or the first item output by `divmod(f, g)` is an integer q which when computed inside `\xintfloatexpr..\relax` is not yet rounded to the prevailing float precision; the second item $f-q*g$ is the rounding to float precision of the exact mathematical value evaluated with this exact q . *This behaviour may change in future major release; perhaps q will be rounded and $f-q*g$ will correspond to usage of this rounded q .*

As `\xintfloatexpr` rounds its global result, or rounds operands at each arithmetic operation, it requires special circumstances to show that the q is produced unrounded. Either as in the above example or this one with comparison operators:

```
\xintDigits := 4\relax
\xintthefloatexpr if(12345678//23==537000, 1, 0), 12345678//23\relax\newline
\xintthefloatexpr if(float(12345678//23)==537000, 1, 0)\relax\par
\xintDigits := 16\relax
```

```
0, 537000
1
```

In the first line, the comparison is done with `floor(12350000/23)=536957` (notice in passing that `12345678//23` was evaluated as `12350000//23` because the operands are first rounded to 4 digits of floating point precision), hence the conditional takes the "False" branch. In the second line the `float` forces rounding of the output to 4 digits, and the conditional takes the "True" branch.

This example shows also that comparison operators in `\xintfloatexpr..\relax` act on unrounded operands.

binomial(x, y) computes binomial coefficients. It returns zero if $y < 0$ or $x < y$ and raises an error if $x < 0$ (or if $x > 99999999$.)

```
\xinttheexpr seq(binomial(20, i), i=0..20)\relax
```

```
1, 20, 190, 1140, 4845, 15504, 38760, 77520, 125970, 167960, 184756, 167960, 125970, 77520,
38760, 15504, 4845, 1140, 190, 20, 1
```

```
\printnumber{\xintthefloatexpr seq(binomial(100, 50+i), i=-5..+5)\relax}%
```

```
6.144847121413618e28, 7.347099819081500e28, 8.441348728306404e28, 9.320655887504988e28, 9.
.891308288780803e28, 1.008913445455642e29, 9.891308288780803e28, 9.320655887504988e28, 8.
441348728306404e28, 7.347099819081500e28, 6.144847121413618e28
```

The arguments must be (expand to) short integers.

pfactorial(a, b) computes partial factorials i.e. `pfactorial(a,b)` evaluates the product $(a+1) \dots b$.

```
\xinttheexpr seq(pfactorial(20, i), i=20..30)\relax
```

```
1, 21, 462, 10626, 255024, 6375600, 165765600, 4475671200, 125318793600, 3634245014400,
109027350432000
```

The arguments must (expand to) short integers. See [subsection 11.36](#) for the behaviour if the arguments are negative.

ndfillraw(TeX-macro, n-uple) The second argument is $[N_1, N_2, \dots, N_k]$. The construct fills an $N_1 \times N_2 \times \dots \times N_k$ hyperrectangular nested list by evaluating the given `macro` as many times as needed.

Do not use!

The expansion result goes directly into internal data and must thus comply with what is expected internally for an individual numeric leaf (at 1.4, [xintfrac](#) raw format worked for `\xintexpr` or `\xintfloatexpr`, but not `\xintiexpr`, and this may have changed since). This is an experimental function serving to generate either constant or random arrays. Attention that **TEX-macro** stands here for any expandable `TeX macro`, and an `\xintexpr`-ession at this location thus requires an explicit `\xinteval` wrapping.

2.4.6. Functions with 3 or 4 arguments

if(cond,yes,no) (twofold-way conditional)

checks if `cond` is true or false and takes the corresponding branch. Any non zero number or fraction is logical true. The zero value is logical false. Both ``branches'' are evaluated (they are not really branches but just numbers). See also the `?` operator.

ifint(x,yes,no) (twofold-way conditional)

checks if `x` is an integer and in that case chooses the ``yes'' branch.
See also `isint()`.

ifone(x,yes,no) (twofold-way conditional)

checks if `x` is equal to one and in that case chooses the ``yes'' branch.
Slightly more efficient than `if(x==1,...)`. See also `isone()`.

ifsgn(cond,<0,=0,>0) (threefold-way conditional)

checks the sign of `cond` and proceeds correspondingly. All three are evaluated. See also the `??` operator.

2.4.7. Functions with an arbitrary number of arguments

The functions `all()`, `any()`, `xor()`, ``+`()`, ``*`()`, `max()`, `min()`, `gcd()`, `lcm()`, `first()`, `last()`, `reversed()` and `len()` work both with "open" and "packed" lists (aka `nutples`).

Since 1.4, when used with a single argument which is a `nutple`, it is automatically unpacked. But from 1.4 to 1.4h these functions could not be used with a single numeric argument: either they had at least two arguments, or only one and it had to be a `nutple`. At 1.4i it is again possible to use them with a lone numeric argument.

In the specific case of `reversed()` with a `nutple` argument the output is then repacked so that the output is a `nutple` if and only if the input was one (the reversal does not propagate to deeper nested `nutple`'s, it applies only at depth one).

Do not use! **qraw(stuff)** It injects directly tokens to represent internally numerical data. Will break at any release modifying the internal data format specifications (which are not always documented).

all(x, y, ...) inserts a logical **AND** in-between its arguments and evaluates the resulting logical assertion (as with all functions, all arguments are evaluated).

```
\xinteval{all(1,1,1), all([1,0,1]), all([1,1,1])}
1, 0, 1
```

any(x, y, ...) inserts a logical **OR** in-between its arguments and evaluates the resulting logical assertion,

```
\xinteval{any(0,0,0), any([1,0,1]), any([0,0,0])}
0, 1, 0
```

xor(x, y, ...) inserts a logical **XOR** in-between its arguments and evaluates the resulting logical assertion,

TOC, Start here, [xintexpr](#), *xintexpr (old doc)*, *xinttrig*, *xintlog*, *xinttools*, *Examples*, *xint bundle*

```
\xinteval{xor(1,1,1), xor([1,0,1]), xor([1,1,1])}
1, 0, 1
```

``+(x, y, ...)` adds (left ticks mandatory):

```
\xinttheexpr `+(1,3,19), `+(1**2,3**2,sqr(19)), `+([1**2,3**2,sqr(19)])\relax
23, 371, 371
```

``*(x, y, ...)` multiplies (left ticks mandatory):

```
\xinttheexpr `*(1,3,19), `*(1^2,3^2,19^2), `*([1^2,3^2,19^2])\relax
57, 3249, 3249
```

`max(x, y, ...)` maximum of the (arbitrarily many) arguments,

```
\xinttheexpr max(1,3,19), min([1,3,19])\relax
19, 1
```

`min(x, y, ...)` minimum of the (arbitrarily many) arguments,

```
\xinttheexpr min(1,3,19), min([1,3,19])\relax
1, 1
```

`gcd(x, y, ...)` computes the positive generator of the fractional ideal of rational numbers $x\mathbb{Z} + y\mathbb{Z} + \dots \subset \mathbb{Q}$. Since 1.4d the output is always in lowest terms.

This example shows how to reduce an n-uple to its primitive part:

```
\xinteval{gcd(7/300, 11/150, 13/60)}\newline
$(7/300, 11/150, 13/60)\to
(\xinteval{subsn(seq(reduce(x/D), x = L), D=gcd(L); L=7/300, 11/150, 13/60)})\newline
\xintexpr gcd([7/300, 11/150, 13/60])\relax\par
1/300
(7/300, 11/150, 13/60) → (7, 22, 65)
1/300
```

MEMO Perhaps a future release will provide a `primpart()` function as built-in functionality.

In case of strict integers, using a `\xintiexpr... \relax` wrapper is advantageous as the integer-only `gcd()` is more efficient. As `\xintiexpr` accepts only strict integers, doing this may require wrapping the argument in `num()`.

`lcm(x, y, ...)` computes the positive generator of the fractional ideal of rational numbers $x\mathbb{Z} \cap y\mathbb{Z} \cap \dots \subset \mathbb{Q}$.

```
\xinttheexpr lcm([7/300, 11/150, 13/60])\relax
```

1001/30 As for `gcd()`, since 1.4d the output is always in lowest terms. For strict integers it is slightly advantageous to use a sub `\xintiexpr`-ession.

`first(x, y, ...)` first item of the list or nutple argument:

```
\xintiexpr first([last(-7..3), [58, 97..105]])\relax
3
```

`last(x, y, ...)` last item of the list or nutple argument:

```
\xintiexpr last([-7..3, 58, first(97..105)])\relax
97
```

`reversed(x, y, ...)` reverses the order of the comma separated list or inside a nutple:

TOC, Start here, `xintexpr`, *xintexpr (old doc)*, *xinttrig*, *xintlog*, *xinttools*, *Examples*, *xint bundle*

```
\xintiieval{reversed(reversed(1..5), reversed([1..5]))}  
[5, 4, 3, 2, 1], 1, 2, 3, 4, 5
```

The above is correct as `xintexpr` functions may produce oples and this is the case here.

`len(x, y, ...)` computes the number of items in a comma separated list or inside a nutple (at first level only: it is not a counter of leaves).

```
\xinttheiexpr len(37.5), len(1..50, [101..150], 1001..1050), len([1..10])\relax  
1, 101, 10
```

`zip(*nutples)` behaves similarly to the Python function of the same name: i.e. it produces an ople of nutples, where the *i*-th nutple contains the *i*-th element from each of the argument nutples. The ople ends when the shortest input nutple is exhausted. With a single nutple argument, it returns an ople of 1-nutples. With no arguments, it returns the empty ople.

As there is no exact match in `xintexpr` of the concept of “iterator” object,¹² there is a significant difference here that (for example) the `zip(x,x,x)` Python idiom to cluster the iterator `x` into successive chunks of length 3 does not apply. Consider for this reason even the name of the function as work-in-progress, susceptible to change.

unstable?

```
\xintiieval{zip([1..9], [0, 1, 2], [11..29], [111..139])}  
[1, 0, 11, 111], [2, 1, 12, 112], [3, 2, 13, 113]
```

See also `\xintthespaceseparated` for some possible usage in combination with `flat()`.

2.4.8. Functions requiring dummy variables

The pseudo-functions `subs()`, `seq()`, `subsm()`, `subsn()`, `iter()`, `add()`, `mul()`, `rseq()`, `iterr()`, `rrseq()`, `iterr()`, `ndseq()`, `ndmap()`, `ndfillraw()` use delimited macros for some tasks:

- for all of them, whenever a `<varname>=` chunk must be parsed into a (non-assigned) variable name, then the equal sign must be visible,
- and if the syntax is with `,<varname>=` the initial comma also must be visible (spaces do not matter),
- for all of them but `ndmap()` and `ndfillraw()` the final closing parenthesis must be visible.

Although delimited macros involving commas are used to locate `,<varname>=` this is done in a way silently ignoring commas located inside correctly balanced parentheses. Thus, as the examples will show, nesting works as expected.

The semi-colons involved in the syntax may arise from expansion alone. For `rseq()`, `iter()`, `rrseq()` and `iterr()` the `,<varname>=` part may also be created from the expansion which will generate the initial comma separated values delimited by a semi-colon.

Prior to 1.4, semi-colons needed to be braced or otherwise hidden when located in an expression parsed by `\xintdefvar` or `\xintdeffunc`, to not be confused with the expression terminator.

`seq()`, `rseq()`, `iter()`, `rrseq()`, `iterr()` and also `add()`, `mul()`, but not `subs()` admit the `omit`, `abort`, and `break()` keywords. This is a new feature at 1.4 for `add()` and `mul()`.

In the case of a potentially infinite list generated by the `<integer>++` syntax, use of `abort` or of `break()` is mandatory, naturally.

All lowercase and uppercase Latin letters are pre-configured for usage as dummy variables. In Unicode engines one can use `\xintnewdummy` to turn any letter into a usable dummy variable.

And since 1.4, `\xintnewdummy` works (in all engines) to turn a multi-letter word into a dummy variable. In the descriptions, `varname` stands for such a dummy variable, either single-letter or word.

¹² Speaking of iterators, I have some ideas about this: as `\xintexpr` does not have the global expression in its hands it is difficult to organize globally expandably the idea of iterator, but locally via syntax like the one for `seq()` this is feasible. When one thinks about it, `seq()` is closely related to the iterator idea.

subs(expr, varname=values) for variable substitution.

```
\xinttheexpr subs(subs(seq(x*z,x=1..10),z=y^2),y=10)\relax\newline
100, 200, 300, 400, 500, 600, 700, 800, 900, 1000
```

Attention that `xz` generates an error, one must use explicitly `x*z`, else the parser expects a variable with name `xz`.

`subs()` is useful when defining macros for which some argument will be used more than once but may itself be a complicated expression or macro, and should be evaluated only once, for matters of efficiency. But `subs()` is helpless in function definitions: all places where a variable is substituted will receive the complete recipe to compute the variable, rather than evaluate only once.

One should rather define auxiliary functions to compute intermediate results. Or one can use `seq()`. See the documentation of [\xintdeffunc](#).

add(expr, varname=values) addition

```
\xintiexpr add(x^3,x=1..20), add(x(x+1), x=1,3,19)\relax\newline
\xintiexpr add(x^3, x = 1..[2]..20)\relax\newline          % add only odd cubes
\xintiexpr add((odd(x)){x^3}{omit}, x = 1..20)\relax\par % add only odd cubes
44100, 394
19900
19900
```

At 1.4 (fixed at 1.4a), the keywords `omit` (as in example above), `abort` and `break()` are allowed. The meaning of `break()` is specific: its argument serves as last operand for the addition, not as ultimate value.

```
\xintiexpr add((x>10){break(1000)}{x}, x = 1..15)\relax
1055
```

The `@` special variable holds the so-far accumulated value. Initially its value is zero.

```
\xintiexpr add(1 + @, i=1..10)\relax % iterates x <- 2x+1
1023
```

See ``+`()` for syntax simply adding items of a list without usage of a dummy variable.

mul(expr, varname=values) multiplication

```
\xintiexpr mul(x^2, x = 1, 3, 19, 37..50)\relax
2171846653848741108521227980217211108720640000000
```

The `@` special variable holds the so-far accumulated value. Initially its value is one.

At 1.4 (fixed at 1.4a), the keywords `omit`, `abort` and `break()` are allowed. The meaning of `break()` is specific: its argument serves as last operand for the multiplication, not as ultimate value.

```
\xintieval{mul((i==100){break(i^4)}{i}, i = 98, 99, 100)}
97020000000
```

See ``*`()` for syntax without a dummy variable.

seq(expr, varname=values) comma separated values generated according to a formula

```
\xintiexpr seq(x(x+1)(x+2)(x+3),x=1..10), ``(seq(3x+2,x=1..10))\relax
24, 120, 360, 840, 1680, 3024, 5040, 7920, 11880, 17160, 1162274713600
```

```
\smallskip
\leavevmode\vbox{\xintthealign\xintiexpr [seq([seq(i^2+j^2, i=0..j)], j=0..10)]\relax}
```

```
[[ 0 ],
 [ 1, 2 ],
 [ 4, 5, 8 ],
 [ 9, 10, 13, 18 ],
 [ 16, 17, 20, 25, 32 ],
 [ 25, 26, 29, 34, 41, 50 ],
 [ 36, 37, 40, 45, 52, 61, 72 ],
 [ 49, 50, 53, 58, 65, 74, 85, 98 ],
 [ 64, 65, 68, 73, 80, 89, 100, 113, 128 ],
 [ 81, 82, 85, 90, 97, 106, 117, 130, 145, 162 ],
 [ 100, 101, 104, 109, 116, 125, 136, 149, 164, 181, 200 ]]
```

rseq(initial value; expr, varname=values) recursive sequence, @ for the previous value.

```
\printnumber {\xintthefloatexpr subs(rseq (1; @/2+y/2@, i=1..10),y=1000)\relax }\newline
1, 500.5, 251.2490009990010, 127.6145581634591, 67.72532736082604, 41.24542607499115, 32.7
4526934448864, 31.64201586865079, 31.62278245070105, 31.62277660168434, 31.62277660168379
```

Attention: in the example above $y/2@$ is interpreted as $y/(2*@)$. With versions 1.2c or earlier it would have been interpreted as $(y/2)*@$.

In case the initial stretch is a comma separated list, @ refers at the first iteration to the whole list. Use parentheses at each iteration to maintain this ``nuple''. For example:

```
\printnumber{\xintthefloatexpr rseq(1,10^6;
(sqrt@[0]*@[1]),(@@[0]+@[1])/2), i=1..7)\relax }
1, 1e6, 1000, 500000.5, 22360.69095533499, 250500.25, 74842.22521066670, 136430.4704776675,
101048.3052657827, 105636.3478441671, 103316.8617608946, 103342.3265549749, 103329.593373
4841, 103329.5941579348, 103329.5937657094, 103329.5937657095
```

Prior to 1.4 the above example had to be written with `[@]`. This is still possible (@ stands for an ople with two items, bracketing then extracting is like extracting directly), but it is leaner to drop the extra “packing”.

iter(initial value; expr, varname=values) is exactly like `rseq`, except that it only prints the last iteration.

`iter()` is convenient to handle compactly higher order iterations. We can illustrate its use with an expandable (!) implementation of the Brent-Salamin algorithm for the computation of π :

```
\xintDigits:= 87\relax % we target 84 digits, and use 3 guard digits
\xintdeffloatfunc BS(a, b, t, p):= 0.5*(a+b), sqrt(a*b), t-p*sqrt(a-b),
\xintiexpr 2p\relax;

\xinteval
{trunc(% I feel truncation is better than rounding to display decimals of  $\pi$ 
\xintfloatexpr
iter(1, sqrt(0.5), 1, 1; % initial values
% this 43 is 84/2 + 1
(@@[0]-@[1]<2e-43)?% stopping criteria; takes into account that the
% exit computation (break() argument) doubles
% number of exact digits (roughly)
{break(sqrt(@@[0]+@[1])/@[2])} % ... do final computation,
{BS(@)}, % else do iteration
i=1++) % This generates infinite iteration. The i is not used.
\relax
% this 83 is 84 - 1 (there is a digit known to be 3 actually, before decimal mark)
, 83)% closing parenthesis of trunc()
```



```
}...% some dots following end of \xinteval argument
\xintDigits:=16\relax
```

3.14159265358979323846264338327950288419716939937510582097494459230781640628620899862...

You can try with `\xintDigits:=1004\relax` and `2e-501` in place of `\xintDigits:=87\relax` and `2e-43`, but be patient for some seconds for the result. Of course don't truncate the final result to only 83 fractional decimal digits but 1000... and better to wrap the whole thing in `\message` or `\immediate\write128` or `\edef` because it will then run in the right margin.

Prior to 1.4 the above example had to use notation such as `[@][0]`; this would still work but `@[0]` is leaner.

rrseq(initial values; expr, varname=values) recursive sequence with multiple initial terms. Say, there are K of them. Then `@1`, ..., `@4` and then `@@(n)` up to $n=K$ refer to the last K values. Notice the difference with `rseq()` for which `@` refers to a list of items in case the initial value is a list and not a single item.¹³ Using `rrseq()` with `@1` etc... accessors may be perhaps a bit more efficient than using `rseq()` with a list as starting value and constructs such as `@[0]`, `@[1]` (or rather `@[-1]`, `@[-2]` to mimick what `@1`, `@2`, `@3`, `@4` and `@@(integer)` do in `rrseq()`.

```
\xinttheiexpr rrseq(0,1; @1+@2, i=2..30)\relax
```

0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584, 4181, 6765, 10946, 17711, 28657, 46368, 75025, 121393, 196418, 317811, 514229, 832040

```
\xinttheiexpr rseq(1; 2@, i=1..10)\relax
```

1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024

```
\xinttheiexpr rseq(1; 2@+1, i=1..10)\relax
```

1, 3, 7, 15, 31, 63, 127, 255, 511, 1023, 2047

```
\xinttheiexpr rseq(2; @(@+1)/2, i=1..5)\relax
```

2, 3, 6, 21, 231, 26796

```
\xinttheiexpr rrseq(0,1,2,3,4,5; @1+@2+@3+@4+@@(5)+@@(6), i=1..20)\relax
```

0, 1, 2, 3, 4, 5, 15, 30, 59, 116, 229, 454, 903, 1791, 3552, 7045, 13974, 27719, 54984, 109065, 216339, 429126, 851207, 1688440, 3349161, 6643338

I implemented an `Rseq` which at all times keeps the memory of *all* previous items, but decided to drop it as the package was becoming big.

iterr(initial values; expr, varname=values) same as `rrseq` but does not print any value until the last K .

```
\xinttheiexpr iterr(0,1; @1+@2, i=2..5, 6..10)\relax
```

```
% the iterated over list is allowed to have disjoint defining parts.
```

55

subsm(expr, var1=value1; var2=value2; ...; varN=valueN[;]) Simultaneous substitutions. The assigned values must not involve the variables. An optional final semi-colon is allowed.

```
\xintiieval{subsm(x+2y+3z+4t, x=1; y=10; z=100; t=1000;)}\relax
```

4321

subsn(expr, var1=value1; var2=value2; ...; varN=valueN[;]) Simultaneous substitutions. The assigned values may involve all variables located further to its right. An optional final semi-colon is allowed.

```
\xintiieval{subsn(x+y+z+t, x=20y; y=20z; z=20t; t=1)}\relax
```

¹³ Prior to 1.4, one could use `@` in `rrseq()` and `iterr()` as an alias to `@1`. This undocumented feature is dropped and `@` will break `rrseq()` and `iterr()`.

8421

ndmap(function, values1; values2;; valuesN[;]) Construction of a nested list (a priori having N dimensions) from function values. The function must be an N -variable function (or a function accepting arbitrarily many arguments), but it is not constrained to produce only scalar values. Only in the latter case is the output really an N -dimensional “*ndlist*” type object. An optional final semi-colon in the input before the closing parenthesis is allowed.

```
\xintdeffunc foo(a,b,c,d) = a+b+c+d;
\begin{multicols}{2}
\xintthealign\xintexpr ndmap(foo, 1000,2000,3000; 100,200,300; 10,20,30; 1,2,3)\relax
\end{multicols}
```

```
[[[ [ 1111, 1112, 1113 ], [ 2231, 2232, 2233 ]],
 [ 1121, 1122, 1123 ], [ 2311, 2312, 2313 ]],
 [ 1131, 1132, 1133 ]], [ 2321, 2322, 2323 ]],
 [[ [ 1211, 1212, 1213 ], [ 2331, 2332, 2333 ]]],
 [ 1221, 1222, 1223 ], [[ [ 3111, 3112, 3113 ]],
 [ 1231, 1232, 1233 ]], [ 3121, 3122, 3123 ]],
 [[ [ 1311, 1312, 1313 ], [ 3131, 3132, 3133 ]]],
 [ 1321, 1322, 1323 ], [[ [ 3211, 3212, 3213 ]],
 [ 1331, 1332, 1333 ]]], [ 3221, 3222, 3223 ]],
 [[ [ 2111, 2112, 2113 ], [ 3231, 3232, 3233 ]]],
 [ 2121, 2122, 2123 ], [[ [ 3311, 3312, 3313 ]],
 [ 2131, 2132, 2133 ]]], [ 3321, 3322, 3323 ]],
 [[ [ 2211, 2212, 2213 ], [ 3331, 3332, 3333 ]]],
 [ 2221, 2222, 2223 ],
```

ndseq(expr, var1=values1; var2=values2;; varN = valuesN[;]) Constructs a nested list (a priori having N dimensions) from substitutions in an expression involving N (dummy) variables. The expression is not constrained to produce only scalar values. Only in the latter case is the output really an N -dimensional “*ndlist*” type object. An optional final semi-colon in the input before the closing parenthesis is allowed.

```
\begin{multicols}{2}
\xintthealign\xintexpr ndseq(a+b+c+d, a=1000,2000,3000; b=100,200,300;
c=10,20,30; d=1,2,3;)\relax
\end{multicols}% in case of page break, this makes amusing zigzag rendering
```

```
[[[ [ 1111, 1112, 1113 ], [ 2231, 2232, 2233 ]],
 [ 1121, 1122, 1123 ], [ 2311, 2312, 2313 ]],
 [ 1131, 1132, 1133 ]], [ 2321, 2322, 2323 ]],
 [[ [ 1211, 1212, 1213 ], [ 2331, 2332, 2333 ]]],
 [ 1221, 1222, 1223 ], [[ [ 3111, 3112, 3113 ]],
 [ 1231, 1232, 1233 ]], [ 3121, 3122, 3123 ]],
 [[ [ 1311, 1312, 1313 ], [ 3131, 3132, 3133 ]]],
 [ 1321, 1322, 1323 ], [[ [ 3211, 3212, 3213 ]],
 [ 1331, 1332, 1333 ]]], [ 3221, 3222, 3223 ]],
 [[ [ 2111, 2112, 2113 ], [ 3231, 3232, 3233 ]]],
 [ 2121, 2122, 2123 ], [[ [ 3311, 3312, 3313 ]],
 [ 2131, 2132, 2133 ]]], [ 3321, 3322, 3323 ]],
 [[ [ 2211, 2212, 2213 ], [ 3331, 3332, 3333 ]]],
 [ 2221, 2222, 2223 ],
```

Recursions may be nested, with `@@@(n)` giving access to the values of the outer recursion... and there is even `@@@@(n)` to access the outer outer recursion but I never tried it!

The following keywords are recognized:

abort it is a pseudo-variable which indicates to stop here and now.

omit it is a pseudo-variable which says to omit this value and go to next one.

break(stuff) says to abort and insert `stuff` as last value.

`<integer>++` serves to generate a potentially infinite list. In conjunction with an `abort` or `break()` this is often more efficient than iterating over a pre-established list of values.

```
\xinttheiexpr iter(1;(@>10^40)?{break(@)}{2@},i=1++)\relax
10889035741470030830827987437816582766592 is the smallest power of 2 with at least forty one
digits.
```

The `i=<integer>++` syntax (any letter is allowed in place of `i`) works only in the form `<letter>_i=<integer>++`, something like `x=10,17,30++` is not legal. The `<integer>` must be a TeX-allowable integer.

```
First Fibonacci number at least |2^31| and its index
% we use iterr to refer via @1 and @2 to the previous and previous to previous.
\xinttheiexpr iterr(0,1; (@1>=2^31)?{break(@1, i)}{@2+@1}, i=1++)\relax
First Fibonacci number at least 2^31 and its index 2971215073, 47. If one also wants the pre-
vious Fibonacci number one only has to use break(@2, @1, i) in the above example.
```

2.5. Generators of arithmetic progressions

- `a..b` constructs the `small` integers from the ceil `[a]` to the floor `[b]` (possibly a decreasing sequence): one has to be careful if using this for algorithms that `1..0` for example is not empty or `1` but expands to `1, 0`. Again, `a..b` can not be used with `a` and `b` greater than $2^{31} - 1$. Also, only about at most `5000` integers can be generated (this depends upon some TeX memory settings).

The `..` has lower precedence than the arithmetic operations.

```
\xintexpr 1.5+0.4..2.3+1.1\relax; \xintexpr 1.9..3.4\relax; \xintexpr 2..3\relax
2, 3; 2, 3; 2, 3
```

The step of replacing `a` by its ceil and `b` by its floor is a kind of silly overhead, but `a` and `b` are allowed to be themselves the result of computations and there is no notion of “int” type in `\xinteval`. The solution is, when `a` and `b` are given explicit integers to temporarily switch to the `\xintiexpr` parser:

```
\xintexpr \xintiexpr 1..10\relax\relax
1, 2, 3, 4, 5, 6, 7, 8, 9, 10
```

On the other hand integers from `\xintexpr 1..10\relax` are already in raw `xintfrac` format for example `3/1[0]` which speeds up their usage in the macros internally involved in computations... thus perhaps what one gains on one side is lost on the other side.

- `a..[d]..b` generates “real” numbers along arithmetic progression of reason `d`. It does not replace `a` by its ceil, nor `b` by its floor. The generated list is empty if `b-a` and `d` are of opposite signs; if `d=0` or if `a=b` the list expands to single element `a`.

```
\xintexpr 1.5..[1.01]..11.23\relax
1.5, 2.51, 3.52, 4.53, 5.54, 6.55, 7.56, 8.57, 9.58, 10.59
```

At 1.4, this generator behaves in `\xintfloatexpr` exactly as in `\xintexpr`, i.e. *exactly*. This is breaking change.

```
\xintDigits:=6;
\xintexpr\xintfloatexpr 100..[1.23456]..110\relax\relax
\xintDigits:=16;
```

100, 101.23456, 102.46912, 103.70368, 104.93824, 106.1728, 107.40736, 108.64192, 109.87648

This demonstration embedded the float expression in the exact parser only to avoid the rounding to the prevailing precision on output, thus we can see that internally additions are done exactly and not with 6 digits mantissas (in this example).

2.6. Python slicing and indexing of one-dimensional sequences

We denote here by *list* or *sequence* a general *opple*, either given as a variable or explicitly. In the former case the parentheses are optional.¹⁴

- `(list)[n]` returns the $n+1$ th item if $n \geq 0$. If $n < 0$ it enumerates items from the tail. Items are numbered as in Python, the first element corresponding to $n=0$.

```
\xintexpr (0..10)[6], (0..10)[-1], (0..10)[23*18-22*19]\relax
6, 10, 7
```

This also works for singleton *opples* which are in fact a *number*:

```
\xintexpr (7)[0], (7)[-1], 9, (7)[-2], 9\relax
7, 7, 9, 9
```

In the example above the parentheses serve to disambiguate from the raw `xintfrac` format such as `7[-1]` which, although discouraged, is accepted on input. And we used a trick to show that `(7)[-2]` returns `nil`.

The behaviour changes for singleton *opples* which are not *numbers*. They are thus *nutples*, or equivalently they are the bracketing (bracing, packing) of another *opple*. In this case, the meaning of the syntax for item indexing is, as in Python, item extraction:

```
\xintexpr [0,1,2,3,4,5][2], [0,1,2,3,4,5][-3]\relax\newline
\xintexpr [0,[1,2,3,4,5],6][1][-1]\relax
2, 3
5
```

- `(list)[:n]` produces the first n elements if $n > 0$, or suppresses the last $|n|$ elements if $n < 0$.

```
\xintiexpr (0..10)[:6]\relax\ and \xintiexpr (0..10)[: -6]\relax
0, 1, 2, 3, 4, 5 and 0, 1, 2, 3, 4
```

As above, the meaning change for *nutples* and fits with expectations from Python regarding its sequence types:

```
\xintiexpr [0..10][:6]\relax\ and \xintiexpr [0..10][: -6]\relax
[0, 1, 2, 3, 4, 5] and [0, 1, 2, 3, 4]
```

- `(list)[n:]` suppresses the first n elements if $n > 0$, or extracts the last $|n|$ elements if $n < 0$.

```
\xintiexpr (0..10)[6:]\relax\ and \xintiexpr (0..10)[-6:]\relax
6, 7, 8, 9, 10 and 5, 6, 7, 8, 9, 10
```

As above, the meaning change for *nutples* and fit with expectations from Python with *tuple* or *list* types:

```
\xintiexpr [0..10][6:]\relax\ and \xintiexpr [0..10][-6:]\relax
[6, 7, 8, 9, 10] and [5, 6, 7, 8, 9, 10]
```

¹⁴ Even for an “open list”, if it is given as a *variable* then the indexing or slicing will not apply to its last item but to itself as an entity.

TOC, Start here, [xintexpr](#), *xintexpr (old doc)*, *xinttrig*, *xintlog*, *xinttools*, *Examples*, *xint bundle*

- Finally, `(list)[a:b]` also works according to the Python ``slicing'' rules (inclusive of negative indices). Notice though that stepping is currently not supported.

```
\xinttheiexpr (1..20)[6:13]\relax\ = \xinttheiexpr (1..20)[6-20:13-20]\relax
\newline
\xinttheiexpr [1..20][6:13]\relax\ = \xinttheiexpr [1..20][6-20:13-20]\relax
7, 8, 9, 10, 11, 12, 13 = 7, 8, 9, 10, 11, 12, 13
[7, 8, 9, 10, 11, 12, 13] = [7, 8, 9, 10, 11, 12, 13]
```

- It is naturally possible to execute such slicing operations one after the other (the syntax is simplified compared to before 1.4):

```
\xintexpr (1..50)[13:37][10:-10]\relax\newline
\xintexpr (1..50)[13:37][10:-10][-1]\relax
24, 25, 26, 27
27
```

2.7. NumPy like nested slicing and indexing for arbitrary oples and nuples

I will give one illustrative example and refer to the NumPy documentation for more.

Notice though that our interpretation of the syntax is more general than NumPy's concepts (of basic slicing/indexing):

- slicing and itemizing apply also to non-bracketed objects i.e. *oples*,
- the leaves do not have to be all at the same depth,
- there are never any out-of-range index errors: out-of-range indices are silently ignored.

```
\begin{multicols}{3}
\xintdefvar myArray = ndseq(a+b+c, a=100,200,300; b=40,50,60; c=7,8,9);
myArray = \xintthealign\xintexpr myArray\relax
\columnbreak
mySubArray = \xintthealign\xintexpr myArray[0:2,0:2,0:2]\relax
myExtractedSubArray = \xintthealign\xintexpr myArray[0:2,0:2,0:2][0]\relax
\columnbreak
myExtractedSubArray = \xintthealign\xintexpr myArray[0:2,0:2,0:2][0,1]\relax
\noindent
firstExtractedScalar = \xintexpr myArray[0:2,0:2,0:2][0,1,0]\relax\newline
secondExtractedScalar = \xintexpr myArray[0,1,0]\relax\par
\end{multicols}
```

```
myArray =          mySubArray =          myExtractedSubArray =
[[[ 147, 148, 149 ],      [[ [ 147, 148 ],      [ 157, 158 ]
 [ 157, 158, 159 ],      [ 157, 158 ]],
 [ 167, 168, 169 ]],     [[ 247, 248 ],
 [ 247, 248, 249 ],      [ 257, 258 ]]]
 [ 257, 258, 259 ],
 [ 267, 268, 269 ]],     myExtractedSubArray =
[[ 347, 348, 349 ],      [[ 147, 148 ],
 [ 357, 358, 359 ],      [ 157, 158 ]]
 [ 367, 368, 369 ]]]
```

As said before, *stepping* is not yet implemented. Also the NumPy extension to Python for item selection (i.e. via a **tuple** of comma separated indices) is not yet implemented.

2.8. Tacit multiplication

Tacit multiplication (insertion of a `*`) applies when the parser is currently either scanning the digits of a number (or its decimal part or scientific part, or hexadecimal input), or is looking for an infix operator, and:

- (1.) encounters a count or *dimen* or *skip register* or variable or an ϵ -TeX expression, or
- (2.) encounters a sub-`\xintexpression`, or
- (3.) encounters an opening parenthesis, or
- (4.) encounters a letter (which is interpreted as signaling the start of either a variable or a function name), or
- (5.) (of course, only when in state "looking for an operator") encounters a digit.

!!!!ATTENTION!!!!

Explicit digits prefixing a variable, or a function, whose name starts with an `e` or `E` will trap the parser into trying to build a number in scientific notation. So the `*` must be explicitly inserted.

```
\xintdefiivar e := (2a+4b+6d+N)/:7;%
\xintdefiivar f := (c+11d+22*e)//451;% 22e would raise errors
I don't think I will fix this anytime soon...
```

For example, if `x`, `y`, `z` are variables all three of `(x+y)z`, `x(y+z)`, `(x+y)(x+z)` will create a tacit multiplication.

Furthermore starting with release 1.2e, whenever tacit multiplication is applied, in all cases it always ``ties'' more than normal multiplication or division, but still less than power. Thus `x/2y` is interpreted as `x/(2y)` and similarly for `x/2max(3,5)` but `x^2y` is still interpreted as `(x^2)*y` and `2n!` as `2*n!`.

```
\xintdefvar x:=30;\xintdefvar y:=5;%
\xinttheexpr (x+y)x, x/2y, x^2y, x!, 2x!, x/2max(x,y)\relax
1050, 30/10, 4500, 265252859812191058636308480000000, 530505719624382117272616960000000,
30/60
```

Since 1.2q tacit multiplication is triggered also in cases such as `(1+2)5` or `10!20!30!`.

```
\xinttheexpr (10+7)5, 4!4!, add(i, i=1..10)10, max(x, y)100\relax
85, 576, 550, 3000
```

The ``tie more'' rule applies to all cases of tacit multiplication. It impacts only situations with a division operator as the last seen operator, as multiplication is mathematically associative.

```
\xinttheexpr 1/(3)5, (1+2)/(3+4)(5+6), 2/x(10), 2/10x,
3/y\xintiexpr 5+6\relax, 1/x(y)\relax\
differ from\newline\xinttheexpr 1/3*5, (1+2)/(3+4)*(5+6), 2/x*(10), 2/10*x,
3/y*\xintiexpr 5+6\relax, 1/x*(y)\relax\par
1/15, 3/77, 2/300, 2/300, 3/55, 1/150 differ from
5/3, 33/7, 20/30, 60/10, 33/5, 5/30
```

Note that `y\xinttheiexpr 5+6\relax` would have tried to use a variable with name `y11` rather than doing `y*11`: tacit multiplication works only in front of sub-`\xintexpressions`, not in front of `\xinttheexpressions` which are unlocked into explicit digits.

Here is an expression whose meaning is completely modified by the ``tie more'' property of tacit multiplication:

```
\xintdeffunc e(z):=1+z(1+z/2(1+z/3(1+z/4)));
will be parsed as
\xintdeffunc e(z):=1+z*(1+z/(2*(1+z/(3*(1+z/4)))));
```

which is not at all the presumably hoped for:

```
\xintdeffunc e(z):=1+z*(1+(z/2)*(1+(z/3)*(1+(z/4))));
```

2.9. User defined variables

Since release 1.1 it is possible to make an assignment to a variable name and let it be known to the parsers of *xintexpr*. Since 1.2p simultaneous assignments are possible. Since 1.4 simultaneous assignments are possible with a right-hand-side being a *nutple* which will be automatically unpacked.

```
\xintdefvar myPi:=3.141592653589793238462643;%
$myPi = \xinteval{myPi}$\newline % (there is already built-in Pi variable)
\xintdefvar x_1, x_2, x_3 := 10, 20, 30;%
$x_1 = \xinteval{x_1}, x_2 = \xinteval{x_2}, x_3 = \xinteval{x_3}$\newline
\xintdefvar x_1, x_2, x_3 := [100, 200, 300];%
$x_1 = \xinteval{x_1}, x_2 = \xinteval{x_2}, x_3 = \xinteval{x_3}$\par
```

myPi = 3.141592653589793238462643

x₁ = 10, x₂ = 20, x₃ = 30

x₁ = 100, x₂ = 200, x₃ = 300

Simultaneous assignments with more variables than values do not raise an error but simply set the extra variables to the *nil* value.

```
\xintdefiivar a, b, c := [1, 2];% will be automatically unpacked
The value of a is \xinteval{a}, the one of b is \xinteval{b} and
the one of c is \xinteval{c}.
```

The value of a is 1, the one of b is 2 and the one of c is .

```
\xintdefiivar a, b, c := 314;%
The value of a is \xinteval{a}, the one of b is \xinteval{b} and
the one of c is \xinteval{c}.
```

The value of a is 314, the one of b is and the one of c is .

Notice that *nil* variables must be used with caution as they break arithmetic operations if used as operands to them. And they are not the same as the *None* variables, which can also be input as `[]`.

Simultaneous assignments with less variables than values do not raise an error but set the last variable to be the ople concatenating the remaining values.

```
\xintdefiivar seq := 1..10;%
\xintdefiivar a, seq := seq;%
\xintdefiivar b, seq := seq;%
\xintdefiivar c, d, seq := seq;%
The value of a is \xinteval{a}, the one of b is \xinteval{b}, the one of c is \xinteval{c},
the one of d is \xinteval{d}, the one of seq is \xinteval{seq}.
```

The value of a is 1, the one of b is 2, the one of c is 3, the one of d is 4, the one of seq is 5, 6, 7, 8, 9, 10.

In the above we define a variable *seq* but there is a built-in function *seq()*. It is indeed allowed to use the same name for both a variable and a function.¹⁵ But for safety we will unassign *seq* now:

```
\xintunassignvar{a}\xintunassignvar{b}\xintunassignvar{c}\xintunassignvar{d}%
\xintunassignvar{seq}%
```

Single letter names *a..z* and *A..Z* are pre-declared by the package for use as a special type of variables called ``dummy variables'`. Unassigning them restores this initial meaning. See further `\xintunassignvar` and `\xintnewdummy`. Since 1.4 even assigned variables can be used in the call signatures of function declarations.

Regarding the manipulation of an “open list” as above, there is no way to obtain with only one use of the variable both its last item and the reduction of the variable to its truncated self. One can do rather:

¹⁵ But until a bugfix added at release 1.4i, some built-in function names (those implementing syntax with dummy variables, and the so-called “pseudo”-functions) were fragile under such overloading.

```
\xintdefiivar mylist := 1..10;%
\xintdefiivar z, mylist := last(mylist), mylist[:-1];%
The value of z is \xinteval{z} and mylist is now \xinteval{mylist}.\par
```

The value of `z` is 10 and `mylist` is now 1, 2, 3, 4, 5, 6, 7, 8, 9.

This uses twice `mylist` and is about the same as doing it in two steps:

```
\xintdefiivar w := last(mylist);%
\xintdefiivar mylist := mylist[:-1];%
The value of w is \xinteval{w} and mylist is now \xinteval{mylist}.%
\xintunassignvar{z}\xintunassignvar{w}\xintunassignvar{mylist}\par
```

The value of `w` is 9 and `mylist` is now 1, 2, 3, 4, 5, 6, 7, 8.

It is recommended generally speaking to work with “closed (i.e. bracketed) lists” because only them and numbers can be arguments to functions (but see [\xintdeffunc](#) and the notion of variadic last argument). For more on the Python-like slicing used above see [subsection 2.6](#) and [subsection 2.13.4](#). For more information relative to variables versus arguments see [subsection 2.13.6](#).

- For catcodes issues (particularly, for the semi-colon used to delimit the fetched expression), see the discussion of [\xintexprSafeCatcodes](#).
- Both syntaxes `\xintdefvar foo := <expr>;` and `\xintdefvar foo = <expr>;` are accepted.
- Spaces in the variable name or around the equal sign are removed and are immaterial.
- The variable names are expanded in an `\edef` (and stripped of spaces). Example:

```
\xintdefvar x\xintListWithSep{, x}{\xintSeq{0}{10}} := seq(2**i, i = 0..10);%
This defines x0, x1, . . . , x10 for future usage.
```

Legal variable names are composed of letters, digits, `_` and `@` and characters. A variable name must start with a letter. Variable names starting with a `@` or `_` are reserved for internal usage.¹⁶

As `x_1x_2` or even `x_1x` are licit variable names, and as the parser does not trace back its steps, input syntax must be `x_1*x_2` if the aim is to multiply such variables.

Using `\xintdefvar`, `\xintdefiivar`, or `\xintdeffloatvar` means that the variable value will be computed using respectively `\xintexpr`, `\xintiexpr` or `\xintfloatexpr`. It can then be used in all three parsers, as long as the parser understands the format. Currently this means that variables using `\xintdefvar` or `\xintdeffloatvar` can be used freely either with `\xintexpr` or `\xintfloatexpr` but not with `\xintiexpr`, and variables defined via `\xintdefiivar` can be used in all parsers.

When defining a variable with `\xintdeffloatvar` it (or generally speaking its numerical leaves) is rounded to `\xinttheDigits` precision. So the variable holds the same value as would be printed via `\xintfloateval` for the same computation.

Prior to 1.4e, this was the case only if the variable definition actually involved some computation.

However the `\xintfloatexpr.\relax` wrapper by itself induces no rounding. If it is encountered in the typesetting flow, the print-out will be rounded to `\xinttheDigits` precision, but this is an effet of behaving like `\xintfloateval` in this context.

```
% Since 1.4e, \xintdeffloatvar always rounds (to \xinttheDigits)
\xintdeffloatvar e:=2.7182818284590452353602874713526624977572470936999595749669676;%
1) \xintexpr e\relax\newline % shows the recorded value: it is rounded
\xintunassignvar{e}
2) \xintfloatexpr % when used in typesetting flow, acts like \xintfloateval:
   2.7182818284590452353602874713526624977572470936999595749669676
\relax\newline % the print-out is rounded.
3) \xintexpr
   \xintfloatexpr
```

¹⁶ The process of variable declaration does not check that these rules are met, and breakage will arise on use, if rules are not followed. For example, prior to 1.4g, using a variable (illegally) declared with a name starting with a (normal, catcode 8) `_` triggered an infinite loop.


```

2.7182818284590452353602874713526624977572470936999595749669676
\relax
\relax\newline
%
% but we can see via the \xintexpr wrapper all the digits were there rounding
% can be forced using an extra 0+, the float() function, or the [D] option.
% tidbit: comparison operators do not pre-round, so 1.2345678 is not same as
% (1.2345678+0) in low precision.
%
\begingroup\xintDigits:=4;%
4) \xintifboolfloatexpr{1.2345 == 1.23456}
   {\error}{Different! Comparisons do not pre-round to Digits precision.}\newline
5) \xintifboolfloatexpr{1.2345 == 1.2345 + 0}
   {\error}{Different! Right hand side rounded from operation,
           left hand side not rounded.}\par
\endgroup

```

- 1) 2.718281828459045
- 2) 2.718281828459045
- 3) 2.7182818284590452353602874713526624977572470936999595749669676
- 4) Different! Comparisons do not pre-round to Digits precision.
- 5) Different! Right hand side rounded from operation, left hand side not rounded.

After issuing `\xintverbosetrue` the values of defined variables are written out to the log (and terminal). As in this example:

```

Package xintexpr Info: (on line 1)
  Variable myPi defined with value {3141592653589793238462643[-24]}.
Package xintexpr Info: (on line 2)
  Variable x_1 defined with value {10}.
Package xintexpr Info: (on line 2)
  Variable x_2 defined with value {20}.
Package xintexpr Info: (on line 2)
  Variable x_3 defined with value {30}.
Package xintexpr Info: (on line 3)
  Variable List defined with value {0}{1}{3}{6}{10}{15}{21}{28}{36}{45}{55}
.
Package xintexpr Info: (on line 4)
  Variable Nuple defined with value {{0}{1}{9}{36}{100}{225}{441}{784}{1296}
}{2025}{3025}}.
Package xintexpr Info: (on line 5)
  Variable FourthPowers defined with value {{0}{1}{81}{1296}{10000}{50625}{
194481}{614656}{1679616}{4100625}{9150625}}.

```

2.9.1. \xintunassignvar

Variable declarations obey the current scope. To let a (multi-letter) name be unknown to (all parsers of) `xintexpr` without waiting the end of the scope one issues `\xintunassignvar{<variable>}`.



In the special case of `\xintunassignvar{<letter>}`, the effect is different, as it is synonymous with `\xintnewdummy{<letter>}`: the (catcode 11) `<letter>` recovers or acquires meaning as a dummy variable in the current scope.

```

\xintunassignvar{e}%
% overwriting a dummy letter
\xintdefvar i := 3;%
\xinttheiexpr add(i, i = 1..10)\relax\ ("i" has the fixed value 3)\newline
\xintunassignvar{i}% back to normal

```

```
\xinttheiexpr add(i, i = 1..10)\relax ("i" is again a dummy variable)\par
30 ("i" has the fixed value 3)
55 ("i" is again a dummy variable)
```

Under `\xintglobaldefstrue` regime the effect of `\xintunassignvar` is global.

2.9.2. `\xintnewdummy`

Any catcode 11 character can serve as a dummy variable, via this declaration:

```
\xintnewdummy{<character>}
```

For example with XeTeX or LuaTeX the following works:

```
% use a Unicode engine
\input xintexpr.sty
\xintnewdummy ξ% or any other letter character !
\xinttheexpr add(ξ, ξ=1..10)\relax
\bye
```

Under `\xintglobaldefstrue` regime the effect of `\xintnewdummy` is global.

Starting with 1.4, it is allowed to use `\xintnewdummy` with multi-letter names (obeying the condition for being a variable name).

2.9.3. `\xintensuredummy`, `\xintrestorevariable`

Use

```
\xintensuredummy{<character>}
...
... code using the (catcode 11) character as a dummy variable
...
\xintrestorevariable{<character>}
```

if other parts need the letter as an assigned variable name. For example `xinttrig` being written at high level needs a few genuine dummy variables, and it uses `\xintensuredummy` to be certain everything is ok.

2.10. User defined functions

2.10.1 <code>\xintdeffunc</code>	42
2.10.2 <code>\xintdefiifunc</code>	45
2.10.3 <code>\xintdeffloatfunc</code>	45
2.10.4 <code>\xintdefufunc</code> , <code>\xintdefiifunc</code> , <code>\xintdeffloatufunc</code>	45
2.10.5 Using the same name for both a variable and a function	47
2.10.6 <code>\xintunassignexprfunc</code> , <code>\xintunassigniifunc</code> , <code>\xintunassignfloatexprfunc</code>	47
2.10.7 <code>\ifxintverbose</code> conditional	49
2.10.8 <code>\ifxintglobaldefs</code> conditional	49
2.10.9 <code>\xintNewFunction</code>	49

2.10.1. `\xintdeffunc`

Here is an example:

```
\xintdeffunc
  Rump(x,y):=1335 y^6/4 + x^2 (11 x^2 y^2 - y^6 - 121 y^4 - 2) + 11 y^8/2 + x/2y;
```

(notice the numerous tacit multiplications in this expression; and that `x/2y` is interpreted as `x/(2y)`.)

- The ending semi-colon is allowed to be of active catcode, as `\xintdeffunc` temporarily resets catcodes before parsing the expression. But this will fail if the whole thing is inside a macro definition. Then the used semi-colon must be the standard one.
- Semi-colons used inside the expression need not be hidden inside braces. (new with 1.4)
- The colon before the equal sign is optional and its (reasonable) catcode does not matter.

Here are a few important items (bookmark this for reading again later once you have gained experience in using this interface...):

- The function names are composed of letters, digits, underscores or @ signs. A function name must start with a letter. It may be a single letter (see [subsubsection 2.10.5](#)).
- The variable names used in the function signature may be multi-letter words. It is also allowed for them to already be in use for previously declared variables. Their meanings will get restored for usage after the function declaration.
- A function can be declared with at most nine arguments. It can be declared as a function with no arguments.
- If in the function declaration the last argument is prefixed by `*`, it stands for a `nutple` which will gather all arguments of the function call beyond the first positional ones. See [subsubsection 2.13.6](#) for additional explanations on such “variadic” arguments.
- Recursive definitions are possible; for them to not generate error or fall in infinite loops, the use of the short-circuit conditionals `?` and `??` is mandatory.
- If a function is used in another definition it will check if it is applied to numerical arguments and if this is the case will expand fully.
- The previous item has an exception for functions with no arguments; they never expand immediately in other function definitions (else they would be almost like variables). This provides a way to define functions with parameters: simply let their definition use some functions with no arguments.
- A function declared via `\xintdeffunc` remains `unknown` to `\xintfloatexpr` (or `\xintfloateval`). See `\xintdeffloatfunc`, `\xintdefiifunc`. One can use the same formula in a new definition, but if one wants the expansion to execute in a parser independent way, one can transfer a function with scalar values like this:

```
\xintdeffloatfunc foo(x) := float_dgt(\xintexpr foo(x)\relax);
```

The `float_dgt()` wrapper (which was renamed at 1.4e) is in order for the float variant to produce an already-rounded value, possibly speeding-up usage if used as input for other functions. Using `float()` here would work the same but the produced function would incorporate a routine to check (at time of use, because at time of definition it is impossible to tell what will be the case) if it is applied to one or two arguments.

- And in the reverse direction one can do:

```
\xintdeffunc bar(x) := \xintfloatexpr bar(float(x))\relax;
```

With this the transplanted float-function will expand in `\xintexpr` as it would have in `\xintfloatexpr`, i.e. using float operations; this is different from declaring the function again with the same expression as used for the original, as it would have then been parsed with a

Conclusion: if some `\xintdeffunc` break, check if it does not fit the above criterion before reporting... and recall `\xintNewFunction` is your friend. It has the big advantage of declaring a function for all parsers simultaneously!

A special note on `subs()`: it is and has always been hopeless in `\xintdeffunc` context. All it does (if it works at all) after being malaxed by `\xintdeffunc` is to copy over at the indicated places the *recipe* to compute something. Thus at every location where that something is needed it will be evaluated from scratch again. Yes, this is disappointing. But... on the other hand the more general `seq()` does work, or pretends to work. Let me illustrate to make things clear. We start with this:

```
\xintverbosetrue
\xintdeffunc foo(x,y,z) = subs(S + S^2, S = x+y+z);
\xintdeffunc bar(x,y,z) = seq(S + S^2, S = x+y+z);
\xintexpr foo(100,10,1), bar(100,10,1)\relax
\xintverbosefalse
```

12432, 12432

It produces in the log:

```
Package xintexpr Info: (on line 2)
  Function foo for \xintexpr parser associated to \XINT_expr_userfunc_foo with
  h meaning macro:#1#2#3->{\xintAdd {\xintAdd {\xintAdd {#1}{#2}}{#3}}{\xintPow {
\xintAdd {\xintAdd {#1}{#2}}{#3}}{2}}}
Package xintexpr Info: (on line 3)
  Function bar for \xintexpr parser associated to \XINT_expr_userfunc_bar with
  h meaning macro:#1#2#3->\expanded \bgroup \expanded {\unexpanded {\XINT_expr_se
q:_b {\xintbareeval S + S^2\relax !S}}{\xintAdd {\xintAdd {#1}{#2}}{#3}}^}
```

Even without understanding all details one sees that in the first case the `\xintAdd {\xintAdd {#1}{#2}}{#3}` appears twice, and in the second case only once. But in the second case we have a yet to evaluate expression. So the second approach is not much different in its effect than using the more simple-minded `\xintNewFunction`. Besides one gets a feeling why the function arguments can not appear in the expression but only in the iterated over values, because there is no way to understand what `x`, `y`, `z` are supposed to mean without adding extra structure showing they map to `#1`, `#2`, `#3`.

The above remarks apply to `subsm()` and `subsn()`. Even if they do work in `\xintdeffunc` context (warning, testing at 1.4 release has remained minimal), they will not bring added efficiency if the substituted values are to be used multiple times. They may still be useful to visually simplify the input of a big expression by expressing it in terms of smaller constituents.

Another workaround if one wants genuine (not “macro”-) functions for some expression where the same thing is used multiple times is to define helper functions computing the intermediate data. One can see illustrations of this in the code source of `xinttrig` (or in the matrix multiplication example at the end of this chapter).

2.10.2. `\xintdefiifunc`

With `\xintdeffunc` the created function is known by the `\xintexpr` parser only. For usage in the `\xintiexpr` parser, it is required to use `\xintdefiifunc`.

2.10.3. `\xintdeffloatfunc`

With `\xintdeffunc` the created function is known by the `\xintexpr` parser only. For usage in the `\xintfloatexpr` parser, it is required to use `\xintdeffloatfunc`.

Note: the optional argument `[Q]` accepted by `\xintfloatexpr` does not work with `\xintdeffloatfunc`. It is still possible to wrap the expression in `float(expression,Q)`, if it evaluates to a scalar.

2.10.4. `\xintdefufunc`, `\xintdefiifunc`, `\xintdeffloatufunc`

This allows to define so-called “Universal functions”. This is terminology borrowed from NumPy.

Here is an example:

```
\xintdefiiivar Array = ndmap(lcm, 1..5; 1..10; 1..10);
Array = \xintthealign\xintiexpr Array\relax
\xintdefiiufunc foo(x) = x^3;
\begin{figure}[htbp]
\caption{Output of a universal function acting on an array}\label{fig:ufunc}
\centeredline{\$\vcenter{\xintthealign\xintiexpr foo(Array)\relax}\$}
\end{figure}
See \autopageref{fig:ufunc} for the output.
```

Array =

```
[[ [ 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 ],
  [ 2, 2, 6, 4, 10, 6, 14, 8, 18, 10 ],
  [ 3, 6, 3, 12, 15, 6, 21, 24, 9, 30 ],
  [ 4, 4, 12, 4, 20, 12, 28, 8, 36, 20 ],
  [ 5, 10, 15, 20, 5, 30, 35, 40, 45, 10 ],
  [ 6, 6, 6, 12, 30, 6, 42, 24, 18, 30 ],
  [ 7, 14, 21, 28, 35, 42, 7, 56, 63, 70 ],
  [ 8, 8, 24, 8, 40, 24, 56, 8, 72, 40 ],
  [ 9, 18, 9, 36, 45, 18, 63, 72, 9, 90 ],
  [ 10, 10, 30, 20, 10, 30, 70, 40, 90, 10 ]],
[[ [ 2, 2, 6, 4, 10, 6, 14, 8, 18, 10 ],
  [ 2, 2, 6, 4, 10, 6, 14, 8, 18, 10 ],
  [ 6, 6, 6, 12, 30, 6, 42, 24, 18, 30 ],
  [ 4, 4, 12, 4, 20, 12, 28, 8, 36, 20 ],
  [ 10, 10, 30, 20, 10, 30, 70, 40, 90, 10 ],
  [ 6, 6, 6, 12, 30, 6, 42, 24, 18, 30 ],
  [ 14, 14, 42, 28, 70, 42, 14, 56, 126, 70 ],
  [ 8, 8, 24, 8, 40, 24, 56, 8, 72, 40 ],
  [ 18, 18, 18, 36, 90, 18, 126, 72, 18, 90 ],
  [ 10, 10, 30, 20, 10, 30, 70, 40, 90, 10 ]],
[[ [ 3, 6, 3, 12, 15, 6, 21, 24, 9, 30 ],
  [ 6, 6, 6, 12, 30, 6, 42, 24, 18, 30 ],
  [ 3, 6, 3, 12, 15, 6, 21, 24, 9, 30 ],
  [ 12, 12, 12, 12, 60, 12, 84, 24, 36, 60 ],
  [ 15, 30, 15, 60, 15, 30, 105, 120, 45, 30 ],
  [ 6, 6, 6, 12, 30, 6, 42, 24, 18, 30 ],
  [ 21, 42, 21, 84, 105, 42, 21, 168, 63, 210 ],
  [ 24, 24, 24, 24, 120, 24, 168, 24, 72, 120 ],
  [ 9, 18, 9, 36, 45, 18, 63, 72, 9, 90 ],
  [ 30, 30, 30, 60, 30, 30, 210, 120, 90, 30 ]],
[[ [ 4, 4, 12, 4, 20, 12, 28, 8, 36, 20 ],
  [ 4, 4, 12, 4, 20, 12, 28, 8, 36, 20 ],
  [ 12, 12, 12, 12, 60, 12, 84, 24, 36, 60 ],
  [ 4, 4, 12, 4, 20, 12, 28, 8, 36, 20 ],
  [ 20, 20, 60, 20, 20, 60, 140, 40, 180, 20 ],
  [ 12, 12, 12, 12, 60, 12, 84, 24, 36, 60 ],
  [ 28, 28, 84, 28, 140, 84, 28, 56, 252, 140 ],
  [ 8, 8, 24, 8, 40, 24, 56, 8, 72, 40 ],
  [ 36, 36, 36, 36, 180, 36, 252, 72, 36, 180 ],
  [ 20, 20, 60, 20, 20, 60, 140, 40, 180, 20 ]],
[[ [ 5, 10, 15, 20, 5, 30, 35, 40, 45, 10 ],
  [ 10, 10, 30, 20, 10, 30, 70, 40, 90, 10 ]],
```

```
[ 15, 30, 15, 60, 15, 30, 105, 120, 45, 30 ],
[ 20, 20, 60, 20, 20, 60, 140, 40, 180, 20 ],
[ 5, 10, 15, 20, 5, 30, 35, 40, 45, 10 ],
[ 30, 30, 30, 60, 30, 30, 210, 120, 90, 30 ],
[ 35, 70, 105, 140, 35, 210, 35, 280, 315, 70 ],
[ 40, 40, 120, 40, 40, 120, 280, 40, 360, 40 ],
[ 45, 90, 45, 180, 45, 90, 315, 360, 45, 90 ],
[ 10, 10, 30, 20, 10, 30, 70, 40, 90, 10 ]]
```

See [page 48](#) for the output.

The function can be applied to any nested structure:

```
\xintiexpr foo([1, [2, [3, [4, [5, 6, 7, 8, 9, 10]]]])\relax
```

```
[1, [8, [27, [64, [125, 216, 343, 512, 729, 1000]]]]]
```

It must be defined as function acting on scalars, but its value type is not constrained.

```
\xintdefiivar Array = [1..10];
\xintdefiifunc foo(x) = [1..x];
\xintthealign\xintiexpr foo(Array)\relax
```

```
[[ 1 ],
 [ 1, 2 ],
 [ 1, 2, 3 ],
 [ 1, 2, 3, 4 ],
 [ 1, 2, 3, 4, 5 ],
 [ 1, 2, 3, 4, 5, 6 ],
 [ 1, 2, 3, 4, 5, 6, 7 ],
 [ 1, 2, 3, 4, 5, 6, 7, 8 ],
 [ 1, 2, 3, 4, 5, 6, 7, 8, 9 ],
 [ 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 ]]
```

It is even allowed to produce oples and act on oples:

```
\xintdefiivar Ople = 1..10;
\xintdefiifunc bar(x) = x, x^2, x^3;
\xintiexpr bar(Ople)\relax
```

```
1, 1, 1, 2, 4, 8, 3, 9, 27, 4, 16, 64, 5, 25, 125, 6, 36, 216, 7, 49, 343, 8, 64, 512, 9, 81, 729, 10,
100, 1000
```

2.10.5. Using the same name for both a variable and a function

It is licit to overload a variable name (all Latin letters are predefined as dummy variables) with a function name and vice versa. The parsers will decide from the context if the function or variable interpretation must be used (dropping various cases of tacit multiplication as normally applied).

```
\xintdefiifunc f(x):=x^3;
\xinttheiexpr add(f(f),f=100..120)\relax\newline
\xintdeffunc f(x,y):=x^2+y^2;
\xinttheexpr mul(f(f(f),f(f,f)),f=1..10)\relax
\xintunassigniexprfunc{f}\xintunassignexprfunc{f}%
```

```
28205100
```

```
186188134867578885427848806400000000
```

2.10.6. \xintunassignexprfunc, \xintunassigniexprfunc, \xintunassignfloatexprfunc

Function names can be unassigned via `\xintunassignexprfunc{<name>}`, `\xintunassigniexprfunc{<name>}`, and `\xintunassignfloatexprfunc{<name>}`.

```
\xintunassignexprfunc{e}
\xintunassignexprfunc{f}
```

Figure 1: Output of a universal function acting on an array

```

[[[ 1,      8,      27,      64,      125,      216,      343,      512,      729,      1000    ],
 [ 8,      8,      216,      64,      1000,      216,      2744,      512,      5832,      1000    ],
 [ 27,     216,     27,      1728,     3375,      216,      9261,      13824,     729,      27000   ],
 [ 64,     64,     1728,      64,      8000,      1728,     21952,      512,     46656,      8000    ],
 [ 125,    1000,    3375,     8000,      125,     27000,     42875,     64000,     91125,     1000    ],
 [ 216,    216,    216,     1728,     27000,      216,     74088,     13824,     5832,     27000   ],
 [ 343,    2744,    9261,     21952,     42875,     74088,      343,     175616,    250047,    343000   ],
 [ 512,    512,    13824,     512,     64000,     13824,    175616,      512,    373248,     64000   ],
 [ 729,    5832,     729,     46656,     91125,     5832,    250047,    373248,     729,     729000   ],
 [ 1000,   1000,   27000,     8000,     1000,     27000,    343000,    64000,     729000,     1000   ]],
 [[ 8,      8,      216,      64,      1000,      216,      2744,      512,      5832,      1000    ],
 [ 8,      8,      216,      64,      1000,      216,      2744,      512,      5832,      1000    ],
 [ 216,    216,    216,     1728,     27000,      216,     74088,     13824,     5832,     27000   ],
 [ 64,     64,     1728,      64,      8000,      1728,     21952,      512,     46656,      8000    ],
 [ 1000,   1000,   27000,     8000,     1000,     27000,    343000,    64000,     729000,     1000   ],
 [ 216,    216,    216,     1728,     27000,      216,     74088,     13824,     5832,     27000   ],
 [ 2744,   2744,   74088,     21952,    343000,     74088,     2744,     175616,    2000376,    343000   ],
 [ 512,    512,    13824,     512,     64000,     13824,    175616,      512,    373248,     64000   ],
 [ 5832,   5832,   5832,     46656,     729000,     5832,    2000376,    373248,     5832,     729000   ],
 [ 1000,   1000,   27000,     8000,     1000,     27000,    343000,    64000,     729000,     1000   ]],
 [[ 27,     216,     27,      1728,     3375,      216,      9261,      13824,     729,      27000   ],
 [ 216,    216,    216,     1728,     27000,      216,     74088,     13824,     5832,     27000   ],
 [ 27,     216,     27,      1728,     3375,      216,      9261,      13824,     729,      27000   ],
 [ 1728,   1728,   1728,     1728,    216000,     1728,    592704,     13824,     46656,    216000   ],
 [ 3375,   27000,   3375,    216000,     3375,     27000,    1157625,    1728000,     91125,     27000   ],
 [ 216,    216,    216,     1728,     27000,      216,     74088,     13824,     5832,     27000   ],
 [ 9261,   74088,   9261,    592704,    1157625,     74088,     9261,    4741632,    250047,    9261000   ],
 [ 13824,  13824,  13824,    13824,    1728000,    13824,    4741632,    13824,    373248,    1728000   ],
 [ 729,    5832,     729,     46656,     91125,     5832,    250047,    373248,     729,     729000   ],
 [ 27000,  27000,  27000,    216000,     27000,     27000,    9261000,    1728000,    729000,     27000   ]],
 [[ 64,     64,     1728,      64,      8000,      1728,     21952,      512,     46656,      8000    ],
 [ 64,     64,     1728,      64,      8000,      1728,     21952,      512,     46656,      8000    ],
 [ 1728,   1728,   1728,     1728,    216000,     1728,    592704,     13824,     46656,    216000   ],
 [ 64,     64,     1728,      64,      8000,      1728,     21952,      512,     46656,      8000    ],
 [ 8000,   8000,   216000,     8000,     8000,    216000,    2744000,    64000,    5832000,     8000    ],
 [ 1728,   1728,   1728,     1728,    216000,     1728,    592704,     13824,     46656,    216000   ],
 [ 21952,  21952,  592704,     21952,    2744000,    592704,     21952,    175616,    16003008,    2744000   ],
 [ 512,    512,    13824,     512,     64000,     13824,    175616,      512,    373248,     64000   ],
 [ 46656,  46656,  46656,     46656,    5832000,    46656,    16003008,    373248,     46656,    5832000   ],
 [ 8000,   8000,   216000,     8000,     8000,    216000,    2744000,    64000,    5832000,     8000   ]],
 [[ 125,    1000,    3375,     8000,      125,     27000,     42875,     64000,     91125,     1000    ],
 [ 1000,   1000,   27000,     8000,     1000,     27000,    343000,    64000,     729000,     1000    ],
 [ 3375,   27000,   3375,    216000,     3375,     27000,    1157625,    1728000,     91125,     27000    ],
 [ 8000,   8000,   216000,     8000,     8000,    216000,    2744000,    64000,    5832000,     8000    ],
 [ 125,    1000,    3375,     8000,      125,     27000,     42875,     64000,     91125,     1000    ],
 [ 27000,  27000,  27000,    216000,     27000,     27000,    9261000,    1728000,    729000,     27000    ],
 [ 42875,  343000,  1157625,    2744000,     42875,    9261000,     42875,    21952000,    31255875,    343000    ],
 [ 64000,  64000,  1728000,     64000,     64000,    1728000,    21952000,    64000,    46656000,     64000    ],
 [ 91125,  729000,  91125,    5832000,     91125,    729000,    31255875,    46656000,     91125,    729000    ],
 [ 1000,   1000,   27000,     8000,     1000,     27000,    343000,    64000,     729000,     1000   ]]]

```


Warning: no check is done to avoid undefining built-in functions...

2.10.7. `\ifxintverbose` conditional

With `\xintverbosetrue` the meanings of the functions (or rather their associated macros) will be written to the log. For example the `Rump` declaration above generates this in the log file:

```
Function Rump for \xintexpr parser associated to \XINT_expr_userfunc_Rump with meaning macro:#1#2->{\xintAdd {\xintAdd {\xintAdd {\xintDiv {\xintMul {1335}{\xintPow {#2}{6}}}{4}}{\xintMul {\xintPow {#1}{2}}{\xintSub {\xintSub {\xintSub {\xintMul {11}{\xintMul {\xintPow {#1}{2}}{\xintPow {#2}{2}}}}{\xintPow {#2}{6}}}{\xintMul {121}{\xintPow {#2}{4}}}{2}}}{\xintDiv {\xintMul {11}{\xintPow {#2}{8}}}{2}}{\xintDiv {#1}{\xintMul {2}{#2}}}}
```

The meanings written out to the log for more complicated functions may sometimes use the same character at different locations but with different catcodes.

It may thus be impossible to retokenize it (even after having removed the extra spaces from the added line breaks).

This is in contrast with variable values which are always output in the log in the benign way, using digits, braces and some characters of catcode 12.

2.10.8. `\ifxintglobaldefs` conditional

If true user defined variables (`\xintdefvar`, ...) and functions (`\xintdeffunc`, ..., `\xintNewFunction`) for the expression parsers, as well as macros obtained via `\xintNewExpr` et al. have global scope. If false (default) they have local scope.

2.10.9. `\xintNewFunction`

This is syntactic sugar which allows to use notation of functions for what is nothing more in disguise than a \TeX macro. Here is an example:

```
\xintNewFunction {foo}[3]{add(mul(x+i, i=#1..#2),x=1..#3)}
```

We now have a genuine function `foo()` of three variables which can be used in *all three* parsers.

```
\xintexpr seq(foo(0, 3, j), j= 1..10)\relax
```

```
24, 144, 504, 1344, 3024, 6048, 11088, 19008, 30888, 48048
```

Each time the created “macro-function” `foo()` will be encountered the corresponding replacement text will get inserted as a sub-expression (of the same type as the surrounding one), the macro parameters having been replaced with the (already evaluated) function arguments, and the parser *will then have to parse the expression*. It is very much like a macro substitution, but with parentheses and comma separated arguments (which can be arbitrary expressions themselves).

It differs fundamentally from `\xintdeffunc` as it realizes no pre-parsing whatsoever of the associated sub-expression; using it shortens the input but not the parsing time (which however is most of the time negligible compared to actual numerical computations). Use it for syntax which `\xintdeffunc` does not parse successfully.

2.11. Examples of user defined functions

2.11.1. Example with vectors and matrices

Suppose we want to manipulate 3-dimensional vectors, which will be represented as `nuples` of length 3. And let's add a bit of matrix algebra.

```
\xintdeffunc dprod(V, W) := V[0]*W[0] + V[1]*W[1] + V[2]*W[2];
\xintdeffunc cprod(V, W) := [V[1]*W[2] - V[2]*W[1],
                             V[2]*W[0] - V[0]*W[2],
```

```

                                V[0]*W[1] - V[1]*W[0]];
\xintdeffunc Det3(U, V, W) := dprod(cprod(U, V), W);
\xintdeffunc DetMat(M) = Det3(*M);
\xintdeffunc RowMat(U, V, W) := [U, V, W];
\xintdeffunc ColMat(U, V, W) := [[U[0], V[0], W[0]],
                                [U[1], V[1], W[1]],
                                [U[2], V[2], W[2]]];
\xintdeffunc MatMul(A, B) :=
[[A[0,0]*B[0,0]+A[0,1]*B[1,0]+A[0,2]*B[2,0],
  A[0,0]*B[0,1]+A[0,1]*B[1,1]+A[0,2]*B[2,1],
  A[0,0]*B[0,2]+A[0,1]*B[1,2]+A[0,2]*B[2,2]],
 [A[1,0]*B[0,0]+A[1,1]*B[1,0]+A[1,2]*B[2,0],
  A[1,0]*B[0,1]+A[1,1]*B[1,1]+A[1,2]*B[2,1],
  A[1,0]*B[0,2]+A[1,1]*B[1,2]+A[1,2]*B[2,2]],
 [A[2,0]*B[0,0]+A[2,1]*B[1,0]+A[2,2]*B[2,0],
  A[2,0]*B[0,1]+A[2,1]*B[1,1]+A[2,2]*B[2,1],
  A[2,0]*B[0,2]+A[2,1]*B[1,2]+A[2,2]*B[2,2]]];

\xintdefvar vec1, vec2, vec3 := [1, 1, 1], [1, 1/2, 1/4], [1, 1/3, 1/9];
\xintdefvar mat1 = RowMat(vec1, vec2, vec3);
\xintdefvar mat2 = ColMat(vec1, vec2, vec3);
\xintdefvar mat12 = MatMul(mat1,mat2);
\xintdefvar mat21 = MatMul(mat2,mat1);
Some computations (|align| executes multiple times hence we pre-computed!):
\begin{align*}
M_1      &= \vcenter{\xintthealign \xintexpr mat1\relax}&&\quad
M_2 . M_1 = \vcenter{\xintthealign \xintexpr mat21\relax}\jot
M_2      &= \vcenter{\xintthealign \xintexpr mat2\relax}&&\quad
M_1 . M_2 = \vcenter{\xintthealign \xintexpr mat12\relax}
\end{align*}
$$
\det(M_1)      = \xinteval{DetMat(mat1)},\quad
\det(M_1.M_2) = \xinteval{reduce(DetMat(mat12))},\quad
\det(M_2.M_1) = \xinteval{reduce(DetMat(mat21))}
$$

```

Some computations (align executes multiple times hence we pre-computed!):

$$\begin{array}{ll}
 \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1/2 & 1/4 \\ 1 & 1/3 & 1/9 \end{bmatrix}, & \begin{bmatrix} 3 & 11/6 & 49/36 \\ 11/6 & 49/36 & 251/216 \\ 49/36 & 251/216 & 1393/1296 \end{bmatrix}, \\
 M_1 = \begin{bmatrix} 1 & 1/2 & 1/4 \\ 1 & 1/3 & 1/9 \end{bmatrix} & M_2.M_1 = \begin{bmatrix} 11/6 & 49/36 & 251/216 \\ 49/36 & 251/216 & 1393/1296 \end{bmatrix} \\
 \\
 \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1/2 & 1/3 \\ 1 & 1/4 & 1/9 \end{bmatrix}, & \begin{bmatrix} 3 & 7/4 & 13/9 \\ 7/4 & 21/16 & 43/36 \\ 13/9 & 43/36 & 91/81 \end{bmatrix}, \\
 M_2 = \begin{bmatrix} 1 & 1/2 & 1/3 \\ 1 & 1/4 & 1/9 \end{bmatrix} & M_1.M_2 = \begin{bmatrix} 7/4 & 21/16 & 43/36 \\ 13/9 & 43/36 & 91/81 \end{bmatrix} \\
 \\
 \det(M_1) = -1/18, \quad \det(M_1.M_2) = 1/324, \quad \det(M_2.M_1) = 1/324
 \end{array}$$

For some hair-raising experience check the `\xintverbosetrue` output in the log... here is an alternative with two (three, counting `dprod()`) helper functions:

```

% annoying that Tr also starts Trace, but Spur is available
% well Sp also starts Spectrum. Big problems.
\xintdeffunc Tr(M) :=
[[M[0,0], M[1,0], M[2,0]],
 [M[0,1], M[1,1], M[2,1]],

```

```
[M[0,2], M[1,2], M[2,2]]];

\xintdeffunc MatMul_a(r1, r2, r3, c1, c2, c3) :=
  [[dprod(r1, c1), dprod(r1, c2), dprod(r1, c3)],
   [dprod(r2, c1), dprod(r2, c2), dprod(r2, c3)],
   [dprod(r3, c1), dprod(r3, c2), dprod(r3, c3)]];

```

```
\xintdeffunc MatMul(A, B) := MatMul_a(*A, *Tr(B));
```

And once we have the transpose and the scalar product of vectors, we can simply use `ndmap()` for a lean syntax (this would extend to arbitrary dimension):

```
\xintdeffunc MatMul(A, B) = ndmap(dprod, *A; *Tr(B));
\xintdefvar mat1212 = MatMul(mat12, mat12);
\begingroup
\def\xintexprPrintOne      #1{\xintTeXFrac{#1}}%
\def\xintexpralignbegin   {\begin{pmatrix}}%
\def\xintexpralignend     {\end{pmatrix}}%
\def\xintexpralignlinesep {\noexpand\\[2\jot]}% counteract an internal \expanded
\def\xintexpraligninnersep {&}%
\let\xintexpralignleftbracket\empty \let\xintexpralignleftsep\empty
\let\xintexpralignrightbracket\empty \let\xintexpralignrightsep\empty
$$ \xintthealign \xintexpr mat1\relax \cdot \xintthealign \xintexpr mat2\relax \cdot
  \xintthealign \xintexpr mat1\relax \cdot \xintthealign \xintexpr mat2\relax =
  \xintthealign \xintexpr mat12\relax ^2 = \xintthealign \xintexpr mat1212\relax$$
$$ \det(M_1\cdot M_2 \cdot M_1 \cdot M_2) = \xinteval{reduce(DetMat(mat1212))}$$
\endgroup

```

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & \frac{1}{2} & \frac{1}{4} \\ 1 & \frac{1}{3} & \frac{1}{9} \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 & 1 \\ 1 & \frac{1}{2} & \frac{1}{3} \\ 1 & \frac{1}{4} & \frac{1}{9} \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 & 1 \\ 1 & \frac{1}{2} & \frac{1}{4} \\ 1 & \frac{1}{3} & \frac{1}{9} \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 & 1 \\ 1 & \frac{1}{2} & \frac{1}{3} \\ 1 & \frac{1}{4} & \frac{1}{9} \end{pmatrix} = \begin{pmatrix} 3 & \frac{7}{4} & \frac{13}{9} \\ \frac{7}{4} & \frac{21}{16} & \frac{43}{36} \\ \frac{13}{9} & \frac{43}{36} & \frac{91}{81} \end{pmatrix}^2 = \begin{pmatrix} \frac{18337}{1296} & \frac{48067}{5184} & \frac{93853}{11664} \\ \frac{48067}{5184} & \frac{128809}{20736} & \frac{253687}{46656} \\ \frac{93853}{11664} & \frac{253687}{46656} & \frac{501289}{104976} \end{pmatrix}$$

$$\det(M_1 \cdot M_2 \cdot M_1 \cdot M_2) = \frac{1}{104976}$$

2.11.2. Example with the Rump test

Let's try out our `Rump()` function:

```
\xinttheexpr Rump(77617,33096)\relax.
```

`-54767/66192`. Nothing problematic for an exact evaluation, naturally!

Thus to test the Rump polynomial (it is not quite a polynomial with its `x/2y` final term) with floats, we *must* also declare `Rump` as a function to be used there:

```
\xintdeffloatfunc
  Rump(x,y):=333.75 y^6 + x^2 (11 x^2 y^2 - y^6 - 121 y^4 - 2) + 5.5 y^8 + x/2y;
```

The numbers are scanned with the current precision, hence as here it is 16, they are scanned exactly in this case. We can then vary the precision for the evaluation.

```
\def\CR{\cr}
\halign
{\tabskiplex
\hfil\bfseries#\&\xintDigits:=\xintloopindex\relax
  \xintthefloatexpr Rump(77617,33096)#\cr
\xintloop [8+1]
\xintloopindex &\relax\CR
\ifnum\xintloopindex<40 \repeat

```

```

}
8 7e29
9 -1e28
10 5e27
11 -3e26
12 4e25
13 3e24
14 3e23
15 -2e22
16 1e21
17 -5e20
18 1.17260394005317863
19 1.00000000000000000001e18
20 -9.9999999999999998827e16
21 1.00000000000000011726e16
22 3.000000000000001172604e15
23 -9.999999999998827396060e13
24 -1.9999999999988273960599e13
25 -1.99999999998827396059947e12
26 1.1726039400531786318588349
27 -5.9999999988273960599468214e10
28 -9.99999988273960599468213681e8
29 2.0000000117260394005317863186e8
30 1.00000011726039400531786318588e7
31 -999998.8273960599468213681411651
32 200001.17260394005317863185883490
33 -9998.82739605994682136814116509548
34 -1998.827396059946821368141165095480
35 -198.82739605994682136814116509547982
36 21.1726039400531786318588349045201837
37 -0.8273960599468213681411650954798162920
38 -0.82739605994682136814116509547981629200
39 -0.827396059946821368141165095479816292000
40 -0.8273960599468213681411650954798162919990

```

2.11.3. Examples of recursive definitions

Recursive definitions *require* using the short-circuit branching operators.

```
\xintdeffunc GCD(a,b):=(b)?{GCD(b,a/:b)}{a};
```

This of course is the Euclidean algorithm: it will be here applied to variables which may be fractions. For example:

```
\xinttheexpr GCD(385/102, 605/238)\relax
55/714
```

There is already a built-in `gcd()` (which accepts arbitrarily many arguments):

```
\xinttheexpr gcd(385/102, 605/238)\relax
55/714
```

Our second example is modular exponentiation:

```
\xintdefiifunc powmod_a(x, m, n) :=
  isone(m)?
    % m=1, return x modulo n
    { x /: n }
  % m > 1 test if odd or even and do recursive call
  { odd(m)? { x*sqr(powmod_a(x, m//2, n)) /: n }
```

```

        {   sqr(powmod_a(x, m//2, n)) /: n }
    }
;
\intdefiifunc powmod(x, m, n) := (m)?{powmod_a(x, m, n)}{1};

```

I have made the definition here for the `\xintiexpr` parser; we could do the same for the `\xintexpr`-parser (but its usage with big powers would quickly create big denominators, think `powmod(1/2, 1000, 1)` for example.)

```

\xinttheiexpr seq(powmod(x, 1000, 128), x=9, 11, 13, 15, 17, 19, 21)\relax\par
65, 97, 33, 1, 1, 33, 97

```

The function assumes the exponent is non-negative (the Python `pow` behaved the same until 3.8 release), but zealous users will add the necessary code for negative exponents, after having defined another function for modular inverse!

If function **A** needs function **B** which needs function **A** start by giving to **B** some dummy definition, define **A**, then define **B** properly. TODO: add some example here...

2.12. Links to some (old) examples within this document

- The utilities provided by `xinttools` (section 6), some completely expandable, others not, are of independent interest. Their use is illustrated through various examples: among those, it is shown in subsection 7.8 how to implement in a completely expandable way the [Quick Sort algorithm](#) and also how to illustrate it graphically. Other examples include some dynamically constructed alignments with automatically computed prime number cells: one using a completely expandable prime test and `\xintApplyUnbraced` (subsection 7.2), another one with `\xintFor*` (subsection 7.6).
- One has also a [computation of primes within an \edef](#) (subsection 6.15), with the help of `\xintiloop`. Also with `\xintiloop` an [automatically generated table of factorizations](#) (subsection 7.5).
- The code for the title page fun with Fibonacci numbers is given in subsection 3.18 with `\xintFor*` joining the game.
- The computations of π and $\log 2$ (subsection 15.11) using `xint` and the computation of the [convergents of e](#) with the further help of the `xintcfrac` package are among further examples.
- Also included, an [expandable implementation of the Brent-Salamin algorithm](#) for evaluating π .
- The [subsection 7.4](#) implements expandably the Miller-Rabin pseudo-primality test.
- The functionalities of `xintexpr` are illustrated with various other examples, in [subsubsection 2.10.1, Functions with dummy variables](#), [subsection 7.1](#) or [Recursive definitions](#).

2.13. Oples and nutples: the 1.4 terminology

Skip this on first reading, else you will never start using the package. SKIP THIS! (understood?)

In this section I will describe a mathematical terminology which models how the parser handles the input syntax with numbers, commas, and brackets, and how it maps internally to \TeX specific concept, particularly braces and macro arguments.

.13.1	Base terminology	54
.13.2	Items (and sub-items) versus elements	55
.13.3	Oples as trees	55
.13.4	Ople slicing and indexing	56
.13.5	Nested slicing of oples	57
.13.6	Function arguments versus variables	57
.13.7	Final words on leaves	58
.13.8	Farewell, thanks for your visit!	58

2.13.1. Base terminology

We start with a set \mathcal{A} of *atoms*, which represent numeric data. In \TeX syntax such *atoms* are always braced, more precisely, currently they look like

`{raw format within \TeX braces}`

The \TeX braces are not set-theoretical braces here, they are simply used for \TeX nical reasons (one could imagine using rather some terminator token, but ultimately support macros for built-in and user defined functions rely on \TeX macros with undelimited parameters, at least so far).

Our category C of “oples” is the smallest collection of *totally ordered finite sets* verifying these properties:

1. The empty set \emptyset is an *ople*, i.e. it belongs to C .
2. Each singleton set $\{0\}$ whose element 0 is either an *atom* $a \in \mathcal{A}$ or an *ople* qualifies as an *ople*.
3. C is stable by concatenation.

Notes:

- We refer to the empty set \emptyset via the variable *nil*.¹⁷
- It is convenient to accept the empty set as being also an *atom*. If this is done, then we may refer to the original *atoms* (elements of \mathcal{A}) as *non empty numerical data*.
- Concatenation is represented in the syntax by the comma. Thus repeated commas are like only one and *nil* is a neutral element.
- A singleton *ople* $\{a\}$ whose single element is a (non-empty) *atom* is called a *number*.¹⁸
- The operation of constructing $\{0\}$ from the *ople* 0 is called *bracing* (set theory, \TeX), or *bracketing* (`xintexpr` input syntax, Python *lists*), or *packing* (as a reverse to Python's unpacking of sequence type objects). In the expression input syntax it corresponds to enclosing 0 within square brackets: $[0]$.
- A braced *ople* is called a *nutple*. Among them $\{\text{nil}\}$ (aka $\{\emptyset\}$) is a bit special. It is called the *none-ple*.¹⁹ It is not *nil*.²⁰

Each *ople* has a *length* which is its cardinality as set. The singleton *oples* are called *one-ples*. There are thus two types of *one-ples*:

- *numbers* $\{a\}$, $a \in \mathcal{A}$,
- *nutples* $\{0\}$, $0 \in C$.

If we consider the empty set *nil* on the same footing as *atoms*, the two types have only one common object which is the *none-ple*. As a rule arithmetic operations will either break or silently convert the *none-ple* to the zero value:

```
\xinteval{3+[], 5^[], 10*[]}
```

`3, 1, 0`. But attention that `\xintiieval` in contrast to `\xinteval` is broken by such inputs.

¹⁷ There is actually a built-in variable with this name. At 1.4, `\xintexpr \relax` is legal and also generates the *nil*.¹⁸ This has to be taken in a general sense, for example with `polexpr`, polynomials are represented by such “numbers”.¹⁹ Prior to version 1.4j of this documentation it was called the *not-ple*.²⁰ There is (experimental) a pre-defined “None” variable which stands for the *none-ple*. It can also be input as `[]`.

2.13.2. Items (and sub-items) versus elements

In order to illustrate these concepts, let us consider how one should interpret notation such as `3,5,7,9` when it arises in an `\xintexpression`:

tempting vocabulary: Each of `3`, `5`, `7`, and `9` is an *item*, or *element* of the (comma separated) *list*. In other terms we have here a list with 4 items.

rigorous vocabulary: each one of `3`, `5`, `7`, `9` stands for an *ople* (of the *one-ple* type) and `3,5,7,9` stands for their *concatenation*.

It is important to understand that in an `\xintexpression`, there is no difference between `3,5,7` and `3,,,,5,,,,,,7`. So the view of the comma as separator is misleading. In other terms, the comma is NOT a separator but the (associative) operator of concatenation of totally ordered sets, and the number `3` for example represents a (singleton) set.

If we want to refer to `3` or `5` or `7` or `9` as “the items of the (open) list `3,5,7,9`” (and probably this documentation already has such utterances, due to legacy reasons from the pre-1.4 internal model), we *must* realize that this clashes with using the word *item* as synonymous to *element* in the set-theoretical sense.

To repeat, any *ople* `0` is a finite totally ordered set: if not the empty set, it has *elements* a_1, \dots, a_k , and the above means that its *items* are the singleton oples (aka one-ples) $I_1 = \{a_1\}, \dots, I_k = \{a_k\}$. Each a_j may be an *atom*, then I_j is a *number*, or a_j is an *ople* (possibly the empty set), then I_j is a *nutple* whose depth is one more than the one of the *ople* a_j .

Thus we can refer to “items” but must then understand they are not “elements”: “items” are “singleton sub-sets”. The cardinality (aka length) of an *ople* is also the number of its items. It would be tempting to use the terminology “sub-item” to keep in mind they are “sub-sets” but this would again create confusion: a *nutple* has only one item which is itself; and we need some terminology to refer to the individual numbers in the *nutple* given in input as `[1,2,3]` for example. It is natural to refer to `1`, `2`, `3` as “sub-items” of `[1,2,3]` as the latter may be an “item” (it is in particular an “item” of itself, the unique one at that).

We distinguish the *oples* of length zero (there is only one, the empty set) or at least two as those which can never be an “item”. Those of length one, the *one-ples*, are exactly those which can be “items”. Among them some may have “sub-items”, they are the *nutples* with the exception of the *none-ple*. And the others do not have “sub-items”, they are the *numbers* and the *none-ple* (whose input syntax is either `[]` or the variable `None`).²¹

2.13.3. Oples as trees

We say that the empty set `nil` and *atoms* are *leaves*.

We associate with any *ople* a tree. The root is the *ople*. In the case of the `nil` *ople*, there is nothing else than the root, which we then consider also a *leaf*. Else the children at top level are the successive *elements* (not “items”!) of the *ople*.²² Among the elements some are *atoms* giving *leaves* of the tree, others are *nutples* which in turn have children. In the special case of the *none-ple* we consider it has a child, which is the empty set and this is why we consider the empty

²¹ A note on the `\xintverbosetrue` regime: for a variable defined to be `3,5,7,9`, it will say that its value is `{3}{5}{7}{9}`, because it does not keep the external set-theoretical braces. The braces here are only TeX braces, and `{3}` is an *atom*. The *number* would be `{{3}}` with the external braces being set-theoretical and also used internally as TeX braces. From the four numbers `{{3}}`, ..., `{9}` concatenation gives `{{3}{5}{7}{9}}`, which is the *ople* `3,5,7,9`. But the log view drops deliberately the external braces. If the variable is defined to be the *nutple* `[3,5,7,9]`, then the log view will be `{{3}{5}{7}{9}}` (up to details on how exactly the numeric quantities are coded) and the actual internal TeX entity will be `{{{{3}}{5}{7}{9}}}`, where the two external layers of braces are both set-theoretical and TeXnical braces. ²² We could also consider a tree for which the children of the root node would be its items and recursively; in that case the leaves would be *numbers* and possibly the `None`. The tree of the `nil` would be the empty tree, the tree of `None` would have a single node and no edges. Such a tree would match the input syntax (of course applying the rule that iterated commas are like only one). The tree which is described in this section matches more directly the internal syntax, hence is more useful to the author, who is also the sole reader who extracts some benefit from reading this documentation once in a while.

set `nil` to be also a potential *leaf*. We then proceed recursively. We thus obtain from the root *ople* a tree whose vertices are either *oples* or *leaves*. Only the empty set `nil` is both a *leaf* and an *ople*.

Considering the empty set `nil` as an *atom* fits with the `xintexpr` internal implementation based on \TeX : `nil` is an empty pair of braces {}, whereas an *atom* is a braced representation of a numeric value using digits and other characters. We construct *oples* by putting one after the other such constituents and bracing them, and then repeating the process recursively.

It has also an impact on the definition of the *depth* (a.k.a as *maximal dimension*) of an *ople*. For example the *ople* $\{\emptyset A_1 A_2\}$ with three elements, among them the empty set and two atoms is said to have depth 1, or to have maximal dimension 1. And $\{\{\emptyset\} A_1 A_2\}$ is of depth 2 because it has a leaf (the empty set) which is a child of a child of the *ople*. NumPy *ndarrays* have a more restricted structure for example $\{\{A_{00} A_{01}\} \{A_{10} A_{11}\}\}$ is a 2-dimensional array, where all leaves are at the same depth. When slicing empties the array from its atoms, NumPy keeps the shape information but prints the array as []. This will not be the case with `xintexpr`, which has no other way to indicate the shape than display it.

```
\xinteval{[[[]],[[]]}
[[[]], [[]]]
\xinteval{[[[0,1],[10,11]][:,:2]}
[[[]], [[]]]
```

2.13.4. Ople slicing and indexing

“Set-theoretical” slicing of an *ople* means replacing it with one of its subsets. This applies also if it is a *number*. Then it can be sliced only to itself or to the empty set (indeed it has only one element, which is an *atom*). Similarly the *none-ple* can only be sliced to give itself or the empty set. And more generally a *nutple* is a singleton so also can only be set-sliced to either the empty set or itself.

`xintexpr` extends “Python-like” slicing to act on *oples*:

- if they are not *nutples* set-theoretical slicing applies,
- if they are *nutples* (only case having a one-to-one correspondence in Python) then the slicing happens *within brackets*: i.e. the *nutple* is unpacked then the set-theoretical slicing is applied, then the result is *repacked* to produce a new *nutple*.

With these conventions the *none-ple* for example is invariant under slicing: unpacking it gives the empty set, which has only the empty set as subset and repacking gives back the *none-ple*. Slicing a general *nutple* returns a *nutple* but now of course in general distinct from the first one.

The input syntax for Python slicing is to postfix a variable or a parenthesized *ople* with `[a:b]`. See [subsection 2.6](#) for more. There are never any out-of-range errors when slicing or indexing. All operations are licit and resolved by the `nil`, a.k.a. empty set.

“Set-theoretical” item indexing of an *ople* means reducing it to a subset which is a singleton. It is thus a special case of set-theoretical slicing (which is the general process of selecting a subset as replacement of a set).

`xintexpr` extends “Python-like” indexing to act on *oples*:

- if they are not *nutples* set-theoretical item indexing applies,
- if they are *nutples* (only case having a one-to-one correspondence in Python) then the meaning becomes *extracting*: i.e. the *nutple* is unpacked then the set-theoretical indexing is applied, but the result is *not repacked*.

For example when applied to the *none-ple* we always obtain the `nil`. Whereas as we saw slicing the *none-ple* always gives back the *none-ple*. Indexing is denoted in the syntax by postfixing by `[N]`. Thus for *nutples* (which are analogous to Python objects), there is genuine difference between the `[N]` extractor and the `[N:N+1]` slicer. But for *oples* which are either `nil`, a *number*, or of length at least 2, there is no difference.

2.13.5. Nested slicing of oples

Nested slicing is a concept from NumPy, which is extended by `xintexpr` to trees of varying depths. We have a chain of slicers and extractors. I will describe only the case of slicers and letting them act on a `nutple`. The first slicer gives back a new `nutple`. The second slicer will be applied to each of one of its remaining elements. However some of them may be `atoms` or the empty set. In the NumPy context all leaves are at the same depth thus this can happen only when we have reached beyond the last dimension (axis). This is not permitted by NumPy and generates an error. `xintexpr` does not generate an error. But any attempt to slice an `atom` or the empty set (as element of its container) removes it. Recall we call them `leaves`. We can not slice leaves. We can only slice non-leaf elements: such items are necessarily `nutples`. The procedure then applies recursively.

If we handle an extractor rather than a slicer, the procedure is similar: we can not extract out of an `atom` or the empty set. They are thus removed. Else we have a `nutple`. It is thus unpacked and replaced by the selected element. This element may be an `atom` or the empty set and any further slicer or extractor will remove them, or it is a `nutple` and the procedure applies with the next slicer/extractor.

`xintexpr` allows to apply such a `[a:b,c:d,N,e:f,...]` chain of slicing/extracting also to an `ople`, which is not a `nutple`. We simply apply the first step as has been described previously and successive steps will only get applied to either `nutples` or `leaves`, the latter getting silently removed by any attempted operation.

2.13.6. Function arguments versus variables

In a function declaration with `\xintdeffunc`, the call signature is parsed as a comma separated list, so here it is not true that repeated commas are like only one: repeated commas are not allowed and will break the function declaration.

When `xintexpr` parses a function call, it first constructs the `ople` which is delimited by the opening and closing parentheses, then it applies the function body, after having mapped the successive items (not the elements) of the parsed `ople` to the variables appearing in the function call signature. Hence the arguments in the call signature stand for `one-ples` (i.e. either `numbers` or `nutples`).

Let me explain why we can not define a function `foo(A,B)` of two `oples`: the function call will evaluate as an `ople` what is enclosed within the parentheses. It is then impossible in general to split this uniquely into two `oples` `A` and `B`, except if for example we know a priori the length of `A`. We could imagine defining a declarative interface for a `foo(A,B)` with `A` preset to have 37 items or at least a pre-defined number of items but this is extraneous layer for a functionality no-one will use.

The alternative would be to consider that declaring `foo(A,B)` means `A` will pick-up always the first item and `B` all the remaining ones, and thus will be an `ople`; here, there are some technical implementation reasons which have dissuaded the author to do this.

In its place, a special syntax `foo(A,*B)` for the declaration of the function is available. It means that `B` stands for the `nutple` which receives as items all arguments in the function call beyond the first one already assigned to `A`.

More generally, the last positional argument in a function declaration can have the form `*⟨argname⟩`. This then means that `⟨argname⟩` represents a `nutple` which will receive as items all arguments in the function call remaining after the earlier positional arguments have been assigned. The declared function body is free to again use the syntax `*⟨argname⟩` which will unpack it and thus produce the `ople` concatenating all such optional arguments.

With `\xintdefvar` one can define a variable with value an `ople` of arbitrary cardinality. Such a variable can be used in a function call, it will then occupy the place of as many arguments as its cardinality (which is its number of elements, hence of its associated items). For example if function `foo` was declared as a function of 5 arguments `f(a,b,c,d,e)` it is legitimate to use it as `f(A,B)` if `A` is an `ople`-valued variable of length three and `B` of length two. The actual arguments

a,b,c,d,e will be made to match the three items of A and the two items of B.

2.13.7. Final words on leaves

In case things were too clear, let's try to add a bit of confusion with an extra word on *leaves*. When we discuss informally (particularly to compare with NumPy) an input such as

```
[[1, 2], [3, 4]]
```

we may well refer to 1, 2, 3, and 4 as being “the leaves of the 2d array”. But obviously we have here numbers and previously we explained that a number is not a *leaf*, its *atom* is. Well, the point here is that we must make a difference between the input form as above and the actual constructed *ople* the parser will obtain out of it. In the input we do have numbers. The comma is a *concatenator*, it is not a separator for enumeration! The *ople* which corresponds to it has a \TeX representation like this:

```
{{{1}{2}}}{3}{4}}
```

where we don't have the *numbers* anymore (which would look like $\{\{1\}\}$, $\{\{2\}\}$, ...) but numeric *atoms* $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$ where the braces are \TeX braces and **not** set-theoretical braces (the other braces are both). Hence we should see the above as the *ople* $\{A_{00}A_{01}\}\{A_{10}A_{11}\}$ with atoms $A_{00} = \{1\}$, ..., being the *leaves* of the tree associated to (or which is) the *ople*.

Numbers may be called the *leaves* of the **input**, but once parsed, the input becomes an *ople* which is (morally) a tree whose leaves are *atoms* (and the empty set). This discussion can also be revisited with footnote 22 in mind.

2.13.8. Farewell, thanks for your visit!

I hope this is clear to everyone. If not, maybe time to say this section is not needed to understand almost all of the manual, but I needed to write it to be able to maintain in future my own software.

2.14. Expansion (for geeks only)

As mentioned already, the parsers are compatible with expansion-only context.

Also, they expand the expression piece by piece: the normal mode of operation of the parsers is to unveil the parsed material token by token. Unveiling is a process combining space swallowing, brace removal (one level generally), and *f-expansion*.

For example a closing parenthesis after some function arguments does not have to be immediately visible, it and the arguments themselves may arise from *f-expansion* (applied before grabbing each successive token). Even the ending `\relax` may arise from expansion. Even though the `\xinteval` user interface means that the package has at some point the entire expression in its hands, it immediately re-inserts it into token stream with an additional postfixed `\relax` and from this point on has lost any ways (a simple-minded delimited macro won't do because the expression is allowed to contain sub-`\xintexpressions`, even nested) to manipulate formally again the whole thing; it can only re-discover it one token at a time.

This general behaviour (which allows much more freedom in assembling expressions than is usually the case with familiar programming languages such as Python, although admittedly that freedom will prove useful only to power- \TeX users and possibly does not have that many significant use cases) has significative exceptions. These exceptions are mostly related to “pseudo”-functions. A “pseudo”-function will grab some of its arguments via delimited macros. For example `subs(expr1,x=expr2)` needs to see the comma, equal sign and closing parenthesis. But it has mechanisms to allow `expr1` and `expr2` to possess their own commas and parentheses.

Inner semi-colons on the other hand currently always can originate from expansion. Defining functions or variables requires a visible semi-colon acting as delimiter of the expression, but inner semi-colons do not need to be hidden within braces or macros.

The expansion stops only when the ending `\relax` has been found (it is then removed from the token stream).

For catcode related matters see `\xintexprSafeCatcodes`.

A word of warning on the bracketed optional argument of respectively `\xintfloatexpr` and `\xintiexpr`. When defining macros which will hand over some argument to one of these two parsers, the argument may potentially start with a left square bracket `[` (e.g. argument could be `[1, 2, 3]`) and this will break the parser. The fix is to use in the macro definition `\xintfloatexpr\empty`. This extra `\empty` token will prevent the parser from thinking there is an optional argument and it will then disappear during expansion.

If comparing to other languages able to handle floating point numbers or big integers, such as Python, one should take into account that what the `xint` packages manipulate are streams of ascii bytes, one per digit. At no time (due to expandability) is it possible to store intermediate results in an arithmetic CPU register; each elementary operation via `\the\numexpr` will output digit tokens (hence as many bytes), not things such as handles to memory locations where some numbers are stored as memory words. The process can never put aside things but can only possibly permute them with upcoming tokens, to use them later, or, via combinations of `\expanded` and `\unexpanded` or some other more antiquated means grab some tokens and shift the expansion to some distant locations to later come back. The process is a never-ending one-dimensional one...

2.15. Known bugs/features (last updated at 1.4m)

`\xinteval{\xintLength{\par\par\par}}` complains about a Runaway argument:

`\xintLength{\par\par\par}` has no issue as `\xintLength` is a `\long` macro but this is not the case of `\xinteval`. Most macros of a non arithmetic nature in `xintkernel` and `xinttools` are declared `\long` but absolutely none in `xintexpr`, and its dependencies `xint`, etc... As a remark in passing, I could not use the \TeX `\item` directly:²³

```
Runaway argument?
{|\xinteval {\xintLength {
! Paragraph ended before \@item was complete.
<to be read again>
      \par
1.4434 \item[{|\xinteval{\xintLength{\par
      \par\par}| complains about a Runawa...
```

This is the reason I guess why everything is a priori `\long` in the \TeX 3 interface except if asked for otherwise (as far as I know).

Although most macros are dealing with inputs which can only be with digits and some other character tokens, it would still be quite some work to chase all top-level ones. Besides, in practice, it does help better locate ill-formed input.

`if(100>0,(100,125),(100,128))` breaks my code: This is a feature. This is a syntax error, as the comma serves to concatenate "oples" (see [subsection 2.13](#)), and parentheses do not create analogs of "tuples", so this input is parsed the same as

```
if(100>0,100,125,100,128)
```

which is an error as `if()` requires exactly three arguments, not five. Use:

```
if(100>0,[100,125],[100,128])
```

which will expand to the "tuple" `[100,125]`.

`\xintdeffunc foo(x):= gcd((x>0)?{[x,125]}{[x,128]})`; creates a broken function: Bug. Normally `gcd()` (and other multi-arguments functions) work both with open lists of arguments or bracketed lists ("nutples") and the above syntax would work perfectly fine in numerical context. But the presence of the `?` breaks in `\xintdeffunc` context the flexibility of `gcd()`.

Currently working alternatives:

²³ For those who wonder my custom `\verb` employs a `\scantokens` approach, so it can be used in the argument of a macro, for example `\footnote {\verb |\verb |}`.

TOC, Start here, [xintexpr](#), [xintexpr \(old doc\)](#), [xinttrig](#), [xintlog](#), [xinttools](#), [Examples](#), [xint bundle](#)

```
\xintdeffunc foo(x) := gcd(if(x>0, [x,125], [x,128]));
\xintdeffunc foo(x) := if(x>0, gcd(x,125), gcd(x,128));
\xintdeffunc foo(x) := if(x>0, gcd([x,125]), gcd([x,128]));
\xintdeffunc foo(x) := gcd((x>0){x,125}{x,128});
\xintdeffunc foo(x) := (x>0){gcd(x,125){gcd(x,128)};
\xintdeffunc foo(x) := (x>0){gcd([x,125]){gcd([x,128])};
```

The same problem will arise with an `??` nested inside `gcd()` or similar functions, in an `\xintdeffunc`.

`\xinteval{0^-0.5}` says "0 raised to power -1" Feature. Half integer exponents are handled via a square-root extraction, so here `xintexpr` wanted to first raise 0 to power -1, as reported.

Comparison operator == crashes with nutples Not yet implemented...

I liked the "broadcasting" `[1..10]^10` syntax, but it was removed at 1.4 Patience... `seq(x^10,x=1..10)` is alternative (add external `[..]` to get a nutple).

`1e\numexpr 5+2\relax` crashes Not clear yet if bug or feature. The syntax accepted in the scientific part is limited, and failure is expected: hitting a `\numexpr` when parsing a number triggers insertion of a tacit multiplication and then `1e` is missing the scientific exponent. The same happens with `1e(2+3)`. Use syntax such as `1e\the\numexpr5+2\relax`, or `1e\xinteval{5+2}` (although here this relies on output format of `\xinteval` using integer notation with no decoration in this case).

`seq(1e-i,i=1..5)` crashes Not clear if bug or feature. Use `seq(1e\xinteval{-i},i=1..5)` or, as a possibly faster way `seq(1e\xintiieval{-i},i=\xintiexpr1..5\relax)`.

`omit/abort` if nested and not last in the sub-expression cause a crash For example `seq(subs((i){i}{abort},t=i)+10, i=-2, -1, 0, 1)` crashes, due to the presence of the `+10`. This is a long-standing limitation, applying ever since `omit/abort` were added to the syntax at 1.1. Even without the `+10` the nested case was broken by a 1.4 regression and got fixed only at 1.4h. The non-nested case `seq((i){i}{abort}+10, i=-2, -1, 0, 1)` works and the "must be last in expression if nested" limitation is currently considered a feature.

`seq([i,i\string ^2], i=1..10)` crashes with Ooops, looks like we are missing a `]`. Aborting! The cause is that the square brackets do not hide the comma from `seq()` parsing. This will probably remain ``wont-fix''. Work-arounds: either use an extra pair of parentheses `seq(([i,i^2]), ...)` or hide the inner comma within braces `seq([i,{i^2}, ...)`.

The list stops here, but there are certainly other pending bugs in my bug-log, and many more I am not yet aware of. In particular it is already mentioned in the `\xintdeffunc` documentation that it can not parse currently the entirety of the available purely numerical syntax, some (documented or not, known or not) limitations apply.

3. The macros of **xintexpr** (ancient documentation, mostly)

.1	The <code>\xintexpr</code> expressions	61	.13	The <code>\xintthespaceseparated</code> macro	70
.2	<code>\numexpr</code> or <code>\dimexpr</code> expressions, count and dimension registers and variables	64	.14	<code>\xintifboolexpr</code> , <code>\xintifboolfloatexpr</code> , <code>\xintifbooliiexpr</code>	71
.3	Catcodes and spaces	64	.15	<code>\xintifsgnexpr</code> , <code>\xintifsgnfloatexpr</code> , <code>\xintifsgniiexpr</code>	71
.4	Expandability, <code>\xintexpro</code>	65	.16	The <code>\xintNewExpr</code> , <code>\xintNewIIExpr</code> , <code>\xint-</code> <code>NewFloatExpr</code> , <code>\xintNewIExpr</code> , and <code>\xint-</code> <code>NewBoolExpr</code> macros	71
.5	<code>\xintDigits*</code> , <code>\xintSetDigits*</code>	66	.17	Analogies and differences of <code>\xintiexpr</code> with <code>\numexpr</code>	72
.6	<code>\xintiexpr</code> , <code>\xinttheiexpr</code>	66	.18	Chaining expressions for expandable algo- rithmics	73
.7	<code>\xintiexpr</code> , <code>\xinttheiexpr</code>	66	.19	When expandability is too much.....	75
.8	<code>\xintboolexpr</code> , <code>\xinttheboolexpr</code>	67	.20	Acknowledgements (2013/05/25)	77
.9	<code>\xintfloatexpr</code> , <code>\xintthefloatexpr</code>	67			
.10	<code>\xinteval</code> , <code>\xintieval</code> , <code>\xintiieval</code> , <code>\xintfloateval</code>	68			
.11	Using an expression parser within another one.....	69			
.12	The <code>\xintthecoords</code> macro	69			

The **xintexpr** package was first released with version 1.07 (2013/05/25) of the **xint** bundle. It was substantially enhanced with release 1.1 from 2014/10/28.

The 1.4 release from 2020/01/31 maintains the same general architecture but needed adapting all the code base for the switch from `\csname` to `\expanded` techniques. On this occasion the mechanism for defining functions was substantially strengthened. The parser core mechanisms were improved too.

The package loads automatically **xintfrac** and **xinttools**.

This section should be trimmed to contain only information not already covered in [section 2](#).

3.1. The `\xintexpr` expressions

x ★ An **xintexpression** is a construct `\xintexpr<expandable_expression>\relax` where the expandable expression is read and completely expanded from left to right.

An `\xintexpr... \relax` must end in a `\relax` (which will be absorbed). Contrarily to a `\numexpr` expression, it is printable as is without a prefix `\the` or `\number` (don't use them with `\xintexpr` this will raise an error).

But one can use `\xintthe` prefix if one does need the explicit digits and other characters as in the final typesetted result.

As an alternative and equivalent syntax to

```
\xintexpr round(<expression>, D)\relax
```

there is

```
\xintiexpr [D] <expression> \relax
```

For $D > 0$ this produces a decimal number with D figures after the decimal mark, which is the rounding of the expression. For $D = 0$ the rounding to an integer is produced. For $D < 0$ (and this was changed at 1.4f), the rounded quotient of the expression by $1e|D|$ is produced.

- the expression may contain arbitrarily many levels of nested parenthesized sub-expressions,
- the expression may contain explicitly or from a macro expansion a sub-expression `\xintexpr. . \relax`, which itself may contain a sub-expressions etc. . .
- to let sub-contents evaluate as a sub-unit it should thus be either
 1. parenthesized,
 2. or a sub-expression `\xintexpr... \relax`.

- to use an expression as argument to macros from `xintfrac`, or more generally to macros which expand their arguments, one must use the `\xinttheexpr...\relax` or `\xintthe\xintexpr...\relax` forms.
- one should not use `\xintthe\xintexpr...\relax` as a sub-constituent of another expression but only the `\xintexpr...\relax` form which is more efficient in this context.
- each `xintexpr` expression, whether prefixed or not with `\xintthe`, is completely expandable and obtains its result in two expansion steps.

The information now following is possibly in need of updates.

- An expression is built the standard way with opening and closing parentheses, infix operators, and (big) numbers, with possibly a fractional part, and/or scientific notation (except for `\xintiexpr` which only admits big integers). All variants work with comma separated expressions. On output each comma will be followed by a space. A decimal number must have digits either before or after the decimal mark.
- As everything gets expanded, the characters `.`, `+`, `-`, `*`, `/`, `^`, `!`, `&`, `|`, `?`, `:`, `<`, `>`, `=`, `(`, `)`, `"`, `]`, `[`, `@` and the comma `,` should not (if used in the expression) be active. For example, the French language in `Babel` system, for pdf_{TeX}, activates `!`, `?`, `;` and `:`. Turn off the activity before expressions using such characters.

Alternatively the macro `\xintexprSafeCatcodes` resets all characters potentially needed by `\xintexpr` to their standard catcodes and `\xintexprRestoreCatcodes` restores the former status.

- Count registers and `\numexpr`-essions are accepted (LaTeX's counters can be inserted using `\value`) natively without `\the` or `\number` as prefix. Also dimen registers and control sequences, skip registers and control sequences (TeX's lengths), `\dimexpr`-essions, `\glueexpr`-essions are automatically unpacked using `\number`, discarding the stretch and shrink components and giving the dimension value in `sp` units (1/65536th of a TeX point). Furthermore, tacit multiplication is implied, when the (count or dimen or glue) register or variable, or the (`\numexpr` or `\dimexpr` or `\glueexpr`) expression is immediately prefixed by a (decimal) number. See [subsection 2.8](#) for the complete rules of tacit multiplication.



- With a macro `\x` defined like this:

```
\def\x {\xintexpr \a + \b \relax} or \edef\x {\xintexpr \a+\b\relax}
```

one may then do `\xintthe\x`, either for printing the result on the page or to use it in some other macros expanding their arguments. The `\edef` does the computation immediately but keeps it in a protected form. Naturally, the `\edef` is only possible if `\a` and `\b` are already defined. With both approaches the `\x` can be inserted in other expressions, as for example (assuming naturally as we use an `\edef` that in the 'yet-to-be computed' case the `\a` and `\b` now have some suitable meaning):

```
\edef\y {\xintexpr \x^3\relax}
```

- There is also `\xintboolexpr ... \relax` and `\xinttheboolexpr ... \relax`.
- See also `\xintifboolexpr` ([subsection 3.14](#)) and the `bool()` and `togl()` functions in [section 2](#). Here is an example. Well in fact the example ended up using only `\xintboolexpr` so it was modified to use `\xintifboolexpr`.

```
\xintdeffunc A(p,q,r) = p && (q || r) ;
\xintdeffunc B(p,q,r) = p || (q && r) ;
\xintdeffunc C(p,q,r) = xor(p, q, r) ;

\centeredline{\normalcolor
\begin{tabular}{ccrcl}
\xintFor* #1 in {{False}{True}} \do {%
\xintFor* #2 in {{False}{True}} \do {%
```

```

\intFor* #3 in {{False}{True}} \do {%
#1 &AND &(#2 &OR #3)&is&\textcolor[named]{OrangeRed}
      {\xintifboolexpr{A(#1,#2,#3)}{true}{false}}\
#1 &OR #2 &AND #3)&is&\textcolor[named]{OrangeRed}
      {\xintifboolexpr{B(#1,#2,#3)}{yes}{no}}\
#1 &XOR #2 &XOR #3 &is&\textcolor[named]{OrangeRed}
      {\xintifboolexpr{C(#1,#2,#3)}{oui}{non}}\
}}
\end{tabular}%
}

```

False	AND	(False	OR	False)	is	false
False	OR	(False	AND	False)	is	no
False	XOR	False	XOR	False	is	non
False	AND	(False	OR	True)	is	false
False	OR	(False	AND	True)	is	no
False	XOR	False	XOR	True	is	oui
False	AND	(True	OR	False)	is	false
False	OR	(True	AND	False)	is	no
False	XOR	True	XOR	False	is	oui
False	AND	(True	OR	True)	is	false
False	OR	(True	AND	True)	is	yes
False	XOR	True	XOR	True	is	non
True	AND	(False	OR	False)	is	false
True	OR	(False	AND	False)	is	yes
True	XOR	False	XOR	False	is	oui
True	AND	(False	OR	True)	is	true
True	OR	(False	AND	True)	is	yes
True	XOR	False	XOR	True	is	non
True	AND	(True	OR	False)	is	true
True	OR	(True	AND	False)	is	yes
True	XOR	True	XOR	False	is	non
True	AND	(True	OR	True)	is	true
True	OR	(True	AND	True)	is	yes
True	XOR	True	XOR	True	is	oui

- See also [\xintifsgnexpr](#).
- There is [\xintfloatexpr](#) ... [\relax](#) where the algebra is done in floating point approximation (also for each intermediate result). Use the syntax [\xintDigits:=N\relax](#) to set the precision. Default: 16 digits.

```
\xintthefloatexpr 2^100000\relax: 9.990020930143845e30102
```

The square-root operation can be used in [\xintexpr](#), it is computed as a float with the precision set by [\xintDigits](#) or by the optional second argument:

```

\xinttheexpr sqrt(2,60)\relax\newline
Here the [60] is to avoid truncation to |\xinttheDigits| of precision on output.
\newline
\printnumber{\xintthefloatexpr [60] sqrt(2,60)\relax}

```

```
1.41421356237309504880168872420969807856967187537694807317668
```

Here the [60] is to avoid truncation to [\xinttheDigits](#) of precision on output.

```
1.41421356237309504880168872420969807856967187537694807317668
```

Floats are quickly indispensable when using the power function, as exact results will easily have hundreds, even thousands of digits.

```
\xintDigits:=48\relax \xintthefloatexpr 2^100000\relax
```

9.99002093014384507944032764330033590980429139054e30102

Only integer and (in `\xintfloatexpr...\relax`) half-integer exponents are allowed.

- if one uses *macros* within `\xintexpr...\relax` one should obviously take into account that the parser will *not* see the macro arguments, hence one cannot use the syntax there, except if the arguments are themselves wrapped as `\xinttheexpr...\relax` and assuming the macro *f-expands* these arguments.

3.2. `\numexpr` or `\dimexpr` expressions, count and dimension registers and variables

Count registers, count control sequences, dimen registers, dimen control sequences (like `\parindent`), skips and skip control sequences, `\numexpr`, `\dimexpr`, `\glueexpr`, `\fontdimen` can be inserted directly, they will be unpacked using `\number` which gives the internal value in terms of scaled points for the dimensional variables: `1pt = 65536sp` (stretch and shrink components are thus discarded).

Tacit multiplication (see [subsection 2.8](#)) is implied, when a number or decimal number prefixes such a register or control sequence. \TeX lengths are skip control sequences and \TeX counters should be inserted using `\value`.

Release 1.2 of the `\xintexpr` parser also recognizes and prefixes with `\number` the `\ht`, `\dp`, and `\wd` \TeX primitives as well as the `\fontcharht`, `\fontcharwd`, `\fontchardp` and `\fontcharic` ε - \TeX primitives.

In the case of numbered registers like `\count255` or `\dimen0` (or `\ht0`), the resulting digits will be re-parsed, so for example `\count255 0` is like `100` if `\the\count255` would give `10`. The same happens with inputs such as `\fontdimen6\font`. And `\numexpr 35+52\relax` will be exactly as if `87` as been encountered by the parser, thus more digits may follow: `\numexpr 35+52\relax 000` is like `87000`. If a new `\numexpr` follows, it is treated as what would happen when `\xintexpr` scans a number and finds a non-digit: it does a tacit multiplication.

`\xinttheexpr \numexpr 351+877\relax\numexpr 1000-125\relax\relax{}` is the same as `\xinttheexpr 1228*875\relax`.

`1074500` is the same as `1074500`.

Control sequences however (such as `\parindent`) are picked up as a whole by `\xintexpr`, and the numbers they define cannot be extended extra digits, a syntax error is raised if the parser finds digits rather than a legal operation after such a control sequence.

A token list variable must be prefixed by `\the`, it will not be unpacked automatically (the parser will actually try `\number`, and thus fail). Do not use `\the` but only `\number` with a dimen or skip, as the `\xintexpr` parser doesn't understand `pt` and its presence is a syntax error. To use a dimension expressed in terms of points or other \TeX recognized units, incorporate it in `\dimexpr...\relax`.

Regarding how dimensional expressions are converted by \TeX into scaled points see also [subsection 8.7](#).

3.3. Catcodes and spaces

The main problems are caused by active characters, because `\xintexpr` et al. expand forward whatever comes from token stream; they apply `\string` only in a second step. For example the catcode of `&` from `&&` Boolean disjunction is not really important as long as it is not active, or comment, or escape... or brace... or ignored... in brief, as long as it is reasonable, and in particular whether `@` is of catcode letter or other does not matter.

It is always possible to insert manually the `\string` in the expression before a problematic (but reasonable) character catcode, or even to use `\detokenize` for a big chunk.

3.3.1. `\xintexprSafeCatcodes`

Some problems with active characters can be resolved on the fly by prefixing them by `\string` but some aspects of the parsing done by `\xintexpr` involves delimited macros which need the comma, equality sign and closing parenthesis to have their standard catcodes.

So `\xintexprSafeCatcodes` is provided as a utility to set in one go catcodes of many characters to `\xintexpr`-safely compatible values. This is a non-expandable step as it changes catcodes.

`\xintdefvar`, `\xintdeffunc`, et al., use it, and then they restore catcodes to the prior state via `\xintexprRestoreCatcodes`.

3.3.2. `\xintexprRestoreCatcodes`

Restores the catcodes to the state prevailing at the time of the last executed `\xintexprSafeCatcodes` (if located at the same \LaTeX environment or \TeX grouping level).

Prior to 1.4k, in a situation like the following:

```
\xintexprSafeCatcodes
...stuff possibly changing catcodes
\xintexprSafeCatcodes
...stuff possibly changing catcodes
\xintexprSafeCatcodes
...stuff possibly changing catcodes
\xintexprRestoreCatcodes
```

On exit, the catcodes recovered their status as prior to the *first* `\xintexprSafeCatcodes`. Since 1.4k, they are set to what they were prior to the *last* `\xintexprSafeCatcodes`, i.e. the mechanism is now similar to a ``last in, first out'' stack.

Note that no global assignments are made so the behaviour can be modified by usage of \TeX groups or \LaTeX environments: e.g. if an `\xintexprSafeCatcodes` is issued inside a \LaTeX environment it does not have to be paired by `\xintexprRestoreCatcodes` explicitly, the catcode scope is limited by the environment.

Spaces inside an `\xinttheexpr...\relax` should mostly be innocuous (except inside macro arguments).

`\xintexpr` and `\xinttheexpr` are for the most part agnostic regarding catcodes, but the characters in the expression should not be “active” (except on purpose) as everything is expanded along the way, and `\xintexpr` will choke on typesetting related commands. One can use `\string` to prefix a problematic character.

Digits, slash, square brackets, minus sign, in the output from an `\xinttheexpr` are all of catcode 12. For `\xintthefloatexpr` the `e' in the output has its standard catcode ``letter''.

3.4. Expandability, `\xintexpr`

As is the case with all other package macros `\xintexpr` *f-expands* (in two steps) to its final (somewhat protected) result; and `\xinttheexpr` *f-expands* (in two steps) to the chain of digits (and possibly minus sign -, decimal mark ., fraction slash /, scientific e, square brackets [,]) representing the result.

The once expanded `\xintexpr` is `\romannumeral0\xintexpr`. And there are similarly `\xintiexpr`, `\xintiexpr` and `\xintfloatexpr`. For an example see [subsection 3.18](#).

An expression can only be legally finished by a `\relax` token, which will be absorbed.

It is quite possible to nest expressions among themselves; for example, if one needs inside an `\xintiexpr...\relax` to do some computations with fractions, rounding the final result to an integer, one just has to insert `\xintiexpr...\relax`. The functioning of the infix operators will not be in the least affected from the fact that the outer ``environment'' is the `\xintiexpr` one.

3.5. `\xintDigits*`, `\xintSetDigits*`

These starred variants of `\xintDigits` and `\xintSetDigits` execute `\xintreloadxinttrig` and `\xintreloadxintlog`.

3.6. `\xintiexpr`, `\xinttheiexpr`

- x ★ Equivalent to doing `\xintexpr round(...)\relax` (more precisely, `round` is applied to each leaf item of the `ople` independently of its depth).

Intermediate calculations are exact, only the final output gets rounded. Half integers are rounded towards $+\infty$ for positive numbers and towards $-\infty$ for negative ones.

An optional parameter `D` within brackets, immediately after `\xintiexpr` is allowed: it instructs (for `D>0`) the expression to do its final rounding to the nearest value with that many digits after the decimal mark, i.e. `\xintiexpr [D] <expression>\relax` is equivalent (in case of a single expression) to `\xintexpr round(<expression>, D)\relax`.

`\xintiexpr [0] ...` is the same as `\xintiexpr ...` and rounds to an integer.

The case of negative `D` gives quantization to an integer multiple of $1e^{-D}$. This was modified at 1.4, and the produced value is now the rounded quotient by $1e^{-D}$ (i.e. no trailing zeros nor scientific exponent in the output).

If truncation rather than rounding is needed one can use `\xintexpr trunc(...)\relax` for truncation to an integer or `\xintexpr trunc(...,D)\relax` for quantization to an integer multiple or $1e^D$ (if `D>0`, for `D<0` the analog would be `trunc(...)/1e^{-D}`). But this works only for a single scalar value.

When defining a macro doing something such as `\xintiexpr #1\relax`, it is recommended to rather use `\xintiexpr\empty #1\relax`, as the `#1` may start with a `[` which without the `\empty` would be interpreted by `\xintiexpr` as the start of the optional `[D]`.

3.7. `\xintiexpr`, `\xinttheiexpr`

- x ★ This variant does not know fractions. It deals almost only with long integers. Comma separated lists of expressions are allowed.

It maps `/` to the *rounded* quotient. The operator `//` is, like in `\xintexpr...\relax`, mapped to *truncated* division. The Euclidean quotient (which for positive operands is like the truncated quotient) was, prior to release 1.1, associated to `/`. The function `quo(a,b)` can still be employed.

The `\xintiexpr`-essions use the ``ii'` macros for addition, subtraction, multiplication, power, square, sums, products, Euclidean quotient and remainder.

The `round`, `trunc`, `floor`, `ceil` functions are still available, and are about the only places where fractions can be used, but `/` within, if not somehow hidden will be executed as integer rounded division. To avoid this one can wrap the input in `qfrac`: this means however that none of the normal expression parsing will be executed on the argument.

To understand the illustrative examples, recall that `round` and `trunc` have a second (non negative) optional argument. In a normal `\xintexpr`-essions, `round` and `trunc` are mapped to `\xintRound` and `\xintTrunc`, in `\xintiexpr`-essions, they are mapped to `\xintiRound` and `\xintiTrunc`.

```
\xinttheiexpr 5/3, round(5/3,3), trunc(5/3,3), trunc(\xintDiv {5}{3},3),
trunc(\xintRaw {5/3},3)\relax} are problematic, but
%
\xinttheiexpr 5/3, round(qfrac(5/3),3), trunc(qfrac(5/3),3), floor(qfrac(5/3)),
```

```
ceil(qfrac(5/3))\relax{} work!
```

2, 2000, 2000, 2000, 2000 are problematic, but 2, 1667, 1666, 1, 2 work!

On the other hand decimal numbers and scientific numbers can be used directly as arguments to the `num`, `round`, or any function producing an integer.

Scientific numbers will be represented with as many zeroes as necessary, thus one does not want to insert `num(1e100000)` for example in an `\xintiiexpression`!

```
\xinttheiexpr num(13.4567e3)+num(10000123e-3)\relax
% should (num truncates) compute 13456+10000
```

23456

The `reduce` function is not available and will raise an error. The `frac` function also. The `sqr` function is mapped to `\xintiiSqrt` which gives a truncated square root. The `sqrtr` function is mapped to `\xintiiSqrtR` which gives a rounded square root.

One can use the Float macros if one is careful to use `num`, or `round` etc. . . on their output.

```
\xinttheiexpr \xintFloatSqrt [20]{2},
\xintFloatSqrt [20]{3}\relax % no operations
```

14142135623730950488[-19], 17320508075688772935[-19] The above went through because no actual operations were carried out. But it is dangerous because the `printer` for `\xinttheiexpr` could choke on such values. By default however it does nothing.

In the next example there will be an addition. So we first apply `round` to get integers (the second argument of `round` and `trunc` tells how many digits from after the decimal mark one should keep.)

```
\xinttheiexpr round(\xintFloatSqrt [20]{2},19) +
round(\xintFloatSqrt [20]{3},19)\relax
```

31462643699419723423

The whole point of `\xintiiexpr` is to gain some speed in *integer-only* algorithms, and the above explanations related to how to nevertheless use fractions therein are a bit peripheral. We observed (2013/12/18) of the order of 30% speed gain when dealing with numbers with circa one hundred digits (1.2: this info may be obsolete).

3.8. `\xintboolexpr`, `\xinttheboolexpr`

x ★ Equivalent to doing `\xintexpr ...\relax` and returning `true` if the result does not vanish, and `false` if the result is zero. As `\xintexpr`, this can be used on comma separated lists of expressions, and even bracketed lists.

It can be customized, one only needs to modify the following:

```
\def\xintboolexprPrintOne#1{\xintiiifNotZero{#1}{true}{false}}%
```

Not only are `true` and `false` usable in input, also `True` and `False` are pre-declared variables.

There is quirk in case it is used as a sub-expression: the boolean expression needs at least one logic operation else the value is not standardized to 1 or 0, for example we get from

```
\xinttheexpr \xintboolexpr 1.23\relax\relax\newline
```

1.23

which is to be compared with

```
\xinttheboolexpr 1.23\relax
```

true

3.9. `\xintfloatexpr`, `\xintthefloatexpr`

x ★ `\xintfloatexpr...\relax` is exactly like `\xintexpr...\relax` but with the four binary operations and the power function mapped to `\xintFloatAdd`, `\xintFloatSub`, `\xintFloatMul`, `\xintFloatDiv` and `\xintFloatPower`, respectively.²⁴

²⁴ Since 1.2f the `^` handles half-integer exponents, contrarily to `\xintFloatPower`.

The target precision for the computation is from the current setting of `\xintDigits`. Comma separated lists of expressions are allowed.

An optional parameter within brackets `[Q]` is allowed at the very start of the expression:

- if positive it instructs the macro to round the result to that many digits of precision. It thus makes sense to employ it only if this parameter is less than the `\xinttheDigits` precision.
- if negative it means to trim off that many digits (of course, in the sense of rounding the values to shorter mantissas). Don't use it to trim all digits (or more than all)!

Since 1.2f all float operations first round their arguments; a parsed number is not rounded prior to its use as operand to such a float operation.

```
\xintDigits:=36\relax
\xintthefloatexpr (1/13+1/121)*(1/179-1/173)/(1/19-1/18)\relax
0.00564487459334466559166166079096852897
\xintthefloatexpr\xintexpr (1/13+1/121)*(1/179-1/173)/(1/19-1/18)\relax\relax
0.00564487459334466559166166079096852912
```

The latter is the rounding of the exact result. The former one has its last three digits wrong due to the cumulative effect of rounding errors in the intermediate computations, as compared to exact evaluations.

I recall here from subsection 8.2 that with release 1.2f the float macros for addition, subtraction, multiplication and division round their arguments first to `P` significant places with `P` the asked-for precision of the output; and similarly the power macros and the square root macro. This does not modify anything for computations with arguments having at most `P` significant places already.

When defining a macro doing something such as `\xintfloatexpr #1\relax`, it is recommended to rather use `\xintfloatexpr\empty #1\relax`, as the `#1` may start with a `[` which without the `\empty` would be interpreted by `\xintfloatexpr` as the start of the optional `[Q]`.

3.10. `\xinteval`, `\xintieval`, `\xintiieval`, `\xintfloateval`

x ★ `\xinteval` is an *f*-expandable macro which is basically defined in such a way that `\xinteval{<expression>}` behaves like `\xinttheexpr{<expression>}\relax`. It expands completely in two steps and delivers its output using digits, the dot `.` as decimal separator, the letter `e` for scientific notation, the slash `/` for fractions, as well as commas in case of multi-items expression and square brackets `[` and `]` for nesting.

x ★ `\xintieval` is similarly related to `\xinttheiexpr`. It admits an optional argument `[D]` which may be located in the expected location from conventions of `\TeXe` macros with optional argument, but had been long constrained (until 1.4k) to be inside the braces at the start of the expression.

```
\xintieval[7]{355/113} = \xintieval{[7]355/113}
3.1415929 = 3.1415929
```

When defining a macro doing something such as `\xintieval{#1}`, it is recommended to rather use `\xintieval{\empty #1}`, as the `#1` may start with a `[` which without the `\empty` would be interpreted by `\xintieval` as the start of the optional `[D]`.

x ★ `\xintiieval` is similarly related to `\xinttheiexpr`.

x ★ `\xintfloateval` is similarly related to `\xintthefloatexpr`. It admits an optional argument `[Q]` which may be located either outside (since 1.4k) or inside the braces.

```
\xintfloateval [7]{355/113} = \xintfloateval{[7] 355/113}
```

3.141593 = 3.141593

When negative, the optional argument tells how many digits to remove from the prevailing precision:

```
\xintfloateval[-2]{355/113}=
\xintfloateval{[-2]355/113} has \xinttheDigits\ minus 2 digits.
```

3.1415929203540= 3.1415929203540 has 16 minus 2 digits.

When defining a macro doing something such as `\xintfloateval{#1}`, it is recommended to rather use `\xintfloateval{\empty #1}`, as the `#1` may start with a `[` which without the `\empty` would be interpreted by `\xintfloateval` as the start of the optional `[Q]`.

The `\xinteval` et al. macros have the advantage that they can be nested within things such as the `\num` macro of `siunitx`, which would choke on some of the syntax elements which are allowed inside `\xinttheexpr... \relax`, such as brackets. Hidden inside the braces as argument to `\xinteval` et al., they do not cause a breakage of `\num` anymore.

And regarding `\xintieval` and `\xintfloateval` one sees here the advantage that their respective optional arguments can (and had to until 1.4k) be located within the braces too...

3.11. Using an expression parser within another one

This was already illustrated before. In the following:

```
\xintfloatexpr \xintexpr add(1/i, i=1234..1243)\relax ^100\relax
```

5.136088460396579e-210, the inner sum is computed exactly. Then it will be rounded to `\xinttheDigits` significant digits, and then its power will be evaluated as a float operation. One should avoid the "`\xintthe`" parsers in inner positions as this induces digit by digit parsing of the inner computation result by the outer parser. Here is the same computation done with floats all the way:

```
\xintfloatexpr add(1/i, i=1234..1243)^100\relax
```

5.136088460396643e-210

Not surprisingly this differs from the previous one which was exact until raising to the 100th power.

The fact that the inner expression occurs inside a bigger one has nil influence on its behaviour. There is the limitation though that the outputs from `\xintexpr` and `\xintfloatexpr` can not be used directly in `\xinttheiexpr` integer-only parser. But one can do:

```
\xintiexpr round(\xintfloatexpr 3.14^10\relax)\relax % or trunc
```

93174

3.12. The `\xintthecoords` macro

It converts (in two expansion steps) the expansion result of `\xintfloatexpr` (or `\xintexpr` or `\xintiexpr`) into the `(a, b) (c, d) ...` format for list of coordinates as expected by the `TikZ coordinates` syntax.

```
\begin{figure}[htbp]
\centering\begin{tikzpicture}[scale=10]\xintDigits:=8\relax
\clip (-1.1,-.25) rectangle (.3,.25);
\draw [blue] (-1.1,0)--(1,0);
\draw [blue] (0,-1)--(0,+1);
\draw [red] plot[smooth] coordinates {%
%% \xintthecoords converts output of next expression into the
%% (x1, y1) (x2, y2) ...
%% format
\xintthecoords\xintfloatexpr
```

```

%% This syntax -1+[0..4]/2 is currently dropped at xint 1.4
%% seq((x^2-1,mul(x-t,t=-1+[0..4]/2)),x=-1.2..[0.1]..+1.2)\relax
%% Use this:
seq((x^2-1,mul(x-t,t=seq(-1+u/2, u=0..4))),x=-1.2..[0.1]..+1.2)
\relax
};
\end{tikzpicture}
\caption{Coordinates with \csbxint{thecoords}.}
\end{figure}

```




Figure 2: Coordinates with `\xintthecoords`.

Unstable! It is currently undecided how `\xintthecoords` should handle bracketed data. Currently, it (or `TikZ`) will break if the input contains nested structures. One can use it with `flat()` which removes all nesting. And in combination with `zip()` it is easy to plot data given by some mechanism in separate lists of x- and y-coordinates (see an example in next section)

3.13. The `\xintthespaceseparated` macro

It converts (in two expansion steps) the expansion result of `\xintfloatexpr` (or `\xintexpr` or `\xintiexpr`) into the space separated format suitable for usage with `PS-Tricks \listplot` macro.

Here is for example some syntax (the replacement text of `\foo`, which is used here only to show that indeed complete expansion is attained in two steps) which can be used as argument to `\listplot`. Using 4 fractional decimal digits is sufficient when unit is the centimeter (it gives a fixed point precision of one micron, amply enough for plots...).

```

\oodef\foo{%
\xintthespaceseparated
\xintiexpr[4]\xintfloatexpr seq((i, log10(i)), i=1..[0.5]..10)\relax\relax
}\meaning\foo

```

```

macro:->1.0000 0 1.5000 0.1761 2.0000 0.3010 2.5000 0.3979 3.0000 0.4771 3.5000 0.5441 4.0000
0.6021 4.5000 0.6532 5.0000 0.6990 5.5000 0.7404 6.0000 0.7782 6.5000 0.8129 7.0000 0.8451 7.5000
0.8751 8.0000 0.9031 8.5000 0.9294 9.0000 0.9542 9.5000 0.9777 10.0000 1.0000

```

Here we don't really need the inner `\xintfloatexpr... \relax` because the `log10()` function works the same in the exact parser `\xintexpr` but in general this is recommended.

Unstable! It is currently undecided how `\xintthespaceseparated` should handle bracketed data. Currently, it (or `\listplot`) will break if the input contains nested structures. One can use it with `flat()` which removes all nesting. And in combination with `zip()` it is easy to plot data given by some mechanism in separate lists of x- and y-coordinates.

```

% let's imagine we have something like this
\def\Xcoordinates{1, 3, 5, 7, 9}

```

```
\def\Ycoordinates{1, 9, 25, 49, 81}
% then:
|\xintthespaceseparated\xintexpr flat(zip([\Xcoordinates], [\Ycoordinates]))\relax
is suitable to use as argument to |\listplot|, as it expands to
\xintthespaceseparated\xintexpr flat(zip([\Xcoordinates], [\Ycoordinates]))\relax
\xintthespaceseparated\xintexpr flat(zip([\Xcoordinates], [\Ycoordinates]))\relax is suitable
to use as argument to \listplot, as it expands to 1 1 3 9 5 25 7 49 9 81
```

3.14. `\xintifboolexpr`, `\xintifboolfloatexpr`, `\xintifbooliiexpr`

xnn ★ `\xintifboolexpr{<expr>}{<YES>}{<NO>}` does `\xinttheexpr<expr>\relax` and then executes the `<YES>` or the `<NO>` branch depending on whether the outcome was non-zero or zero. Thus one can read *if bool expr* as meaning *if not zero*:

if `<expr>`-ession does not vanish do `<YES>` else do `<NO>`

The expression is not limited to using only comparison operators and Boolean logic (`<`, `>`, `==`, `!=`, `&&`, `||`, `all()`, `any()`, `xor()`, `bool()`, `toql()`, ...), it can be the most general computation.

xnn ★ `\xintifboolfloatexpr{<expr>}{<YES>}{<NO>}` does `\xintthefloatexpr<expr>\relax` and then executes the `<YES>` or the `<NO>` branch depending on whether the outcome was non zero or zero.

xnn ★ `\xintifbooliiexpr{<expr>}{<YES>}{<NO>}` does `\xinttheiiexpr<expr>\relax` and then executes the `<YES>` or the `<NO>` branch depending on whether the outcome was non zero or zero.

The expression argument must be a single one, comma separated sub-expressions will cause low-level errors.

3.15. `\xintifsgnexpr`, `\xintifsgnfloatexpr`, `\xintifsgniiexpr`

xnnn ★ `\xintifsgnexpr{<expr>}{<<0>}{<=0>}{<>0>}` evaluates the `\xintexpression` and chooses the branch corresponding to its sign.

xnnn ★ `\xintifsgnfloatexpr{<expr>}{<<0>}{<=0>}{<>0>}` evaluates the `\xintfloatexpression` and chooses the branch corresponding to its sign.

xnnn ★ `\xintifsgniiexpr{<expr>}{<<0>}{<=0>}{<>0>}` evaluates the `\xintiiexpression` and chooses the branch corresponding to its sign.

The expression argument must be a single one, comma separated sub-expressions will cause low-level errors.

3.16. The `\xintNewExpr`, `\xintNewIIExpr`, `\xintNewFloatExpr`, `\xintNewIExpr`, and `\xintNewBoolExpr` macros

`\xintNewExpr` macro is used as:

```
\xintNewExpr{\myformula}[n]{<stuff>}, where
```

- `<stuff>` will be inserted inside `\xinttheexpr . . . \relax`,
- `n` is an integer between zero and nine, inclusive, which is the number of parameters of `\myformula`,
- the placeholders `#1`, `#2`, ..., `#n` are used inside `<stuff>` in their usual rôle,^{25 26}
- the `[n]` is mandatory, even for `n=0`.²⁷
- the macro `\myformula` is defined without checking if it already exists, \TeX users might prefer to do first `\newcommand*\myformula {}` to get a reasonable error message in case `\myformula` already exists,

²⁵ if `\xintNewExpr` is used inside a macro, the `#`'s must be doubled as usual. ²⁶ the `#`'s will in practice have their usual catcode, but category code other `#`'s are accepted too. ²⁷ there is some use for `\xintNewExpr[0]` compared to an `\edef` as `\xintNewExpr` has some built-in catcode protection.

- the protection against active characters is done automatically (as long as the whole thing has not already been fetched as a macro argument and the catcodes correspondingly already frozen).

It (if it succeeds) will be a completely expandable macro entirely built-up using `\xintAdd`, `\xintSub`, `\xintMul`, `\xintDiv`, `\xintPow`, etc. . . as corresponds to the expression written with the infix operators. Macros created by `\xintNewExpr` can thus be nested.

```
\xintNewFloatExpr \FA [2]{(#1+#2)^10}
\xintNewFloatExpr \FB [2]{sqrt(#1*#2)}
\begin{enumerate}[nosep]
  \item \FA {5}{5}
  \item \FB {30}{10}
  \item \FA {\FB {30}{10}}{\FB {40}{20}}
\end{enumerate}
```

1. 1e10
2. 17.32050807568877
3. 3.891379490446502e16

The documentation is much shortened here because `\xintNewExpr` and `\xintdeffunc` are very much related one with the other.

ATTENTION!

The original spirit of `\xintNewExpr` was to define a (possibly very big) macro using only `xintfrac`, and this means in particular that it must be used only with arguments compatible with the `xintfrac` input format.

Thus an `\xintexpr` declared variable has no chance to work, it must be wrapped explicitly in `\xinteval{...}` to be fetched as argument to a macro constructed by `\xintNewExpr`.

They share essentially the same limitations.

Notice though that `\xintNewFloatExpr` accepts and recognizes the optional argument `[Q]` of `\xintfloatexpr`, contrarily to `\xintdeffloatfunc`. Use an `\empty` in case the contents are not known in advance.

Historical note: prior to 1.4, `xintexpr` used a `\csname..\endcsname` encapsulation technique which impacted the string pool memory. The `\xintNewExpr` was designed as a method to pre-parse the expression and produce one single, gigantic, nested usage of the relevant `xintfrac` macros. This way, only those macros were expanded which had nil impact on the \TeX string pool.

Later on it was found that this mechanism could be employed to define functions. Basically underneath 98% of `\xintNewExpr` and `\xintdeffunc` are using the same shared code.

3.17. Analogies and differences of `\xintiexpr` with `\numexpr`

`\xintiexpr..\relax` is a parser of expressions knowing only (big) integers. There are, besides the enlarged range of allowable inputs, some important differences of syntax between `\numexpr` and `\xintiexpr` and variants:

- Contrarily to `\numexpr`, the `\xintiexpr` parser will stop expanding only after having encountered (and swallowed) a mandatory `\relax` token.
- In particular, spaces between digits (and not only around infix operators or parentheses) do not stop `\xintiexpr`, contrarily to the situation with `numexpr`: `\the\numexpr 7 + 3 5\relax` expands (in one step)²⁸ to `105\relax`, whereas `\xintthe\xintiexpr 7 + 3 5\relax` expands (in two steps) to `42`.²⁹

²⁸ The `\numexpr` triggers continued expansion after the space following the 3 to check if some operator like + is upstream. But after having found the 5 it treats it as an end-marker. ²⁹ Since 1.21 one can also use the underscore `_` to separate digits for readability of long numbers.

- Inside an `\edef`, an expression `\xintiexpr...\relax` get fully evaluated, whereas `\numexpr` without `\the` or `\number` prefix would not, if not itself embedded in another `\the\numexpr` or similar context.
- (ctd.) The private format to which `\xintiexpr...\relax` (et al.) evaluates may use `\xintthe` prefix to turn into explicit digits, (for example in arguments to some macros which expand their arguments). The `\the` TeX primitive prefix would not work here.
- (ctd.) One can embed a `\numexpr...\relax` (with its `\relax!`) inside an `\xintiexpr...\relax` without `\the` or `\number`, but the reverse situation requires usage of `\xintthe` or `\xinteval` user interface,
- `\numexpr -(1)\relax` is illegal. In contrast `\xintiexpr -(1)\relax` is perfectly legal and gives the expected result (what else ?).
- `\numexpr 2\cnta\relax` is illegal (with `\cnta` a `\count` register.) In contrast `\xintiexpr 2\cnta\relax` is perfectly legal and will do the tacit multiplication.
- `\the\numexpr` or `\number\numexpr` expands in one step, but `\xintthe\xintiexpr` or `\xinttheiexpr` needs two steps.

3.18. Chaining expressions for expandable algorithmics

We will see in this section how to chain `\xintexpr`-essions with `\expandafter`'s, like it is possible with `\numexpr`. For this it is convenient to use `\romannumeral0\xintexpr` which is the once-expanded form of `\xintexpr`, as we can then chain using only one `\expandafter` each time.

For example, here is the code employed on the title page to compute (expandably, of course!) the 1250th Fibonacci number:

```
\catcode`_ 11
\def\Fibonacci #1{% \Fibonacci{N} computes F(N) with F(0)=0, F(1)=1.
  \expandafter\Fibonacci_a\expandafter
    {\the\numexpr #1\expandafter}\expandafter
    {\romannumeral0\xintiexpro 1\expandafter\relax\expandafter}\expandafter
    {\romannumeral0\xintiexpro 1\expandafter\relax\expandafter}\expandafter
    {\romannumeral0\xintiexpro 1\expandafter\relax\expandafter}\expandafter
    {\romannumeral0\xintiexpro 0\relax}}
%
\def\Fibonacci_a #1{%
  \ifcase #1
    \expandafter\Fibonacci_end_i
  \or
    \expandafter\Fibonacci_end_ii
  \else
    \ifodd #1
      \expandafter\expandafter\expandafter\Fibonacci_b_ii
    \else
      \expandafter\expandafter\expandafter\Fibonacci_b_i
    \fi
  \fi {#1}%
}% * signs are omitted from the next macros, tacit multiplications
\def\Fibonacci_b_i #1#2#3{\expandafter\Fibonacci_a\expandafter
  {\the\numexpr #1/2\expandafter}\expandafter
  {\romannumeral0\xintiexpro sqrt(#2)+sqrt(#3)\expandafter\relax
    \expandafter}\expandafter
  {\romannumeral0\xintiexpro (2#2-#3)#3\relax}%
```

```

}% end of Fibonacci_b_i
\def\Fibonacci_b_ii #1#2#3#4#5{\expandafter\Fibonacci_a\expandafter
  {\the\numexpr (#1-1)/2\expandafter}\expandafter
  {\romannumeral0\xintiipro sqrt(#2)+sqrt(#3)\expandafter\relax
    \expandafter}\expandafter
  {\romannumeral0\xintiipro (2#2-#3)#3\expandafter\relax\expandafter}\expandafter
  {\romannumeral0\xintiipro #2#4+#3#5\expandafter\relax\expandafter}\expandafter
  {\romannumeral0\xintiipro #2#5+#3(#4-#5)\relax}%
}% end of Fibonacci_b_ii
%
% code as used on title page:
%\def\Fibonacci_end_i #1#2#3#4#5{\xintthe#5}
%\def\Fibonacci_end_ii #1#2#3#4#5{\xinttheiipro #2#5+#3(#4-#5)\relax}
%
% new definitions:
\def\Fibonacci_end_i #1#2#3#4#5{{#4}{#5}}% {F(N+1)}{F(N)} in \xintexpr format
\def\Fibonacci_end_ii #1#2#3#4#5%
  {\expandafter
  {\romannumeral0\xintiipro #2#4+#3#5\expandafter\relax
    \expandafter}\expandafter
  {\romannumeral0\xintiipro #2#5+#3(#4-#5)\relax}}% idem.
% \FibonacciN returns F(N) (in encapsulated format: needs \xintthe for printing)
\def\FibonacciN {\expandafter\xint_secondoftwo\romannumeral-`0\Fibonacci}%
\catcode`_ 8

```

The macro `\Fibonacci` produces not one specific value $F(N)$ but a pair of successive values $\{F(N), F(N+1)\}$ which can then serve as starting point of another routine devoted to compute a whole sequence $F(N), F(N+1), F(N+2), \dots$. Each of $F(N)$ and $F(N+1)$ is kept in the encapsulated internal `xintexpr` format.

`\FibonacciN` produces the single $F(N)$. It also keeps it in the private format; thus printing it will need the `\xintthe` prefix.

Here a code snippet which checks the routine via a `\message` of the first 51 Fibonacci numbers (this is not an efficient way to generate a sequence of such numbers, it is only for validating `\FibonacciN`).

```

\def\Fibo #1.{\xintthe\FibonacciN {#1}}%
\message{\xintloop [0+1] \expandafter\Fibo\xintloopindex.,
  \ifnum\xintloopindex<49 \repeat \xintthe\FibonacciN{50}.}

```

The way we use `\expandafter`'s to chain successive `\xintiipro` evaluations is exactly analogous to what is possible with `\numexpr`. The various `\romannumeral0\xintiipro` could very well all have been `\xintiipro`'s but then we would have needed `\expandafter\expandafter\expandafter` each time.

There is a difference though: `\numexpr` does *NOT* expand inside an `\edef`, and to force its expansion we must prefix it with `\the` or `\number` or `\romannumeral` or another `\numexpr` which is itself prefixed, etc. . . .

But `\xintexpr`, `\xintiipro`, . . . , expand fully in an `\edef`, with the completely expanded result encapsulated in a private format.

Using `\xintthe` as prefix is necessary to print the result (like `\the` or `\number` in the case of `\numexpr`), but it is not necessary to get the computation done (contrarily to the situation with `\numexpr`).

Our `\Fibonacci` expands completely under *f-expansion*, so we can use `\fdef` rather than `\edef` in a situation such as

```
\fdef \X {\FibonacciN {100}} ,
```

but it is usually about as efficient to employ `\edef`. And if we want

```
\edef \Y {(\FibonacciN{100},\FibonacciN{200})} ,
```

then `\edef` is necessary.

Allright, so let's now give the code to generate $\{F(N)\{F(N+1)\{F(N+2)\}\dots\}$, using `\Fibonacci` for the first two and then using the standard recursion $F(N+2)=F(N+1)+F(N)$:

```
\catcode`_ 11
\def\FibonacciSeq #1#2{%#1=starting index, #2>#1=ending index
  \expandafter\Fibonacci_Seq\expandafter
  {\the\numexpr #1\expandafter}\expandafter{\the\numexpr #2-1}%
}%
\def\Fibonacci_Seq #1#2{%
  \expandafter\Fibonacci_Seq_loop\expandafter
  {\the\numexpr #1\expandafter}\romannumeral0\Fibonacci {#1}{#2}%
}%
\def\Fibonacci_Seq_loop #1#2#3#4{% standard Fibonacci recursion
  {#3}\unless\ifnum #1<#4 \Fibonacci_Seq_end\fi
  \expandafter\Fibonacci_Seq_loop\expandafter
  {\the\numexpr #1+1\expandafter}\expandafter
  {\romannumeral0\xintiipro #2+#3\relax}{#2}{#4}%
}%
\def\Fibonacci_Seq_end\fi\expandafter\Fibonacci_Seq_loop\expandafter
  #1\expandafter #2#3#4{\fi {#3}}%
\catcode`_ 8
```

This `\FibonacciSeq` macro is completely expandable but it is not *f-expandable*.

This is not a problem in the next example which uses `\xintFor*` as the latter applies repeatedly full expansion to what comes next each time it fetches an item from its list argument. Thus `\xintFor*` still manages to generate the list via iterated full expansion.

```
\newcounter{myindex}% not "index", which would overwrite theindex environment!
% (many have probably been bitten by this trap)
\tabskip 1ex
\ldef\Fibxxx{\FibonacciN {30}}%
\setcounter{myindex}{30}%
\vbox{\halign{\bfseries#.\hfil&#\hfil &\hfil #\cr
  \xintFor* #1 in {\FibonacciSeq {30}{59}}\do
  {\themyindex &\xintthe#1 &
  \xintiiRem{\xintthe#1}{\xintthe\Fibxxx}\stepcounter{myindex}\cr }}%
}\vrule
\vbox{\halign{\bfseries#.\hfil&#\hfil &\hfil #\cr
  \xintFor* #1 in {\FibonacciSeq {60}{89}}\do
  {\themyindex &\xintthe#1 &
  \xintiiRem{\xintthe#1}{\xintthe\Fibxxx}\stepcounter{myindex}\cr }}%
}\vrule
\vbox{\halign{\bfseries#.\hfil&#\hfil &\hfil #\cr
  \xintFor* #1 in {\FibonacciSeq {90}{119}}\do
  {\themyindex &\xintthe#1 &
  \xintiiRem{\xintthe#1}{\xintthe\Fibxxx}\stepcounter{myindex}\cr }}%
}%
```

This produces the Fibonacci numbers from $F(30)$ to $F(119)$, and computes also all the congruence classes modulo $F(30)$. The output has been put in a `float`, which appears on the following page. I leave to the mathematically inclined readers the task to explain the visible patterns. . . ;-).

3.19. When expandability is too much

Let's use the macros of [subsection 3.18](#) related to Fibonacci numbers. Notice that the 47th Fibonacci number is `2971215073` thus already too big for \TeX and $\varepsilon\text{-}\TeX$.

30.	832040	0	60.	1548008755920	0	90.	2880067194370816120	0
31.	1346269	514229	61.	2504730781961	1	91.	4660046610375530309	514229
32.	2178309	514229	62.	4052739537881	1	92.	7540113804746346429	514229
33.	3524578	196418	63.	6557470319842	2	93.	12200160415121876738	196418
34.	5702887	710647	64.	10610209857723	3	94.	19740274219868223167	710647
35.	9227465	75025	65.	17167680177565	5	95.	31940434634990099905	75025
36.	14930352	785672	66.	27777890035288	8	96.	51680708854858323072	785672
37.	24157817	28657	67.	44945570212853	13	97.	83621143489848422977	28657
38.	39088169	814329	68.	72723460248141	21	98.	135301852344706746049	814329
39.	63245986	10946	69.	117669030460994	34	99.	218922995834555169026	10946
40.	102334155	825275	70.	190392490709135	55	100.	354224848179261915075	825275
41.	165580141	4181	71.	308061521170129	89	101.	573147844013817084101	4181
42.	267914296	829456	72.	498454011879264	144	102.	927372692193078999176	829456
43.	433494437	1597	73.	806515533049393	233	103.	1500520536206896083277	1597
44.	701408733	831053	74.	1304969544928657	377	104.	2427893228399975082453	831053
45.	1134903170	610	75.	2111485077978050	610	105.	3928413764606871165730	610
46.	1836311903	831663	76.	3416454622906707	987	106.	6356306993006846248183	831663
47.	2971215073	233	77.	5527939700884757	1597	107.	10284720757613717413913	233
48.	4807526976	831896	78.	8944394323791464	2584	108.	16641027750620563662096	831896
49.	7778742049	89	79.	14472334024676221	4181	109.	26925748508234281076009	89
50.	12586269025	831985	80.	23416728348467685	6765	110.	43566776258854844738105	831985
51.	20365011074	34	81.	37889062373143906	10946	111.	70492524767089125814114	34
52.	32951280099	832019	82.	61305790721611591	17711	112.	114059301025943970552219	832019
53.	53316291173	13	83.	99194853094755497	28657	113.	184551825793033096366333	13
54.	86267571272	832032	84.	160500643816367088	46368	114.	298611126818977066918552	832032
55.	139583862445	5	85.	259695496911122585	75025	115.	483162952612010163284885	5
56.	225851433717	832037	86.	420196140727489673	121393	116.	781774079430987230203437	832037
57.	365435296162	2	87.	679891637638612258	196418	117.	1264937032042997393488322	2
58.	591286729879	832039	88.	1100087778366101931	317811	118.	2046711111473984623691759	832039
59.	956722026041	1	89.	1779979416004714189	514229	119.	3311648143516982017180081	1

Some Fibonacci numbers together with their residues modulo $F(30)=832040$

The `\FibonacciN` macro found in [subsection 3.18](#) is completely expandable, it is even *f*-expandable. We need a wrapper with `\xintthe` prefix

```
\def\theFibonacciN{\xintthe\FibonacciN}
```

to print in the document or to use within `\message` (or `\TeX typeout`) to write to the log and terminal.

The `\xintthe` prefix also allows its use it as argument to the `xint` macros: for example if we are interested in knowing how many digits $F(1250)$ has, it suffices to issue `\xintLen {\theFibonacciN {1250}}` (which expands to `261`). Or if we want to check the formula $\gcd(F(1859), F(1573)) = F(\gcd(1859, 1573)) = F(143)$, we only need³⁰

```

\xintiiGCD{\theFibonacciN{1859}}{\theFibonacciN{1573}}=%
\theFibonacciN{\xintiiGCD{1859}{1573}}$

```

which produces:

```
343358302784187294870275058337 = 343358302784187294870275058337
```

The `\theFibonacciN` macro expanded its `\xintiiGCD{1859}{1573}` argument via the services of `\nu`, `mexpr`: this step allows only things obeying the \TeX bound, naturally! (but $F(2147483648)$ would be rather big anyhow...).

This is very convenient but of course it repeats the complete evaluation each time it is done. In practice, it is often useful to store the result of such evaluations in macros. Any `\edef` will break expandability, but if the goal is at some point to print something to the `dvi` or `pdf` output, and not only to the `log` file, then expandability has to be broken one day or another!

³⁰ The `\xintiiGCD` macro is provided by both the `xintgcd` package (since 1.0) and by the `xint` package (since 1.3d).

Hence, in practice, if we want to print in the document some computation results, we can proceed like this and avoid having to repeat identical evaluations:

```
\begingroup
\def\A {1859} \def\B {1573}
\edef\X {\theFibonacciN\A} \edef\Y {\theFibonacciN\B}
\edef\GCDAB {\xintiiGCD\A\B}\edef\Z {\theFibonacciN\GCDAB}
\edef\GCDXY{\xintiiGCD\X\Y}
The identity  $\gcd(F(\A),F(\B))=F(\gcd(\A,\B))$  can be checked via evaluation
of both sides:  $\gcd(F(\A),F(\B))=\gcd(\printnumber\X,\printnumber\Y)=$ 
 $\printnumber{\GCDXY} = F(\gcd(\A,\B)) = F(\GCDAB) =\printnumber\Z$.\par
% some further computations involving \A, \B, \X, \Y
\endgroup % closing the group removes assignments to \A, \B, ...
% or choose longer names less susceptible to overwrite something.
% Note: there is no LaTeX \newcommand which would be to \edef like
% \newcommand is to \def$ 
```

The identity $\gcd(F(1859),F(1573)) = F(\gcd(1859, 1573))$ can be checked via evaluation of both sides: $\gcd(F(1859), F(1573)) = \gcd(14405827913044251198771689151504042869913161495023481014226$
 $68636701088272597575494722482437753529619459794869227357628882216309358018264080851775319974$
 $25695605529435028861585245173725088673642222849290822895245583889495442192655760412999290255$
 $65979711337876105452217623490841529979811413199660087517689703410997520079993610707576019520$
 $876324584695551467505894985013610208598628752325727241, 244384192519511857332827945977626199$
 $85399024815706192326053609007840133940367432124452232789599095158695811031891779769058032741$
 $51632595307616686661013725200866754096569888951010022888016831459347310131566517721593249344$
 $79863439947937119575876654476582795890928239007031319713554812200493864453132952484774727316$
 $6471511289078393) = 343358302784187294870275058337 = F(\gcd(1859, 1573)) = F(143) = 3433583027841$
 $87294870275058337.$

One may legitimately ask the author: why expandability to such extremes, for things such as big fractions or floating point numbers (even continued fractions...) which anyhow can not be used directly within \TeX 's primitives such as `\ifnum`? Why insist on a concept which is foreign to the vast majority of \TeX users and even programmers?

I have no answer: it made definitely sense at the start of `xint` (see [subsection 8.13](#)) and once started I could not stop.

3.20. Acknowledgements (2013/05/25)

I was greatly helped in my preparatory thinking, prior to producing such an expandable parser, by the commented source of the `l3fp` package, specifically the `l3fp-parse.dtx` file (in the version of April-May 2013; I think there was in particular a text called ``roadmap'' which was helpful). Also the source of the `calc` package was instructive, despite the fact that here for `\xintexpr` the principles are necessarily different due to the aim of achieving expandability.

4. The **xinttrig** package

<p>.1 \xintreloadxinttrig..... 78</p> <p>.2 Constants..... 78</p> <p>.3 Functions..... 79</p>		<p>.4 Important implementation notes..... 80</p> <p>.5 Some example evaluations 80</p>
---	--	--

This package provides trigonometric functions for use with **xintexpr**. The sole macro is `\xintreloadxinttrig`.

This package was first included in release 1.3e (2019/04/05) of **xintexpr**. It is automatically loaded by **xintexpr**.

At 1.4e (2021/05/05) the accuracy was significantly increased: formerly the high-level user interface used to define the functions had as consequences that intermediate steps of the computations could not operate with guard digits, and as a result the last two digits were most of the time off (at least the last one). Now, computations are done internally in extended precision, and the accuracy is high up to the last digits, with faithful rounding and high probability of correct rounding. And the maximal number of digits was raised slightly to 62 digits.

At 8 digits a special, faster, mode is used, which is less accurate. But faster.

Acknowledgements: I finally decided to release some such functions under friendly pressure of Jürgen GILG and Thomas SÖLL, let them both be thanked here.

Jürgen passed away in 2022. I will miss our friendship which was born and grew from numerous and regular exchanges on topics not limited to this package or even the T_EX world. Let's now continue to "take care and keep motivated"!

4.1. \xintreloadxinttrig

The library is loaded automatically by **xintexpr** at start-up. It is then configured for 16 digits.

To work for example with 48 digits, execute `\xintSetDigits*{48}` or `\xintDigits*:=48`; (the ending `;` can be replaced by a `\relax` in case of problems due to it being active, e.g. with ~~T_EX~~ and some languages).

{ With the non-starred variant `\xintDigits:=48`; it is needed to issue `\xintreloadxinttrig` to recalibrate the functions provided by the library (and the exponential/logarithm functions will only be updated if also `\xintreloadxintlog` is used).

4.2. Constants

Their values (with more digits) get incorporated into the trigonometrical functions at the time of their definitions during loading or reloading of the package. They are left free to use, or modified, or `\xintunassignvar`'d, as this will have no impact whatsoever on the functions.

twoPi what could that be?

threePiover2

Pi

Piover2

oneRadian this is one radian in degrees: $180/\pi$

oneDegree this is one degree in radian: $\pi/180$

4.3. Functions

4.3.1. Direct trigonometry

With the variable in radians:

sin(x) sine

cos(x) cosine

tan(x) tangent

cot(x) cotangent

sec(x) secant

csc(x) cosecant

With the variable in degrees:

sind(x) sine

cosd(x) cosine

tand(x) tangent

cotd(x) cotangent

secd(x) secant

cscd(x) cosecant

Only available with the variable in radians:

tg(x) tangent

cotg(x) cotangent

sinc(x) cardinal sine $\text{sinc}(x) = \sin(x)/x$

4.3.2. Inverse trigonometry

With the value in radians:

asin(x) arcsine

acos(x) arccosine

atan(x) arctangent

Arg(x, y) the main branch of the argument of the complex number $x+iy$, from $-\pi$ (excluded) to π (included). As the output is rounded $-\text{Pi}$ is a possible return value.

pArg(x, y) the branch of the argument of the complex number $x+iy$ with values going from 0 (included) to 2π (excluded). Inherent rounding makes twoPi a possible return value.

atan2(y, x) it is **Arg(x, y)**. Note the reversal of the arguments, this seems to be the most frequently encountered convention across languages.

With the value in degrees:

asind(x) arcsine

acosd(x) arccosine

atand(x) arctangent

Argd(x, y) the main branch of the argument of the complex number $x+iy$, from -180 (excluded) to 180 (included). Inherent rounding of output can cause -180 to be returned.

pArgd(x, y) the branch of the argument of the complex number $x+iy$ with values going from 0 (included) to 360 (excluded). Inherent rounding of output can cause 360 to be returned.

atan2d(y, x) it is **Argd(x, y)**. Note the reversal of the arguments, this seems to be the most frequently encountered convention across languages.

4.3.3. Conversion functions (optional definitions left to user decision)

Python provides functions **degrees()** and **radians()**. But as most of the **xinttrig** functions are already defined for the two units, I felt this was not really needed. It is a oneliner to add them:

```
\xintdeffloatfunc radians(x) := x * oneDegree;
\xintdeffloatfunc degrees(x) := x * oneRadian;
\xintdeffunc radians(x) := float_dgt(x * oneDegree);
\xintdeffunc degrees(x) := float_dgt(x * oneRadian);
```

The **float_dgt()** does a float rounding to **\xinttheDigits** precision (recall that ***** is mapped to exact multiplication in **\xintdeffunc**).

4.4. Important implementation notes

- Currently, **xint** is lacking some dedicated internal representation of floats which means that most operations re-parse the digit tokens of their arguments to count them. . . this does not contribute to efficiency (you can load the module under **\xintverbosetrue** regime and see how the nested macros look like and get an idea of how many times some rather silly re-counting of mantissa lengths will get done!)
- One should not overwrite some function names which are employed as auxiliaries; refer to [sourcexint.pdf](#).
- Floats with large exponents are integers and are multiple of **1000**; hence modulo **360** all such 'angles' are multiple of **40** degrees. Needless to say that considering usage of the **sind()** and **cosd()** functions with such large float numbers is meaningless.
- See [sourcexint.pdf](#) for some comments on limitations of the range reduction implementation.

4.5. Some example evaluations

```
\xintDigits* := 48\relax
Digits at \xinttheDigits:\newline
$sind(17)\approx\xintfloateval{sind(17)}$\newline
$cosd(17)\approx\xintfloateval{cosd(17)}$\newline
$tand(17)\approx\xintfloateval{tand(17)}$\newline
$sind(43)\approx\xintfloateval{sind(43)}$\newline
$cosd(43)\approx\xintfloateval{cosd(43)}$\newline
$tand(43)\approx\xintfloateval{tand(43)}$\newline
$asind(0.3)\approx\xintfloateval{asind(0.3)}$\newline
$acosd(0.3)\approx\xintfloateval{acosd(0.3)}$\newline
$atand(3)\approx\xintfloateval{atand(3)}$\newline
```



```

$tan(atan(7))\approx\xintfloateval{tan(atan(7))}$\newline
$asind(sind(25))\approx\xintfloateval{asind(sind(25))}$\par\medskip
\noindent\xintDigits* := 24\relax
Digits at \xinttheDigits:\newline
$sind(17)\approx\xintfloateval{sind(17)}$\newline
$cosd(17)\approx\xintfloateval{cosd(17)}$\newline
$tand(17)\approx\xintfloateval{tand(17)}$\newline
$sind(43)\approx\xintfloateval{sind(43)}$\newline
$cosd(43)\approx\xintfloateval{cosd(43)}$\newline
$tand(43)\approx\xintfloateval{tand(43)}$\newline
$asind(0.3)\approx\xintfloateval{asind(0.3)}$\newline
$acosd(0.3)\approx\xintfloateval{acosd(0.3)}$\newline
$atand(3)\approx\xintfloateval{atand(3)}$\newline
$tan(atan(7))\approx\xintfloateval{tan(atan(7))}$\newline
$asind(sind(25))\approx\xintfloateval{asind(sind(25))}$\par
\xintDigits* := 16\relax

```

Digits at 48:

```

sind(17) ≈ 0.292371704722736728097468695377143252664687186183
cosd(17) ≈ 0.956304755963035481338650816618418962009410343991
tand(17) ≈ 0.305730681458660355734541958996550716146250221387
sind(43) ≈ 0.681998360062498500442225784711125580340433802762
cosd(43) ≈ 0.731353701619170483287543608275622403378396544763
tand(43) ≈ 0.932515086137661705612185627426186654353537299494
asind(0.3) ≈ 17.4576031237220922902460457924449418216636440147
acosd(0.3) ≈ 72.5423968762779077097539542075550581783363559853
atand(3) ≈ 71.5650511770779893515721937204532946712042142996
tan(atan(7)) ≈ 7.00000000000000000000000000000000000000000000000005
asind(sind(25)) ≈ 25

```

Digits at 24:

```

sind(17) ≈ 0.292371704722736728097469
cosd(17) ≈ 0.956304755963035481338651
tand(17) ≈ 0.305730681458660355734542
sind(43) ≈ 0.681998360062498500442226
cosd(43) ≈ 0.731353701619170483287544
tand(43) ≈ 0.932515086137661705612186
asind(0.3) ≈ 17.4576031237220922902460
acosd(0.3) ≈ 72.5423968762779077097540
atand(3) ≈ 71.5650511770779893515722
tan(atan(7)) ≈ 7
asind(sind(25)) ≈ 25

```

5. The **xintlog** package

This package provides logarithms, exponentials and fractional powers for use with **xintexpr**.

This package was first included in release 1.3e (2019/04/05) of **xintexpr**. It is automatically loaded by **xintexpr**.

At release 1.4e (2021/05/05) it was substantially extended to cover usage with mantissas of up to 62 digits.

At **Digits** set to 8 or less, the old faster but less accurate macros based on **poormanlog** are used. These macros compute logarithms and exponentials with about 8 or 9 nine digits of *fixed point* precision.


TeX-hackers note: There is thus, for **Digits=8** or less a systematic loss of rounding precision in the floating point sense for logarithms of inputs close to 1: e.g. `log10(1.0011871)` is produced as `5.15245e-4` which stands for `0.000515145` having indeed 9 correct fractional digits, but only 6 correct digits in the floating point sense. Situation is worse for `log()` as it applies a conversion factor and does not remove the trailing junk digits, which we don't have for `log10()`. Check [sourcexint.pdf](#) and **poormanlog** README for more info.

.1	<code>\xintreloadxintlog</code>	82		.3	Some information on how powers are com-	
.2	Functions	82			puted.....	82

5.1. `\xintreloadxintlog`

The library is loaded automatically by **xintexpr** at start-up. It is then configured for 16 digits.

To work for example with 48 digits, execute `\xintSetDigits*{48}` or `\xintDigits*:=48`; (the ending `;` can be replaced by a `\relax` in case of problems due to it being active, e.g. with \TeX and some languages).

 With the non-starred variant `\xintDigits:=48`; it is needed to issue `\xintreloadxintlog` to recalibrate the functions provided by the library (and the trigonometric functions will only be updated if also `\xintreloadxinttrig` is used).

5.2. Functions

log10(x) logarithm in base 10

pow10(x) fractional powers of 10

log(x) natural logarithm

exp(x) exponential function

pow(x, y) computes x^y either via the formula `pow10(y*log10(x))` (applied with some internally increased accuracy), for **y** neither an integer nor a half-integer; or via the legacy `\xintFloatPower` and `\xintFloatSqrt` macros if the exponent is integer or half-integer. Integer exponents trigger an exact evaluation in `\xinteval` if the output will not exceed (or will only slightly exceed) 10000 digits (separately for numerator and denominator), else the power is computed in the floating point sense.

```
\xintfloateval{log(2), exp(1), 2^(1/3), 2^10000}
0.6931471805599453, 2.718281828459045, 1.259921049894873, 1.995063116880758e3010
```

5.3. Some information on how powers are computed

For powers `a^b` or `a**b` in `\xintfloateval` the following rules apply:

1. a check is made if exponent is integer or half-integer,

- if this is the case legacy `\xintFloatPower` (combined with `\xintFloatSqrt` for half-integer case) are used to evaluate the power (and `a` can be negative if exponent is integer),
- else the power is computed as `pow10(b*log10(a))` (but keeping some extra digits in intermediate evaluations; in particular `b` is not float-rounded, but `a` is).

The reason is that the log/exp approach loses accuracy for very big exponents (say for exponents of the order of 100000000 or more). Here is an example of a precise computation with a very large exponent (184884258895036416):

```

\xintTeXFromSci{\xintfloateval{1.00000001^\xintiexpr 12^16\relax}}$\newline
\xintDigits:=48;%\xintreloadxintlog is not done as log10/pow10 will not be used
\xintTeXFromSci{\xintfloateval{1.00000001^12^16}}$\newline
\xintDigits:=64;%\xintreloadxintlog is not done as log10/pow10 will not be used
\xintTeXFromSci{\xintfloateval{1.00000001^12^16}}$\newline
\xintDigits:=80;%\xintreloadxintlog is not done as log10/pow10 will not be used
\xintTeXFromSci{\xintfloateval{1.00000001^12^16}}$
\xintDigits:=16;%
```

```

1.879985676694948 · 10802942130
1.87998567669494838838184407480229599674641360997 · 10802942130
1.879985676694948388381844074802295996746413609968646474887080800 · 10802942130
1.8799856766949483883818440748022959967464136099686464748870808001110266973999979 · 10802942130
```

- Notes:
- in the case with 16 digits precision, we ensured 12^{16} got computed exactly with all its 18 digits and was not rounded to only 16 digits (and confirmation is that the result matches the second one at 48 digits),
 - the 1.4g right associativity of powers is taken into account to drop parentheses.

As the legacy `\xintFloatPower` and `\xintFloatSqrt` work in arbitrary precision, the result for integer or half-integer exponents is produced with a full-size mantissa, even if `Digits` is more than 62 (as is exemplified above).

In the $10^{(b \cdot \log_{10}(a))}$ branch the mantissa size is limited to the minimum of `Digits` and of 64. Its last digits will start being wrong if `b` becomes about (in absolute value) 100000000. If you really need to compute powers with exponents that large or larger, it is recommended to decompose the exponent as a sum of the nearest integer or half-integer and a fractional part and express the power as a product. This is not done automatically as it would add some overhead in general for some a priori very rare use cases.

In `\xinteval`, this is as in `\xintfloateval` but for one difference: integer exponents will trigger an exact evaluation, as long as:

- the exponent absolute value is at most 9999,
- it is evaluated a priori, based on the length of the input, that the output will have at most 10000 digits (or only a bit more), separately for numerator and denominator.

The check for integrality of exponent is not on its mathematical value but on its internal representation, for speed. So $6/3$ is not recognized as being an integer exponent in `\xinteval`; but in `\xintfloateval`, the $6/3$ will have been computed and recognized as 2. Also 2.00 or $200e-2$ is recognized as an integer in both parsers. Similar remarks apply to half-integer case.

To compute exactly higher powers than 2^{9999} or 9^{9999} or 99^{5000} or 999^{3333} , etc..., use `\xint-iiieval`. See `\xintiPow` for related comments if you don't want to melt your CPU.

If `Digits` is at most 8, logarithms are computed faster but with less accuracy; internally, using 9 fixed point fractional digits. And powers a^b lose accuracy in last digits quickly as `b` rises. Here is what was observed with some random tests:

- for `b` neither integer nor half-integer and $1 < b < 10$, roughly 8 correct digits for between 80% and 90% of cases and in the remaining cases only a 1ulp error.

[TOC](#), [Start here](#), [xintexpr](#), [xintexpr \(old doc\)](#), [xinttrig](#), [xintlog](#), [xinttools](#), [Examples](#), [xint bundle](#)

- for b neither integer nor half-integer and $10^e < b < 10^{e+1}$, roughly $8-e$ digits are correct for about 90% of cases and there is a one unit error in the last of those digits in the remaining cases.

To maintain higher accuracy, split the input as $a^n a^h$ with n integer or half-integer nearest to b . After having considered (and implemented) the method, decision was made to not incorporate it as it would induce serious overhead generally speaking. The a^b with fractional exponent b such that $abs(b) < 10$ are currently computed with at most 1ulp error in the vast majority of cases it seems, which is largely precise enough for plots, and then speed matters most. Larger exponents can be handled (since 1.4f) via manually implementing the splitting trick, as described above.

The documentation of the legacy macro `\xintFloatPower` (which is used for powers with integer and half-integer exponents) explains it has a guaranteed error bound of 0.52ulp, in arbitrary precision. Generally speaking, the math functions added at 1.4e target even smaller errors (but only up to 62 digits), something of the order of 0.505ulp, and in practice they seem to achieve even better than 99% of correct rounding probability (at least in their natural ranges, and it varies according to the value of `Digits`). Perhaps in future I will re-examine whether it is worthwhile to increase a bit the theoretical accuracy of `\xintFloatPower`, as I have not had the time to really measure systematically its practical accuracy, all anecdotal evidence showing it is good.

6. Macros of the **xinttools** package

These utilities used to be provided within the **xint** package; since 1.09g (2013/11/22) they have been moved to an independently usable package **xinttools**, which has none of the **xint** facilities regarding big numbers. Whenever relevant release 1.09h has made the macros `\long` so they accept `\par` tokens on input.

The completely expandable utilities (up to `\xintilooop`) are documented first, then the non expandable utilities.

[section 7](#) gives additional (some quite dated) examples of use of macros of this package.

xinttools is automatically loaded by **xintexpr**.

.1	<code>\xintRevWithBraces</code>	85	.15	<code>\xintilooop</code> , <code>\xintilooopindex</code> , <code>\xintouter-</code>	
.2	<code>\xintZapFirstSpaces</code> , <code>\xintZapLas-</code>			<code>ilooopindex</code> , <code>\xintbreakilooop</code> , <code>\xint-</code>	
	<code>tSpaces</code> , <code>\xintZapSpaces</code> , <code>\xintZapSpacesB</code>	85		<code>breakilooopanddo</code> , <code>\xintilooopskiptonext</code> ,	
.3	<code>\xintCSVtoList</code>	86		<code>\xintilooopskipandredo</code>	94
.4	<code>\xintNthElt</code>	87	.16	<code>\xintApplyInline</code>	97
.5	<code>\xintNthOnePy</code>	88	.17	<code>\xintFor</code> , <code>\xintFor*</code>	98
.6	<code>\xintKeep</code>	88	.18	<code>\xintifForFirst</code> , <code>\xintifForLast</code>	100
.7	<code>\xintKeepUnbraced</code>	89	.19	<code>\xintBreakFor</code> , <code>\xintBreakForAndDo</code>	101
.8	<code>\xintTrim</code>	89	.20	<code>\xintintegers</code> , <code>\xintdimensions</code> , <code>\xin-</code>	
.9	<code>\xintTrimUnbraced</code>	89		<code>trationals</code>	101
.10	<code>\xintListWithSep</code>	90	.21	<code>\xintForpair</code> , <code>\xintForthree</code> , <code>\xintFor-</code>	
.11	<code>\xintApply</code>	90		<code>four</code>	102
.12	<code>\xintApplyUnbraced</code>	91	.22	<code>\xintAssign</code>	103
.13	<code>\xintSeq</code>	91	.23	<code>\xintAssignArray</code>	104
.14	<code>\xintloop</code> , <code>\xintbreakloop</code> , <code>\xintbreak-</code>		.24	<code>\xintDigitsOf</code>	104
	<code>loopanddo</code> , <code>\xintloopskiptonext</code>	91	.25	<code>\xintRelaxArray</code>	105

6.1. `\xintRevWithBraces`

f ★ `\xintRevWithBraces{⟨list⟩}` first does the *f*-expansion of its argument then it reverses the order of the tokens, or braced material, it encounters, maintaining existing braces and adding a brace pair around each naked token encountered. Space tokens (in-between top level braces or naked tokens) are gobbled. This macro is mainly thought out for use on a `⟨list⟩` of such braced material; with such a list as argument the *f*-expansion will only hit against the first opening brace, hence do nothing, and the braced stuff may thus be macros one does not want to expand.

```
\edef\x{\xintRevWithBraces{12345}}
\meaning\x:macro:->{5}{4}{3}{2}{1}
\edef\y{\xintRevWithBraces\x}
\meaning\y:macro:->{1}{2}{3}{4}{5}
```

The examples above could be defined with `\edef`'s because the braced material did not contain macros. Alternatively:

```
\expandafter\def\expandafter\w\expandafter
{\romannumeral0\xintrevwithbraces{{\A}{\B}{\C}{\D}{\E}}}
\meaning\w:macro:->{\E }{\D }{\C }{\B }{\A }
```

n ★ The macro `\xintReverseWithBracesNoExpand` does the same job without the initial expansion of its argument.

6.2. `\xintZapFirstSpaces`, `\xintZapLastSpaces`, `\xintZapSpaces`, `\xintZapSpacesB`

n ★ `\xintZapFirstSpaces{⟨stuff⟩}` does not do any expansion of its argument, nor brace removal of any sort, nor does it alter `⟨stuff⟩` in anyway apart from stripping away all *leading* spaces.

This macro will be mostly of interest to programmers who will know what I will now be talking about. *The essential points, naturally, are the complete expandability and the fact that no brace removal nor any other alteration is done to the input.*

TeX's input scanner already converts consecutive blanks into single space tokens, but `\xintZapFirstSpaces` handles successfully also inputs with consecutive multiple space tokens. However, it is assumed that `<stuff>` does not contain (except inside braced sub-material) space tokens of character code distinct from 32.

It expands in two steps, and if the goal is to apply it to the expansion text of `\x` to define `\y`, then one can do: `\odef\y{\romannumeral0\expandafter\xintzapfirstspaces\expandafter{\x}}` (one can also define a wrapper macro to `\xintZapFirstSpaces` in order to expand once the argument first, but `xinttools` not being a programming layer, it provides no "Generate Variants" facilities).

Other use case: inside a macro which received a parameter `#1`, one can do `\oodef\x{\xintZapFirstSpaces {#1}}`, or, if `#1`, after leading spaces have been stripped can accept `\edef` expansion, one can do `\edef\x{\xintZapFirstSpaces{#1}}`.

```
\xintZapFirstSpaces { \a { \X } { \b \Y } }->\a { \X } { \b \Y } +++
```

- n ★** `\xintZapLastSpaces{<stuff>}` does not do any expansion of its argument, nor brace removal of any sort, nor does it alter `<stuff>` in anyway apart from stripping away all *ending* spaces. The same remarks as for `\xintZapFirstSpaces` apply.

```
\xintZapLastSpaces { \a { \X } { \b \Y } }-> \a { \X } { \b \Y } +++
```

- n ★** `\xintZapSpaces{<stuff>}` does not do any expansion of its argument, nor brace removal of any sort, nor does it alter `<stuff>` in anyway apart from stripping away all *leading* and all *ending* spaces. The same remarks as for `\xintZapFirstSpaces` apply.

```
\xintZapSpaces { \a { \X } { \b \Y } }->\a { \X } { \b \Y } +++
```

- n ★** `\xintZapSpacesB{<stuff>}` does not do any expansion of its argument, nor does it alter `<stuff>` in anyway apart from stripping away all leading and all ending spaces and possibly removing one level of braces if `<stuff>` had the shape `<spaces>{braced}<spaces>`. The same remarks as for `\xintZapFirstSpaces` apply.

```
\xintZapSpacesB { \a { \X } { \b \Y } }->\a { \X } { \b \Y } +++
\xintZapSpacesB { { \a { \X } { \b \Y } } }-> \a { \X } { \b \Y } +++
```

The spaces here at the start and end of the output come from the braced material, and are not removed (one would need a second application for that; recall though that the `xint` zapping macros do not expand their argument).

6.3. `\xintCSVtoList`

- f ★** `\xintCSVtoList{a,b,c...,z}` returns `{a}{b}{c}...{z}`. A *list* is by convention in this manual simply a succession of tokens, where each braced thing will count as one item ('items' are defined according to the rules of TeX for fetching undelimited parameters of a macro, which are exactly the same rules as for \TeX and macro arguments [they are the same things]). The word 'list' in 'comma separated list of items' has its usual linguistic meaning, and then an 'item' is what is delimited by commas.

So `\xintCSVtoList` takes on input a 'comma separated list of items' and converts it into a 'TeX list of braced items'. The argument to `\xintCSVtoList` may be a macro: it will first be *f-expanded*. Hence the item before the first comma, if it is itself a macro, will be expanded which may or may not be a good thing. A space inserted at the start of the first item serves to stop that expansion (and disappears). The macro `\xintCSVtoListNoExpand` does the same job without the initial expansion of the list argument.

- n ★** Apart from that no expansion of the items is done and the list items may thus be completely arbitrary (and even contain perilous stuff such as unmatched `\if` and `\fi` tokens).

Contiguous spaces and tab characters, are collapsed by TeX into single spaces. All such spaces

around commas³¹ are removed, as well as the spaces at the start and the spaces at the end of the list.³² The items may contain explicit `\par`'s or empty lines (converted by the \TeX input parsing into `\par` tokens).

```
\xintCSVtoList { 1 , { 2 , 3 , 4 , 5 } , a , {b,T} U , { c , d } , { {x , y} } }
->{1}{2 , 3 , 4 , 5}{a}{b,T} U { c , d } { {x , y} }
```

One sees on this example how braces protect commas from sub-lists to be perceived as delimiters of the top list. Braces around an entire item are removed, even when surrounded by spaces before and/or after. Braces for sub-parts of an item are not removed.

We observe also that there is a slight difference regarding the brace stripping of an item: if the braces were not surrounded by spaces, also the initial and final (but no other) spaces of the enclosed material are removed. This is the only situation where spaces protected by braces are nevertheless removed.

From the rules above: for an empty argument (only spaces, no braces, no comma) the output is `{ }` (a list with one empty item), for `<opt. spaces>{ }<opt. spaces>`' the output is `{ }` (again a list with one empty item, the braces were removed), for `{ }`' the output is `{ }` (again a list with one empty item, the braces were removed and then the inner space was removed), for `{ }`' the output is `{ }` (again a list with one empty item, the initial space served only to stop the expansion, so this was like `{ }`' as input, the braces were removed and the inner space was stripped), for `{ }`' the output is `{ }` (this time the ending space of the first item meant that after brace removal the inner spaces were kept; recall though that \TeX collapses on input consecutive blanks into one space token), for `{ }`' the output consists of two consecutive empty items `{ }{ }`. Recall that on output everything is braced, a `{ }` is an 'empty' item. Most of the above is mainly irrelevant for every day use, apart perhaps from the fact to be noted that an empty input does not give an empty output but a one-empty-item list (it is as if an ending comma was always added at the end of the input).

```
\def\y{ \a,\b,\c,\d,\e} \xintCSVtoList\y->{\a }{\b }{\c }{\d }{\e }
\def\t {{\if},\ifnum,\ifx,\ifdim,\ifcat,\ifmmode}
\xintCSVtoList\t->{\if }{\ifnum }{\ifx }{\ifdim }{\ifcat }{\ifmmode }
```

The results above were automatically displayed using \TeX 's primitive `\meaning`, which adds a space after each control sequence name. These spaces are not in the actual braced items of the produced lists. The first items `\a` and `\if` were either preceded by a space or braced to prevent expansion. The macro `\xintCSVtoListNoExpand` would have done the same job without the initial expansion of the list argument, hence no need for such protection but if `\y` is defined as `\def\y{\a, \b,\c,\d,\e}` we then must do:

```
\expandafter\xintCSVtoListNoExpand\expandafter {\y}
```

Else, we may have direct use:

```
\xintCSVtoListNoExpand {\if,\ifnum,\ifx,\ifdim,\ifcat,\ifmmode}
->{\if }{\ifnum }{\ifx }{\ifdim }{\ifcat }{\ifmmode }
```

Again these spaces are an artefact from the use in the source of the document of `\meaning` (or rather here, `\detokenize`) to display the result of using `\xintCSVtoListNoExpand` (which is done for real in this document source).

For the similar conversion from comma separated list to braced items list, but without removal of spaces around the commas, there is `\xintCSVtoListNonStripped` and `\xintCSVtoListNonStrippedNoExpand`.

6.4. `\xintNthElt`

$\overset{\text{num}}{x}$ `\xintNthElt{x}{<list>}` gets (expandably) the `x`th item of the `<list>`. A braced item will lose one level of brace pairs. The token list is first *f-expanded*.

³¹ and multiple space tokens are not a problem; but those at the top level (not hidden inside braces) *must* be of character code 32. ³² let us recall that this is all done completely expandably... There is absolutely no alteration of any sort of the item apart from the stripping of initial and final space tokens (of character code 32) and brace removal if and only if the item apart from initial and final spaces (or more generally multiple `char 32` space tokens) is braced.

Items are counted starting at one.

```
\xintNthElt {3}{\agh}\u{zzz}\v{Z}} is zzz
\xintNthElt {3}{\agh}\u{{zzz}}\v{Z}} is {zzz}
\xintNthElt {2}{\agh}\u{{zzz}}\v{Z}} is \u
\xintNthElt {37}{\xintiiFac {100}}=9 is the thirty-seventh digit of 100!.
\xintNthElt {10}{\xintFtoCv {566827/208524}}=1457/536
is the tenth convergent of 566827/208524 (uses xintcfrac package).
\xintNthElt {7}{\xintCSVtoList {1,2,3,4,5,6,7,8,9}}=7
\xintNthElt {0}{\xintCSVtoList {1,2,3,4,5,6,7,8,9}}=9
\xintNthElt {-3}{\xintCSVtoList {1,2,3,4,5,6,7,8,9}}=7
```

If $x=0$, the macro returns the *length* of the expanded list: this is not equivalent to `\xintLength` which does no pre-expansion. And it is different from `\xintLen` which is to be used only on integers or fractions.

If $x<0$, the macro returns the $|x|$ th element from the end of the list. Thus for example $x=-1$ will fetch the last item of the list.

```
\xintNthElt {-5}{\agh}\u{zzz}\v{Z}} is {agh}
```

num
x n ★

The macro `\xintNthEltNoExpand` does the same job but without first expanding the list argument:
`\xintNthEltNoExpand {-4}{\u\v\w T\x\y\z}` is T.

If x is strictly larger (in absolute value) than the length of the list then `\xintNthElt` produces empty contents.

6.5. \xintNthOnePy

num
x f ★

`\xintNthOnePy{x}{\langle list \rangle}` gets (expandably) the x th item of the $\langle list \rangle$, adding a brace pair if there wasn't one.

Attention, items are counted starting at zero. For negative index, behaves as `\xintNthElt`.

If the index is out of range, the empty output is returned. If the input list was empty (had no items) the empty output is returned.

6.6. \xintKeep

num
x f ★

`\xintKeep{x}{\langle list \rangle}` expands the token list argument L and produces a new list, depending on the value of x :

- if $x>0$, the new list contains the first x items from L (counting starts at one.) *Each such item will be output within a brace pair.* Use `\xintKeepUnbraced` if this is not desired. This means that if the list item was braced to start with, there is no modification, but if it was a token without braces, then it acquires them.
- if $x \geq \text{length}(L)$, the new list is the old one with all its items now braced.
- if $x=0$ the empty list is returned.
- if $x<0$ the last $|x|$ elements compose the output in the same order as in the initial list; as the macro proceeds by removing head items the kept items end up in output as they were in input: no added braces.
- if $x \leq -\text{length}(L)$ the output is identical with the input.

`\xintKeepNoExpand` does the same without first *f-expanding* its list argument.

```
\fdef\test {\xintKeep {17}{\xintKeep {-69}{\xintSeq {1}{100}}}\meaning\test\par
\noindent\fdef\test {\xintKeep {7}{\{1\}\{2\}\{3\}\{4\}\{5\}\{6\}\{7\}\{8\}\{9\}}\meaning\test\par
\noindent\fdef\test {\xintKeep {-7}{\{1\}\{2\}\{3\}\{4\}\{5\}\{6\}\{7\}\{8\}\{9\}}\meaning\test\par
\noindent\fdef\test {\xintKeep {7}{123456789}\meaning\test\par
\noindent\fdef\test {\xintKeep {-7}{123456789}\meaning\test\par
```

```
macro: ->{32}{33}{34}{35}{36}{37}{38}{39}{40}{41}{42}{43}{44}{45}{46}{47}{48}
```

```
macro: ->{1}{2}{3}{4}{5}{6}{7}
```

```
macro: ->{3}{4}{5}{6}{7}{8}{9}
```

```
macro: ->{1}{2}{3}{4}{5}{6}{7}
```


macro:->3456789

6.7. \xintKeepUnbraced

Same as `\xintKeep` but no brace pairs are added around the kept items from the head of the list in the case $x > 0$: each such item will lose one level of braces. Thus, to remove braces from all items of the list, one can use `\xintKeepUnbraced` with its first argument larger than the length of the list; the same is obtained from `\xintListWithSep`{\langle list \rangle}. But the new list will then have generally many more items than the original ones, corresponding to the unbraced original items.

For $x < 0$ the macro is no different from `\xintKeep`. Hence the name is a bit misleading because brace removal will happen only if $x > 0$.

`\xintKeepUnbracedNoExpand` does the same without first *f-expanding* its list argument.

```
\fdef\test {\xintKeepUnbraced {10}{\xintSeq {1}{100}}}\meaning\test\par
\noindent\fdef\test {\xintKeepUnbraced {7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}%
\meaning\test\par
\noindent\fdef\test {\xintKeepUnbraced {-7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}%
\meaning\test\par
\noindent\fdef\test {\xintKeepUnbraced {7}{123456789}}\meaning\test\par
\noindent\fdef\test {\xintKeepUnbraced {-7}{123456789}}\meaning\test\par
```

macro:->12345678910

macro:->1234567

macro:->{3}{4}{5}{6}{7}{8}{9}

macro:->1234567

macro:->3456789

6.8. \xintTrim

^{num}
 x *f* ★ `\xintTrim`{ x }{\langle list \rangle} expands the list argument and gobbles its first x elements.

- if $x > 0$, the first x items from L are gobbled. The remaining items are not modified.
- if $x \geq \text{length}(L)$, the returned list is empty.
- if $x = 0$ the original list is returned (with no added braces.)
- if $x < 0$ the last $|x|$ items of the list are removed. *The head items end up braced in the output.* Use `\xintTrimUnbraced` if this is not desired.
- if $x \leq -\text{length}(L)$ the output is empty.

`\xintTrimNoExpand` does the same without first *f-expanding* its list argument.

```
\fdef\test {\xintTrim {17}{\xintTrim {-69}{\xintSeq {1}{100}}}}\meaning\test\par
\noindent\fdef\test {\xintTrim {7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}\meaning\test\par
\noindent\fdef\test {\xintTrim {-7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}\meaning\test\par
\noindent\fdef\test {\xintTrim {7}{123456789}}\meaning\test\par
\noindent\fdef\test {\xintTrim {-7}{123456789}}\meaning\test\par
```

macro:->{18}{19}{20}{21}{22}{23}{24}{25}{26}{27}{28}{29}{30}{31}

macro:->{8}{9}

macro:->{1}{2}

macro:->89

macro:->{1}{2}

6.9. \xintTrimUnbraced

Same as `\xintTrim` but in case of a negative x (cutting items from the tail), the kept items from the head are not enclosed in brace pairs. They will lose one level of braces. The name is a bit misleading because when $x > 0$ there is no brace-stripping done on the kept items, because the macro works simply by gobbling the head ones.

`\xintTrimUnbracedNoExpand` does the same without first *f-expanding* its list argument.

```
\fdef\test {\xintTrimUnbraced {-90}{\xintSeq {1}{100}}}\meaning\test\par
\noindent\fdef\test {\xintTrimUnbraced {7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}%
\meaning\test\par
\noindent\fdef\test {\xintTrimUnbraced {-7}{{1}{2}{3}{4}{5}{6}{7}{8}{9}}}%
\meaning\test\par
\noindent\fdef\test {\xintTrimUnbraced {7}{123456789}}\meaning\test\par
\noindent\fdef\test {\xintTrimUnbraced {-7}{123456789}}\meaning\test\par
```

macro:->12345678910

macro:->{8}{9}

macro:->12

macro:->89

macro:->12

6.10. \xintListWithSep

nf ★ `\xintListWithSep{⟨sep⟩}{⟨list⟩}` inserts the separator `⟨sep⟩` in-between all items of the given list of braced items (or individual tokens). The items are fetched as does `TEX` with undelimited macro arguments, thus they end up unbraced in output. If the `⟨list⟩` is only one (or multiple) space tokens, the output is empty.

The list argument `⟨list⟩` gets *f-expanded* first (thus if it is a macro whose contents are braced items, the first opening brace stops the expansion, and it is as if the macro had been expanded once.) The separator `⟨sep⟩` is not pre-expanded, it ends up as is in the output (if the `⟨list⟩` contained at least two items.)

nn ★ The variant `\xintListWithSepNoExpand` does the same job without the initial expansion of the `⟨list⟩` argument.

```
\edef\foo{\xintListWithSep{, }{123456789{10}{11}{12}}}\meaning\foo\newline
\edef\foo{\xintListWithSep{:}{\xintiiFac{20}}}\meaning\foo\newline
\oodef\FOO{\xintListWithSepNoExpand{\FOO}{\bat\baz\biz\buz}}\meaning\FOO\newline
% a braced item or a space stops the f-expansion:
\oodef\foo{\xintListWithSep{\FOO}{{\bat}\baz\biz\buz}}\meaning\foo\newline
\oodef\foo{\xintListWithSep{\FOO}{ \bat\baz\biz\buz}}\meaning\foo\par
```

macro:->1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12

macro:->2:4:3:2:9:0:2:0:0:8:1:7:6:6:4:0:0:0:0

macro:->\bat \FOO \baz \FOO \biz \FOO \buz

macro:->\bat \FOO \baz \FOO \biz \FOO \buz

macro:->\bat \FOO \baz \FOO \biz \FOO \buz

6.11. \xintApply

ff ★ `\xintApply{⟨macro⟩}{⟨list⟩}` expandably applies the one parameter macro `⟨macro⟩` to each item in the `⟨list⟩` given as second argument and returns a new list with these outputs: each item is given one after the other as parameter to `⟨macro⟩` which is expanded at that time (as usual, *i.e.* fully for what comes first), the results are braced and output together as a succession of braced items (if `⟨macro⟩` is defined to start with a space, the space will be gobbled and the `⟨macro⟩` will not be expanded; it is allowed to have its own arguments, the list items serve as last arguments to `⟨macro⟩`). Hence `\xintApply{⟨macro⟩}{{1}{2}{3}}` returns `{⟨macro⟩{1}}{⟨macro⟩{2}}{⟨macro⟩{3}}` where all instances of `⟨macro⟩` have been already *f-expanded*.

Being expandable, `\xintApply` is useful for example inside alignments where implicit groups make standard loops constructs usually fail. In such situation it is often not wished that the new list elements be braced, see `\xintApplyUnbraced`. The `⟨macro⟩` does not have to be expandable: `\xintApply` will try to expand it, the expansion may remain partial.

The `⟨list⟩` may itself be some macro expanding (in the previously described way) to the list of tokens to which the macro `⟨macro⟩` will be applied. For example, if the `⟨list⟩` expands to some posi-

tive number, then each digit will be replaced by the result of applying `\macro` on it.

```
\def\macro #1{\the\numexpr 9-#1\relax}
\xintApply\macro{\xintiiFac {20}}=7567097991823359999
```

fn ★ The macro `\xintApplyNoExpand` does the same job without the first initial expansion which gave the `<list>` of braced tokens to which `\macro` is applied.

6.12. `\xintApplyUnbraced`

ff ★ `\xintApplyUnbraced{\macro}{<list>}` is like `\xintApply`. The difference is that after having expanded its list argument, and applied `\macro` in turn to each item from the list, it reassembles the outputs without enclosing them in braces. The net effect is the same as doing

```
\xintListWithSep {}{\xintApply {\macro}{<list>}}
```

This is useful for preparing a macro which will itself define some other macros or make assignments, as the scope will not be limited by brace pairs.

```
\def\macro #1{\expandafter\def\csname myself#1\endcsname {#1}}
\xintApplyUnbraced\macro{\elta}{eltb}{eltc}
\begin{enumerate}[nosep,label=(\arabic{*})]
\item \meaning\myselfelta
\item \meaning\myselfeltb
\item \meaning\myselfeltc
\end{enumerate}
```

(1) `macro:->elta`
(2) `macro:->eltb`
(3) `macro:->eltc`


fn ★ The macro `\xintApplyUnbracedNoExpand` does the same job without the first initial expansion which gave the `<list>` of braced tokens to which `\macro` is applied.

6.13. `\xintSeq`

$\left[\begin{smallmatrix} \text{num} \\ \text{X} \end{smallmatrix} \right] \text{X} \text{X} \text{X} \star$ `\xintSeq[d]{x}{y}` generates expandably `{x}{x+d}...` up to and possibly including `{y}` if `d>0` or down to and including `{y}` if `d<0`. Naturally `{y}` is omitted if `y-x` is not a multiple of `d`. If `d=0` the macro returns `{x}`. If `y-x` and `d` have opposite signs, the macro returns nothing. If the optional argument `d` is omitted it is taken to be the sign of `y-x`. Hence `\xintSeq {1}{0}` is not empty but `{1}{0}`. But `\xintSeq [1]{1}{0}` is empty.

The arguments `x` and `y` are expanded inside a `\numexpr` so they may be count registers or a \TeX `\value{countername}`, or arithmetic with such things.

```
\xintListWithSep{,\hskip2pt plus 1pt minus 1pt }{\xintSeq {12}{-25}}
12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0, -1, -2, -3, -4, -5, -6, -7, -8, -9, -10, -11, -12, -13, -14, -15,
-16, -17, -18, -19, -20, -21, -22, -23, -24, -25
\xintiiSum{\xintSeq [3]{1}{1000}}
167167
```

 When the macro is used without the optional argument `d`, it can only generate up to about 5000 numbers, the precise value depends upon some \TeX memory parameter (input save stack).
With the optional argument `d` the macro proceeds differently (but less efficiently) and does not stress the input save stack.

6.14. `\xintloop`, `\xintbreakloop`, `\xintbreakloopaddo`, `\xintloopskiptonext`

☆ `\xintloop<stuff>\if<test>...\repeat` is an expandable loop compatible with nesting. However to break out of the loop one almost always need some un-expandable step. The cousin `\xintiloop` is `\xintloop` with an embedded expandable mechanism allowing to exit from the loop. The iterated macros may contain `\par` tokens or empty lines.

If a sub-loop is to be used all the material from the start of the main loop and up to the end of the entire subloop should be braced; these braces will be removed and do not create a group. The simplest to allow the nesting of one or more sub-loops is to brace everything between `\xintloop` and `\repeat`, being careful not to leave a space between the closing brace and `\repeat`.

As this loop and `\xintilloop` will primarily be of interest to experienced TeX macro programmers, my description will assume that the user is knowledgeable enough. Some examples in this document will be perhaps more illustrative than my attempts at explanation of use.

One can abort the loop with `\xintbreakloop`; this should not be used inside the final test, and one should expand the `\fi` from the corresponding test before. One has also `\xintbreakloopanddo` whose first argument will be inserted in the token stream after the loop; one may need a macro such as `\xint_afterfi` to move the whole thing after the `\fi`, as a simple `\expandafter` will not be enough.

One will usually employ some count registers to manage the exit test from the loop; this breaks expandability, see `\xintilloop` for an expandable integer indexed loop. Use in alignments will be complicated by the fact that cells create groups, and also from the fact that any encountered unexpandable material will cause the TeX input scanner to insert `\endtemplate` on each encountered `&` or `\cr`; thus `\xintbreakloop` may not work as expected, but the situation can be resolved via `\xint_firstofone{&}` or use of `\TAB` with `\def\tab{&}`. It is thus simpler for alignments to use rather than `\xintloop` either the expandable `\xintApplyUnbraced` or the non-expandable but alignment compatible `\xintApplyInline`, `\xintFor` or `\xintFor*`.

As an example, let us suppose we have two macros `\A{<i>}{<j>}` and `\B{<i>}{<j>}` behaving like (small) integer valued matrix entries, and we want to define a macro `\C{<i>}{<j>}` giving the matrix product (*i* and *j* may be count registers). We will assume that `\A[I]` expands to the number of rows, `\A[J]` to the number of columns and want the produced `\C` to act in the same manner. The code is very dispendious in use of `\count` registers, not optimized in any way, not made very robust (the defined macro can not have the same name as the first two matrices for example), we just wanted to quickly illustrate use of the nesting capabilities of `\xintloop`.³³

```
\newcount\rowmax \newcount\colmax \newcount\summax
\newcount\rowindex \newcount\colindex \newcount\sumindex
\newcount\tmpcount
\makeatletter
\def\MatrixMultiplication #1#2#3{%
  \rowmax #1[I]\relax
  \colmax #2[J]\relax
  \summax #1[J]\relax
  \rowindex 1
  \xintloop % loop over row index i
  {\colindex 1
  \xintloop % loop over col index k
  {\tmpcount 0
  \sumindex 1
  \xintloop % loop over intermediate index j
  \advance\tmpcount \numexpr #1\rowindex\sumindex*#2\sumindex\colindex\relax
  \ifnum\sumindex<\summax
  \advance\sumindex 1
  \repeat }%
  \expandafter\edef\csname\string#3{\the\rowindex.\the\colindex}\endcsname
  {\the\tmpcount}%
  \ifnum\colindex<\colmax
  \advance\colindex 1
```

³³ for a more sophisticated implementation of matrix multiplication, inclusive of determinants, inverses, and display utilities, with entries big integers or decimal numbers or even fractions see some code online posted from November 11, 2013.

```

\repeat }%
\ifnum\rowindex<\rowmax
\advance\rowindex 1
\repeat
\expandafter\edef\csname\string#3{I}\endcsname{\the\rowmax}%
\expandafter\edef\csname\string#3{J}\endcsname{\the\colmax}%
\def #3##1{\ifx[##1\expandafter\Matrix@helper@size
\else\expandafter\Matrix@helper@entry\fi #3{##1}}%
}%
\def\Matrix@helper@size #1#2#3{\csname\string#1{#3}\endcsname }%
\def\Matrix@helper@entry #1#2#3%
{\csname\string#1{\the\numexpr#2.\the\numexpr#3}\endcsname }%
\def\A #1{\ifx[#1\expandafter\A@size
\else\expandafter\A@entry\fi {#1}}%
\def\A@size #1#2{\ifx I#23\else4\fi}% 3rows, 4columns
\def\A@entry #1#2{\the\numexpr #1+#2-1\relax}% not pre-computed...
\def\B #1{\ifx[#1\expandafter\B@size
\else\expandafter\B@entry\fi {#1}}%
\def\B@size #1#2{\ifx I#24\else3\fi}% 4rows, 3columns
\def\B@entry #1#2{\the\numexpr #1-#2\relax}% not pre-computed...
\makeatother
\MatrixMultiplication\A\B\C \MatrixMultiplication\C\C\D
\MatrixMultiplication\C\D\E \MatrixMultiplication\C\E\F
\begin{multicols}2
\[\begin{pmatrix}
\A11&\A12&\A13&\A14\\
\A21&\A22&\A23&\A24\\
\A31&\A32&\A33&\A34
\end{pmatrix}
\times
\begin{pmatrix}
\B11&\B12&\B13\\
\B21&\B22&\B23\\
\B31&\B32&\B33\\
\B41&\B42&\B43
\end{pmatrix}
=
\begin{pmatrix}
\C11&\C12&\C13\\
\C21&\C22&\C23\\
\C31&\C32&\C33
\end{pmatrix}\]
\[\begin{pmatrix}
\C11&\C12&\C13\\
\C21&\C22&\C23\\
\C31&\C32&\C33
\end{pmatrix}^2 = \begin{pmatrix}
\D11&\D12&\D13\\
\D21&\D22&\D23\\
\D31&\D32&\D33
\end{pmatrix}\]
\[\begin{pmatrix}

```

```

\C11&\C12&\C13\\
\C21&\C22&\C23\\
\C31&\C32&\C33
\end{pmatrix}^3 = \begin{pmatrix}
\E11&\E12&\E13\\
\E21&\E22&\E23\\
\E31&\E32&\E33
\end{pmatrix}\]
\[\begin{pmatrix}
\C11&\C12&\C13\\
\C21&\C22&\C23\\
\C31&\C32&\C33
\end{pmatrix}^4 = \begin{pmatrix}
\F11&\F12&\F13\\
\F21&\F22&\F23\\
\F31&\F32&\F33
\end{pmatrix}\]
\end{multicols}

```

$$\begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 5 \\ 3 & 4 & 5 & 6 \end{pmatrix} \times \begin{pmatrix} 0 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix} = \begin{pmatrix} 20 & 10 & 0 \\ 26 & 12 & -2 \\ 32 & 14 & -4 \end{pmatrix}$$

$$\begin{pmatrix} 20 & 10 & 0 \\ 26 & 12 & -2 \\ 32 & 14 & -4 \end{pmatrix}^3 = \begin{pmatrix} 20880 & 10160 & -560 \\ 24624 & 11968 & -688 \\ 28368 & 13776 & -816 \end{pmatrix}$$

$$\begin{pmatrix} 20 & 10 & 0 \\ 26 & 12 & -2 \\ 32 & 14 & -4 \end{pmatrix}^2 = \begin{pmatrix} 660 & 320 & -20 \\ 768 & 376 & -16 \\ 876 & 432 & -12 \end{pmatrix}$$

$$\begin{pmatrix} 20 & 10 & 0 \\ 26 & 12 & -2 \\ 32 & 14 & -4 \end{pmatrix}^4 = \begin{pmatrix} 663840 & 322880 & -18080 \\ 781632 & 380224 & -21184 \\ 899424 & 437568 & -24288 \end{pmatrix}$$

6.15. `\xintilooop`, `\xintilooopindex`, `\xintouterilooopindex`, `\xintbreakilooop`, `\xintbreakilooopanddo`, `\xintilooopskiptonext`, `\xintilooopskipandredo`

☆ `\xintilooop[start+delta]<stuff>\if<test> ... \repeat` is a completely expandable nestable loop. complete expandability depends naturally on the actual iterated contents, and complete expansion will not be achievable under a sole *f-expansion*, as is indicated by the hollow star in the margin; thus the loop can be used inside an `\edef` but not inside arguments to the package macros. It can be used inside an `\xintexpr... \relax`. The `[start+delta]` is mandatory, not optional.

This loop benefits via `\xintilooopindex` to (a limited access to) the integer index of the iteration. The starting value `start` (which may be a `\count`) and increment `delta` (*id.*) are mandatory arguments. A space after the closing square bracket is not significant, it will be ignored. Spaces inside the square brackets will also be ignored as the two arguments are first given to a `\numexpr... \relax`. Empty lines and explicit `\par` tokens are accepted.

As with `\xintloop`, this tool will mostly be of interest to advanced users. For nesting, one puts inside braces all the material from the start (immediately after `[start+delta]`) and up to and inclusive of the inner loop, these braces will be removed and do not create a loop. In case of nesting, `\xintouterilooopindex` gives access to the index of the outer loop. If needed one could write on its model a macro giving access to the index of the outer outer loop (or even to the *nth* outer loop).

The `\xintilooopindex` and `\xintouterilooopindex` can not be used inside braces, and generally speaking this means they should be expanded first when given as argument to a macro, and that this macro receives them as delimited arguments, not braced ones. Or, but naturally this will break expandability, one can assign the value of `\xintilooopindex` to some `\count`. Both `\xintilooopindex` and

`\xintouterloopindex` extend to the literal representation of the index, thus in `\ifnum` tests, if it comes last one has to correctly end the macro with a `\space`, or encapsulate it in a `\numexpr..\relax`.

When the repeat-test of the loop is, for example, `\ifnum\xintloopindex<10 \repeat`, this means that the last iteration will be with `\xintloopindex=10` (assuming `delta=1`). There is also `\ifnum \xintloopindex=10 \else\repeat` to get the last iteration to be the one with `\xintloopindex=10`.

One has `\xintbreakiloop` and `\xintbreakiloopanddo` to abort the loop. The syntax of `\xintbreakiloopanddo` is a bit surprising, the sequence of tokens to be executed after breaking the loop is not within braces but is delimited by a dot as in:

```
\xintbreakiloopanddo <afterloop>.etc.. etc... \repeat
```

The reason is that one may wish to use the then current value of `\xintloopindex` in `<afterloop>` but it can't be within braces at the time it is evaluated. However, it is not that easy as `\xintloopindex` must be expanded before, so one ends up with code like this:

```
\expandafter\xintbreakiloopanddo\expandafter\macro\xintloopindex.%
etc.. etc.. \repeat
```

As moreover the `\fi` from the test leading to the decision of breaking out of the loop must be cleared out of the way, the above should be a branch of an expandable conditional test, else one needs something such as:

```
\xint_afterfi{\expandafter\xintbreakiloopanddo\expandafter\macro\xintloopindex.}%
\fi etc..etc.. \repeat
```

There is `\xintloopskiptonext` to abort the current iteration and skip to the next, `\xintloopskipandredo` to skip to the end of the current iteration and redo it with the same value of the index (something else will have to change for this not to become an eternal loop...).

Inside alignments, if the looped-over text contains a `&` or a `\cr`, any un-expandable material before a `\xintloopindex` will make it fail because of `\endtemplate`; in such cases one can always either replace `&` by a macro expanding to it or replace it by a suitable `\firstofone{&}`, and similarly for `\cr`.

As an example, let us construct an `\edef\z{...}` which will define `\z` to be a list of prime numbers:

```
\begingroup
\edef\z
{\xintloop [10001+2]
 {\xintloop [3+2]
  \ifnum\xintouterloopindex<\numexpr\xintloopindex*\xintloopindex\relax
   \xintouterloopindex,
   \expandafter\xintbreakiloop
  \fi
  \ifnum\xintouterloopindex=\numexpr
   (\xintouterloopindex/\xintloopindex)*\xintloopindex\relax
  \else
  \repeat
  }% no space here
 \ifnum \xintloopindex < 10999 \repeat }%
\meaning\z\endgroup
```

macro:->10007, 10009, 10037, 10039, 10061, 10067, 10069, 10079, 10091, 10093, 10099, 10103, 10111, 10133, 10139, 10141, 10151, 10159, 10163, 10169, 10177, 10181, 10193, 10211, 10223, 10243, 10247, 10253, 10259, 10267, 10271, 10273, 10289, 10301, 10303, 10313, 10321, 10331, 10333, 10337, 10343, 10357, 10369, 10391, 10399, 10427, 10429, 10433, 10453, 10457, 10459, 10463, 10477, 10487, 10499, 10501, 10513, 10529, 10531, 10559, 10567, 10589, 10597, 10601, 10607, 10613, 10627, 10631, 10639, 10651, 10657, 10663, 10667, 10687, 10691, 10709, 10711, 10723, 10729, 10733, 10739, 10753, 10771, 10781, 10789, 10799, 10831, 10837, 10847, 10853, 10859, 10861, 10867, 10883, 10889, 10891, 10903, 10909, 10937, 10939, 10949, 10957, 10973, 10979, 10987, 10993, and we should have taken

some steps to not have a trailing comma, but the point was to show that one can do that in an `\edef`! See also [subsection 7.3](#) which extracts from this code its way of testing primality.

Let us create an alignment where each row will contain all divisors of its first entry. Here is the output, thus obtained without any count register:

```
\begin{multicols}2
\tabskiplex \normalcolor
\halign{&\hfil#\hfil\cr
  \xintloop [1+1]
  {\xexpandafter\bfseries\xintloopindex &
  \xintloop [1+1]
  \ifnum\xintouterloopindex=\numexpr
    (\xintouterloopindex/\xintloopindex)*\xintloopindex\relax
  \xintloopindex&\fi
  \ifnum\xintloopindex<\xintouterloopindex\space % CRUCIAL \space HERE
  \repeat \cr }%
  \ifnum\xintloopindex<30
  \repeat
}
\end{multicols}
```

1 1	16 1 2 4 8 16
2 1 2	17 1 17
3 1 3	18 1 2 3 6 9 18
4 1 2 4	19 1 19
5 1 5	20 1 2 4 5 10 20
6 1 2 3 6	21 1 3 7 21
7 1 7	22 1 2 11 22
8 1 2 4 8	23 1 23
9 1 3 9	24 1 2 3 4 6 8 12 24
10 1 2 5 10	25 1 5 25
11 1 11	26 1 2 13 26
12 1 2 3 4 6 12	27 1 3 9 27
13 1 13	28 1 2 4 7 14 28
14 1 2 7 14	29 1 29
15 1 3 5 15	30 1 2 3 5 6 10 15 30

We wanted this first entry in bold face, but `\bfseries` leads to unexpandable tokens, so the `\expandafter` was necessary for `\xintloopindex` and `\xintouterloopindex` not to be confronted with a hard to digest `\endtemplate`. An alternative way of coding:

```
\tabskiplex
\def\firstofone #1{#1}%
\halign{&\hfil#\hfil\cr
  \xintloop [1+1]
  {\bfseries\xintloopindex\firstofone{&}%
  \xintloop [1+1] \ifnum\xintouterloopindex=\numexpr
  (\xintouterloopindex/\xintloopindex)*\xintloopindex\relax
  \xintloopindex\firstofone{&}\fi
  \ifnum\xintloopindex<\xintouterloopindex\space % \space is CRUCIAL
  \repeat \firstofone{\cr}}%
  \ifnum\xintloopindex<30 \repeat }
```


The next utilities are not compatible with expansion-only context.

6.16. `\xintApplyInline`

`o *f` `\xintApplyInline{\macro}{\list}` works non expandably. It applies the one-parameter `\macro` to the first element of the expanded list (`\macro` may have itself some arguments, the list item will be appended as last argument), and is then re-inserted in the input stream after the tokens resulting from this first expansion of `\macro`. The next item is then handled.

This is to be used in situations where one needs to do some repetitive things. It is not expandable and can not be completely expanded inside a macro definition, to prepare material for later execution, contrarily to what `\xintApply` or `\xintApplyUnbraced` achieve.

```
\def\Macro #1{\advance\cnta #1 , \the\cnta}
\cnta 0
0\xintApplyInline\Macro {3141592653}.
```

0, 3, 4, 8, 9, 14, 23, 25, 31, 36, 39. The first argument `\macro` does not have to be an expandable macro.

`\xintApplyInline` submits its second, token list parameter to an *f-expansion*. Then, each *unbraced* item will also be *f-expanded*. This provides an easy way to insert one list inside another. *Braced* items are not expanded. Spaces in-between items are gobbled (as well as those at the start or the end of the list), but not the spaces *inside* the braced items.

`\xintApplyInline`, despite being non-expandable, does survive to contexts where the executed `\macro` closes groups, as happens inside alignments with the tabulation character `&`. This tabular provides an example:

```
\centerline{\normalcolor\begin{tabular}{ccc}
  $N$ & $N^2$ & $N^3$ \\ \hline
  \def\Row #1{ #1 & \xintiiSqr {#1} & \xintiiPow {#1}{3} \\ \hline }%
  \xintApplyInline \Row {\xintCSVtoList{17,28,39,50,61}}
\end{tabular}}\medskip
```

N	N ²	N ³
17	289	4913
28	784	21952
39	1521	59319
50	2500	125000
61	3721	226981

We see that despite the fact that the first encountered tabulation character in the first row close a group and thus erases `\Row` from \TeX 's memory, `\xintApplyInline` knows how to deal with this.

Using `\xintApplyUnbraced` is an alternative: the difference is that this would have prepared all rows first and only put them back into the token stream once they are all assembled, whereas with `\xintApplyInline` each row is constructed and immediately fed back into the token stream: when one does things with numbers having hundreds of digits, one learns that keeping on hold and shuffling around hundreds of tokens has an impact on \TeX 's speed (make this ``thousands of tokens'' for the impact to be noticeable).

One may nest various `\xintApplyInline`'s. For example (see the [table](#) on the next page):

```
\begin{figure*}[ht!]
  \centering\phantomsection\label{float}
  \def\Row #1{#1:\xintApplyInline {\Item {#1}}{0123456789}\ }%
  \def\Item #1#2{&\xintiiPow {#1}{#2}}%
  \centeredline {\begin{tabular}{ccccccccc} &0&1&2&3&4&5&6&7&8&9 \\ \hline
    \xintApplyInline \Row {0123456789}
  \end{tabular}}
```

`\end{figure*}`

	0	1	2	3	4	5	6	7	8	9
0:	1	0	0	0	0	0	0	0	0	0
1:	1	1	1	1	1	1	1	1	1	1
2:	1	2	4	8	16	32	64	128	256	512
3:	1	3	9	27	81	243	729	2187	6561	19683
4:	1	4	16	64	256	1024	4096	16384	65536	262144
5:	1	5	25	125	625	3125	15625	78125	390625	1953125
6:	1	6	36	216	1296	7776	46656	279936	1679616	10077696
7:	1	7	49	343	2401	16807	117649	823543	5764801	40353607
8:	1	8	64	512	4096	32768	262144	2097152	16777216	134217728
9:	1	9	81	729	6561	59049	531441	4782969	43046721	387420489

One could not move the definition of `\Item` inside the tabular, as it would get lost after the first `&`. But this works:

```
\begin{tabular}{cccccccccc}
&&0&&1&&2&&3&&4&&5&&6&&7&&8&&9\\ \hline
\def\Row #1{#1:\xintApplyInline {\&\xintiiPow {#1}}{0123456789}\\ }%
\xintApplyInline \Row {0123456789}
\end{tabular}
```

A limitation is that, contrarily to what one may have expected, the `\macro` for an `\xintApplyInline` can not be used to define the `\macro` for a nested sub-`\xintApplyInline`. For example, this does not work:

```
\def\Row #1{#1:\def\Item ##1{\&\xintiiPow {#1}{##1}}%
\xintApplyInline \Item {0123456789}\\ }%
\xintApplyInline \Row {0123456789} % does not work
```

But see `\xintFor`.

6.17. `\xintFor`, `\xintFor*`

on `\xintFor` is a new kind of for loop.³⁴ Rather than using macros for encapsulating list items, its behaviour is like a macro with parameters: `#1`, `#2`, ..., `#9` are used to represent the items for up to nine levels of nested loops. Here is an example:

```
\xintFor #9 in {1,2,3} \do {%
\xintFor #1 in {4,5,6} \do {%
\xintFor #3 in {7,8,9} \do {%
\xintFor #2 in {10,11,12} \do {%
$$#9\times#1\times#3\times#2=\xintiiPrd{#1}{#2}{#3}{#9}}$}}$}}
```

This example illustrates that one does not have to use `#1` as the first one: the order is arbitrary. But each level of nesting should have its specific macro parameter. Nine levels of nesting is presumably overkill, but I did not know where it was reasonable to stop. `\par` tokens are accepted in both the comma separated list and the replacement text.

T_EXnical notes:

- The `#1` is replaced in the iterated-over text exactly as in general **T_EX** macros or **L^AT_EX** commands. This spares the user quite a few `\expandafter`'s or other tricks needed with loops which have the values encapsulated in macros, like **L^AT_EX**'s `\@for` and `\@tfor`.

³⁴ first introduced with `xint 1.09c` of 2013/10/09.

- `\xintFor` (and `\xintFor*`) isn't purely expandable: one can not use it inside an `\edef`. But it may be used, as will be shown in examples, in some contexts such as \TeX 's `tabular` which are usually hostile to non-expandable loops.
- `\xintFor` (and `\xintFor*`) does some assignments prior to executing each iteration of the replacement text, but it acts purely expandably after the last iteration, hence if for example the replacement text ends with a `\\`, the loop can be used insided a `tabular` and be followed by a `\hline` without creating the dreaded `'Misplaced \noalign'` error.
- As stated in previous item the first iteration follows some non-expandable internal dealings. This means for example that in \TeX , one can not inject a `\multicolumn` in the first iteration. Sometimes one way work around this by injecting father `&\multicolumn` or `\\ \m, ulticolumn`.
- It does not create groups.
- It makes no global assignments.
- The iterated replacement text may close a group which was opened even before the start of the loop (typical example being with `&` in alignments).

```
\begin{tabular}{rccccc}
\hline
\xintFor #1 in {A, B, C} \do {%
#1:\xintFor #2 in {a, b, c, d, e} \do {&(\$ #2 \to #1 $)}\\ }%
\hline
\end{tabular}
```

```
A: (a → A) (b → A) (c → A) (d → A) (e → A)
B: (a → B) (b → B) (c → B) (d → B) (e → B)
C: (a → C) (b → C) (c → C) (d → C) (e → C)
```

- There is no facility provided which would give access to a count of the number of iterations as it is technically not easy to do so it in a way working with nested loops while maintaining the `'expandable after done'` property; something in the spirit of `\xint-loopindex` is possible but this approach would bring its own limitations and complications. Hence the user is invited to update her own count or \TeX counter or macro at each iteration, if needed.
- A `\macro` whose definition uses internally an `\xintFor` loop may be used inside another `\xintFor` loop even if the two loops both use the same macro parameter. The loop definition inside `\macro` must use `##` as is the general rule for definitions done inside macros.
- `\xintFor` is for comma separated values and `\xintFor*` for lists of braced items; their respective expansion policies differ. They are described later.

Regarding `\xintFor`:

- the spaces between the various declarative elements are all optional,
- in the list of comma separated values, spaces around the commas or at the start and end are ignored,
- if an item must contain itself its own commas, then it should be braced, and the braces will be removed before feeding the iterated-over text,
- the list may be a macro, it is expanded only once,
- items are not pre-expanded. The first item should be braced or start with a space if the list is explicit and the item should not be pre-expanded,

- empty items give empty #1's in the replacement text, they are not skipped,
- an empty list executes once the replacement text with an empty parameter value,
- the list, if not a macro, must be braced.

**fn* Regarding `\xintFor*`:

- it handles lists of braced items (or naked tokens),
- it *f-expands* the list,
- and more generally it *f-expands* each naked token encountered before assigning the #1 values (gobbling spaces in the process); this makes it easy to simulate concatenation of multiple lists `\x`, `\y`: if `\x` expands to `{1}{2}{3}` and `\y` expands to `{4}{5}{6}` then `{\x\y}` as argument to `\xintFor*` has the same effect as `{{1}{2}{3}{4}{5}{6}}`.

For a further illustration see the use of `\xintFor*` at the end of [subsection 3.18](#).

- spaces at the start, end, or in-between items are gobbled (but naturally not the spaces inside braced items),
- except if the list argument is a macro (with no parameters), it must be braced.,
- an empty list leads to an empty result.

The macro `\xintSeq` which generates arithmetic sequences is to be used with `\xintFor*` as its output consists of successive braced numbers (given as digit tokens).

```
\xintFor* #1 in {\xintSeq [+2]{-7}{+2}}\do {stuff
  with #1\xintifForLast{\par}{\newline}}
```

```
stuff with -7
stuff with -5
stuff with -3
stuff with -1
stuff with 1
```

When nesting `\xintFor*` loops, using `\xintSeq` in the inner loops is inefficient, as the arithmetic sequence will be re-created each time. A more efficient style is:

```
\edef\innersequence {\xintSeq[+2]{-50}{50}}%
\xintFor* #1 in {\xintSeq {13}{27}} \do
  {\xintFor* #2 in \innersequence \do {stuff with #1 and #2}%
  .. some other macros .. }
```

This is a general remark applying for any nesting of loops, one should avoid recreating the inner lists of arguments at each iteration of the outer loop.

When the loop is defined inside a macro for later execution the # characters must be doubled.³⁵ For example:

```
\def\T{\def\z {}%
  \xintFor* ##1 in {{u}{v}{w}} \do {%
    \xintFor ##2 in {x,y,z} \do {%
      \expandafter\def\expandafter\z\expandafter {\z\sep (##1,##2)} }%
  }%
}%
\T\def\sep {\def\sep{, }}\z
```

```
(u,x), (u,y), (u,z), (v,x), (v,y), (v,z), (w,x), (w,y), (w,z)
```

Similarly when the replacement text of `\xintFor` defines a macro with parameters, the macro character # must be doubled.

The iterated macros as well as the list items are allowed to contain explicit `\par` tokens.

6.18. `\xintifForFirst`, `\xintifForLast`

nn ★ `\xintifForFirst` {YES branch}{NO branch} and `\xintifForLast` {YES branch}{NO branch} execute the YES or NO branch if the `\xintFor` or `\xintFor*` loop is currently in its first, respectively last, iteration.

³⁵ sometimes what seems to be a macro argument isn't really; in `\raisebox{1cm}{\xintFor #1 in {a,b,c}\do {#1}}` no doubling should be done.

Designed to work as expected under nesting (but see frame next.) Don't forget an empty brace pair `{}` if a branch is to do nothing. May be used multiple times in the replacement text of the loop.

Pay attention to these implementation features:

- if an inner `\xintFor` loop is positioned before the `\xintifForFirst` or `\xintifForLast` of the outer loop it will contaminate their settings. This applies also naturally if the inner loop arises from the expansion of some macro located before the outer conditionals.

One fix is to make sure that the outer conditionals are expanded before the inner loop is executed, e.g. this will be the case if the inner loop is located inside one of the branches of the conditional.

Another approach is to enclose, if feasible, the inner loop in a group of its own.

- if the replacement text closes a group (e.g. from a `&` inside an alignment), the conditionals will lose their ascribed meanings and end up possibly undefined, depending whether there is some outer loop whose execution started before the opening of the group.

The fix is to arrange things so that the conditionals are expanded before \TeX encounters the closing-group token.

6.19. `\xintBreakFor`, `\xintBreakForAndDo`

One may immediately terminate an `\xintFor` or `\xintFor*` loop with `\xintBreakFor`.

As it acts by clearing up all the rest of the replacement text when encountered, it will not work from inside some `\if...\fi` without suitable `\expandafter` or swapping technique.

Also it can't be used from inside braces as from there it can't see the end of the replacement text.

There is also `\xintBreakForAndDo`. Both are illustrated by various examples in the next section which is devoted to ```forever''` loops.

6.20. `\xintintegers`, `\xintdimensions`, `\xintrationals`

If the list argument to `\xintFor` (or `\xintFor*`, both are equivalent in this context) is `\xintintegers` (equivalently `\xintegers`) or more generally `\xintintegers[start+delta]` (the whole within braces!)³⁶, then `\xintFor` does an infinite iteration where `#1` (or `#2`, `...`, `#9`) will run through the arithmetic sequence of (short) integers with initial value `start` and increment `delta`, `a` (default values: `start=1`, `delta=1`; if the optional argument is present it must contain both of them, and they may be explicit integers, or macros or count registers). The `#1` (or `#2`, `...`, `#9`) will stand for `\numexpr <opt sign><digits>\relax`, and the literal representation as a string of digits can thus be obtained as `\the#1` or `\number#1`. Such a `#1` can be used in an `\ifnum` test with no need to be postfixed with a space or a `\relax` and one should not add them.

If the list argument is `\xintdimensions` or more generally `\xintdimensions[start+delta]` (within braces!), then `\xintFor` does an infinite iteration where `#1` (or `#2`, `...`, `#9`) will run through the arithmetic sequence of dimensions with initial value `start` and increment `delta`. Default values: `start=0pt`, `delta=1pt`; if the optional argument is present it must contain both of them, and they may be explicit specifications, or macros, or dimen registers, or length macros in \TeX (the stretch and shrink components will be discarded). The `#1` will be `\dimexpr <opt sign><digits>sp\relax`, from which one can get the literal (approximate) representation in points via `\the#1`. So `#1` can be used anywhere \TeX expects a dimension (and there is no need in conditionals to insert a `\relax`, and one should not do it), and to print its value one uses `\the#1`. The chosen representation

³⁶ the `start+delta` optional specification may have extra spaces around the plus sign of near the square brackets, such spaces are removed. The same applies with `\xintdimensions` and `\xintrationals`.

guarantees exact incrementation with no rounding errors accumulating from converting into points at each step.

If the list argument to `\xintFor` (or `\xintFor*`) is `\xinrationals` or more generally `\xintrationals[start+delta]` (*within braces!*), then `\xintFor` does an infinite iteration where `#1` (or `#2`, . . . , `#9`) will run through the arithmetic sequence of `xintfrac` fractions with initial value `start` and increment `delta` (default values: `start=1/1`, `delta=1/1`). This loop works *only with `xintfrac` loaded*. if the optional argument is present it must contain both of them, and they may be given in any of the formats recognized by `xintfrac` (fractions, decimal numbers, numbers in scientific notations, numerators and denominators in scientific notation, etc. . . .), or as macros or count registers (if they are short integers). The `#1` (or `#2`, . . . , `#9`) will be an `a/b` fraction (without a `[n]` part), where the denominator `b` is the product of the denominators of `start` and `delta` (for reasons of speed `#1` is not reduced to irreducible form, and for another reason explained later `start` and `delta` are not put either into irreducible form; the input may use explicitly `\xintIrr` to achieve that).

```
\begingroup\small
\noindent\parbox{\dimexpr\linewidth-3em}{\color[named]{OrangeRed}%
\xintFor #1 in {\xinrationals [10/21+1/21]} \do
{#1=\xintifInt {#1}
{\textcolor{blue}{\xintTrunc{10}{#1}}}
{\xintTrunc{10}{#1}}% display in blue if an integer
\xintifGt {#1}{1.123}{\xintBreakFor}{, }%
}}
\endgroup\smallskip
10/21=0.4761904761,      11/21=0.5238095238,      12/21=0.5714285714,      13/21=0.6190476190,
14/21=0.6666666666,      15/21=0.7142857142,      16/21=0.7619047619,      17/21=0.8095238095,
18/21=0.8571428571,      19/21=0.9047619047,      20/21=0.9523809523,      21/21=1.0000000000,
22/21=1.0476190476, 23/21=1.0952380952, 24/21=1.1428571428
```

The example above confirms that computations are done exactly, and illustrates that the two initial (reduced) denominators are not multiplied when they are found to be equal. It is thus recommended to input `start` and `delta` with a common smallest possible denominator, or as fixed point numbers with the same numbers of digits after the decimal mark; and this is also the reason why `start` and `delta` are not by default made irreducible. As internally the computations are done with numerators and denominators completely expanded, one should be careful not to input numbers in scientific notation with exponents in the hundreds, as they will get converted into as many zeroes.

```
\noindent\parbox{\dimexpr.7\linewidth}{\raggedright
\xintFor #1 in {\xinrationals [0.000+0.125]} \do
{\edef\tmp{\xintTrunc{3}{#1}}%
\xintifInt {#1}
{\textcolor{blue}{\tmp}}
{\tmp}%
\xintifGt {#1}{2}{\xintBreakFor}{, }%
}}\smallskip
0, 0.125, 0.250, 0.375, 0.500, 0.625, 0.750, 0.875, 1.000, 1.125,
1.250, 1.375, 1.500, 1.625, 1.750, 1.875, 2.000, 2.125
```

We see here that `\xintTrunc` outputs (deliberately) zero as 0, not (here) 0.000, the idea being not to lose the information that the truncated thing was truly zero. Perhaps this behaviour should be changed? or made optional? Anyhow printing of fixed points numbers should be dealt with via dedicated packages such as `numprint` or `siunitx`.

6.21. `\xintForpair`, `\xintForthree`, `\xintForfour`

on The syntax is illustrated in this example. The notation is the usual one for `n`-uples, with paren-

theses and commas. Spaces around commas and parentheses are ignored.

```
{\centering\begin{tabular}{cccc}
\XintForpair #1#2 in { ( A , a ) , ( B , b ) , ( C , c ) } \do {%
\XintForpair #3#4 in { ( X , x ) , ( Y , y ) , ( Z , z ) } \do {%
\Biggl(\begin{tabular}{cc}
-#1- & -#3-\ \\
-#4- & -#2-\ \\
\end{tabular}\Biggr)$&\\\noalign{\vskip1\jot}}%
\end{tabular}\}
```

$$\begin{pmatrix} -A- & -X- \\ -x- & -a- \end{pmatrix} \begin{pmatrix} -A- & -Y- \\ -y- & -a- \end{pmatrix} \begin{pmatrix} -A- & -Z- \\ -z- & -a- \end{pmatrix} \\
\begin{pmatrix} -B- & -X- \\ -x- & -b- \end{pmatrix} \begin{pmatrix} -B- & -Y- \\ -y- & -b- \end{pmatrix} \begin{pmatrix} -B- & -Z- \\ -z- & -b- \end{pmatrix} \\
\begin{pmatrix} -C- & -X- \\ -x- & -c- \end{pmatrix} \begin{pmatrix} -C- & -Y- \\ -y- & -c- \end{pmatrix} \begin{pmatrix} -C- & -Z- \\ -z- & -c- \end{pmatrix}$$

`\XintForpair` must be followed by either `#1#2`, `#2#3`, `#3#4`, ..., or `#8#9` with `#1` usable as an alias for `#1#2`, `#2` as alias for `#2#3`, etc ... and similarly for `\XintForthree` (using `#1#2#3` or simply `#1`, `#2#3#4` or simply `#2`, ...) and `\XintForfour` (with `#1#2#3#4` etc...).

Nesting works as long as the macro parameters are distinct among `#1`, `#2`, ..., `#9`. A macro which expands to an `\XintFor` or a `\XintFor(pair,three,four)` can be used in another one with no constraint about using distinct macro parameters.

`\par` tokens are accepted in both the comma separated list and the replacement text.

6.22. `\XintAssign`

`\XintAssign<braced things>\to<as many cs as they are things>` defines (without checking if something gets overwritten) the control sequences on the right of `\to` to expand to the successive tokens or braced items located to the left of `\to`. `\XintAssign` is not an expandable macro.

f-expansion is first applied to the material in front of `\XintAssign` which is fetched as one argument if it is braced. Then the expansion of this argument is examined and successive items are assigned to the macros following `\to`. There must be exactly as many macros as items. No check is done. The macro assignments are done with removal of one level of brace pairs from each item.

After the initial *f-expansion*, each assigned (brace-stripped) item will be expanded according to the setting of the optional parameter.

For example `\XintAssign [e]...` means that all assignments are done using `\edef`. With `[f]` the assignments will be made using `\fdef`. The default is simply to make the definitions with `\def`, corresponding to an empty optional parameter `[]`. Possibilities for the optional parameter are: `[]`, `[g]`, `[e]`, `[x]`, `[o]`, `[go]`, `[oo]`, `[goo]`, `[f]`, `[gf]`. For example `[oo]` means a double expansion.

```
\XintAssign \xintiiDivision{1000000000000}{133333333}\to\Q\R
\meaning\Q\newline
\meaning\R\newline
\XintAssign {\xintiiDivision{1000000000000}{133333333}}\to\X
\meaning\X\newline
\XintAssign [oo]{\xintiiDivision{1000000000000}{133333333}}\to\X
\meaning\X\newline
\XintAssign \xintiiPow{7}{13}\to\SevenToThePowerThirteen
\meaning\SevenToThePowerThirteen\par
```

macro:->7500

macro:->2500

macro:->\xintiiDivision {1000000000000}{133333333}

macro:->{7500}{2500}

macro:->96889010407

Two special cases:

- if after this initial expansion no brace is found immediately after `\xintAssign`, it is assumed that there is only one control sequence following `\to`, and this control sequence is then defined via `\def` (or what is set-up by the optional parameter) to expand to the material between `\xintAssign` and `\to`.
- if the material between `\xintAssign` and `\to` is enclosed in two brace pairs, the first brace pair is removed, then the *f-expansion* is immediately stopped by the inner brace pair, hence `\xintAssign` now finds a unique item and thus defines only a single macro to be this item, which is now stripped of the second pair of braces.

Note: prior to release 1.09j, `\xintAssign` did an `\edef` by default for each item assignment but it now does `\def` corresponding to no or empty optional parameter.

It is allowed for the successive braced items to be separated by spaces. They are removed during the assignments. But if a single macro is defined (which happens if the argument after *f-expansion* does not start with a brace), naturally the scooped up material has all intervening spaces, as it is considered a single item. But an upfront initial space will have been absorbed by *f-expansion*.

```
\def\X{ {a} {b} {c} {d} }\def\Y { u {a} {b} {c} {d} }
\xintAssign\X\to\A\B\C\D
\xintAssign\Y\to\Z
\meaning\A, \meaning\B, \meaning\C, \meaning\D+++\\newline
\meaning\Z+++\\par
```

macro:->a, macro:->b, macro:->c, macro:->d+++

macro:->u {a} {b} {c} {d} +++

As usual successive space characters in input make for a single TeX space token.

6.23. `\xintAssignArray`

`\xintAssignArray`(*braced things*)`\to\myArray` first expands fully what comes immediately after `\xintAssignArray` and expects to find a list of braced things `{A}{B}...` (or tokens). It then defines `\myArray` as a macro with one parameter, such that `\myArray{x}` expands to give the *x*th braced thing of this original list (the argument `{x}` itself is fed to a `\numexpr` by `\myArray`, and `\myArray` expands in two steps to its output). With 0 as parameter, `\myArray{0}` returns the number *M* of elements of the array so that the successive elements are `\myArray{1}`, . . . , `\myArray{M}`.

```
\xintAssignArray \xintBezout {1000}{113}\to\Bez
```

will set `\Bez{0}` to 3, `\Bez{1}` to -20, `\Bez{2}` to 177, and `\Bez{3}` to 1: $-20 \times 1000 + 177 \times 113 = 1$. This macro is incompatible with expansion-only contexts.

`\xintAssignArray` admits an optional parameter, for example `\xintAssignArray [e]` means that the definitions of the macros will be made with `\edef`. The empty optional parameter (default) means that definitions are done with `\def`. Other possibilities: `[]`, `[o]`, `[oo]`, `[f]`. Contrarily to `\xintAssign` one can not use the `g` here to make the definitions global. For this, one should rather do `\xintAssignArray` within a group starting with `\globaldefs 1`.

6.24. `\xintDigitsOf`

fN This is a synonym for `\xintAssignArray`, to be used to define an array giving all the digits of a given (positive, else the minus sign will be treated as first item) number.

```
\xintDigitsOf\xintiiPow {7}{500}\to\digits
```

7^{500} has `\digits{0}`=423 digits, and the 123rd among them (starting from the most significant) is `\digits{123}`=3.

6.25. \xintRelaxArray

`\xintRelaxArray\myArray` (globally) sets to `\relax` all macros which were defined by the previous `\xintAssignArray` with `\myArray` as array macro.

7. Additional (old) examples with [xinttools](#) or [xintexpr](#) or both

Note: [xintexpr.sty](#) automatically loads [xinttools.sty](#).

The examples given here start to feel dated and are currently in need of some rewrite to better illustrate newer features of the package.

.1	More examples with dummy variables.....	106	.5	A table of factorizations	113
.2	Completely expandable prime test	107	.6	Another table of primes.....	114
.3	Another completely expandable prime test	108	.7	Factorizing again	116
.4	Miller-Rabin Pseudo-Primality expandably	110	.8	The Quick Sort algorithm illustrated.....	117

7.1. More examples with dummy variables

These examples were first added to this manual at the time of the 1.1 release (2014/10/29).

Prime numbers are always cool

```
\xinttheiexpr seq((seq((subs((x/:m)?{(m*m>x)?{1}{0}}{-1},m=2n+1))
??{break(0)}{omit}{break(1)},n=1++))?{x}{omit},
x=10001..[2]..10200)\relax
```

Prime numbers are always cool 10007, 10009, 10037, 10039, 10061, 10067, 10069, 10079, 10091, 10093, 10099, 10103, 10111, 10133, 10139, 10141, 10151, 10159, 10163, 10169, 10177, 10181, 10193

The syntax in this last example may look a bit involved (... and it is so I admit). First `x/:m` computes `x modulo m` (this is the modulo with respect to floored division). The `(x)?{yes}{no}` construct checks if `x` (which *must* be within parentheses) is true or false, i.e. non zero or zero. It then executes either the `yes` or the `no` branch, the non chosen branch is *not* evaluated. Thus if `m` divides `x` we are in the second (``false'`) branch. This gives a `-1`. This `-1` is the argument to a `??` branch which is of the type `(y)?{y<0}{y=0}{y>0}`, thus here the `y<0`, i.e., `break(0)` is chosen. This `0` is thus given to another `?` which consequently chooses `omit`, hence the number is not kept in the list. The numbers which survive are the prime numbers.

The first Fibonacci number beyond $|2^{64}|$, $|2^{64}|$, and the index are respectively

```
\xinttheiexpr subs(iterr(0,1;(@>N)?{break(@1,N,i)}{@1+@2},i=1++),N=2^64)\relax.
```

The first Fibonacci number beyond 2^{64} , 2^{64} , and the index are respectively 19740274219868223167, 18446744073709551616, 94.

One more recursion:

```
\def\syr #1{\xinttheiexpr
rseq#1; (@<=1)?{break(i)}{odd(@)?{3@+1}{@//2}},i=0++)\relax}
```

The $3x+1$ problem: `\syr{231}`\par

The $3x+1$ problem: 231, 694, 347, 1042, 521, 1564, 782, 391, 1174, 587, 1762, 881, 2644, 1322, 661, 1984, 992, 496, 248, 124, 62, 31, 94, 47, 142, 71, 214, 107, 322, 161, 484, 242, 121, 364, 182, 91, 274, 137, 412, 206, 103, 310, 155, 466, 233, 700, 350, 175, 526, 263, 790, 395, 1186, 593, 1780, 890, 445, 1336, 668, 334, 167, 502, 251, 754, 377, 1132, 566, 283, 850, 425, 1276, 638, 319, 958, 479, 1438, 719, 2158, 1079, 3238, 1619, 4858, 2429, 7288, 3644, 1822, 911, 2734, 1367, 4102, 2051, 6154, 3077, 9232, 4616, 2308, 1154, 577, 1732, 866, 433, 1300, 650, 325, 976, 488, 244, 122, 61, 184, 92, 46, 23, 70, 35, 106, 53, 160, 80, 40, 20, 10, 5, 16, 8, 4, 2, 1, 127

OK, a final one:

```
\def\syrMax #1{\xinttheiexpr iterr(#1,#1;even(i)?
{(@2<=1)?{break(@1,i//2)}
{odd(@2)?{3@2+1}{@2//2}}
{(@1>@2)?{@1}{@2}},i=0++)\relax }
```

With initial value 1161, the maximal intermediate value and the number of steps needed to reach 1 are respectively `\syrMax{1161}`.\par

With initial value 1161, the maximal intermediate value and the number of steps needed to reach 1 are respectively 190996, 181.

Look at the [Brent-Salamin algorithm implementation](#) for a more interesting recursion.

7.2. Completely expandable prime test

Let us now construct a completely expandable macro which returns 1 if its given input is prime and 0 if not:

```
\def\remainder #1#2{\the\numexpr #1-(#1/#2)*#2\relax }
\def\IsPrime #1%
  {\xintANDof {\xintApply {\remainder {#1}}{\xintSeq {2}{\xintiiSqrt{#1}}}}
```

This uses `\xintiiSqrt` and assumes its input is at least 5. Rather than `xint`'s own `\xintiiRem` we used a quicker `\numexpr` expression as we are dealing with short integers. Also we used `\xintANDof` which will return 1 only if all the items are non-zero. The macro is a bit silly with an even input, ok, let's enhance it to detect an even input:

```
\def\IsPrime #1%
  {\xintiiifOdd {#1}
   {\xintANDof % odd case
     {\xintApply {\remainder {#1}}
      {\xintSeq [2]{3}{\xintiiSqrt{#1}}}%
    }%
   }
  {\xintifEq {#1}{2}{1}{0}}%
 }
```

We used the `xint` expandable tests (on big integers or fractions) in order for `\IsPrime` to be *f-expandable*.

Our integers are short, but without `\expandafter`'s with `\@firstoftwo`, or some other related techniques, direct use of `\ifnum..fi` tests is dangerous. So to make the macro more efficient we are going to use the expandable tests provided by the package `etoolbox`³⁷. The macro becomes:

```
\def\IsPrime #1%
  {\ifnumodd {#1}
   {\xintANDof % odd case
     {\xintApply {\remainder {#1}}{\xintSeq [2]{3}{\xintiiSqrt{#1}}}}
   {\ifnumequal {#1}{2}{1}{0}}}
```

In the odd case however we have to assume the integer is at least 7, as `\xintSeq` generates an empty list if `#1=3` or `5`, and `\xintANDof` returns 1 when supplied an empty list. Let us ease up a bit `\xintANDof`'s work by letting it work on only 0's and 1's. We could use:

```
\def\IsNotDivisibleBy #1#2%
  {\ifnum\numexpr #1-(#1/#2)*#2=0 \expandafter 0\else \expandafter1\fi}
```

where the `\expandafter`'s are crucial for this macro to be *f-expandable* and hence work within the applied `\xintANDof`. Anyhow, now that we have loaded `etoolbox`, we might as well use:

```
\newcommand{\IsNotDivisibleBy}[2]{\ifnumequal{#1-(#1/#2)*#2}{0}{0}{1}}
```

Let us enhance our prime macro to work also on the small primes:

```
\newcommand{\IsPrime}[1] % returns 1 if #1 is prime, and 0 if not
  {\ifnumodd {#1}
   {\ifnumless {#1}{8}
    {\ifnumequal{#1}{1}{0}{1}}% 3,5,7 are primes
    {\xintANDof
     {\xintApply
      {\IsNotDivisibleBy {#1}}{\xintSeq [2]{3}{\xintiiSqrt{#1}}}}%
    }% END OF THE ODD BRANCH
   {\ifnumequal {#1}{2}{1}{0}}% EVEN BRANCH
  }
```

The input is still assumed positive. There is a deliberate blank before `\IsNotDivisibleBy` to use this feature of `\xintApply`: a space stops the expansion of the applied macro (and disappears). This expansion will be done by `\xintANDof`, which has been designed to skip everything as soon as

³⁷ <http://ctan.org/pkg/etoolbox>

it finds a false (i.e. zero) input. This way, the efficiency is considerably improved.

We did generate via the `\xintSeq` too many potential divisors though. Later sections give two variants: one with `\xintiloop` (subsection 7.3) which is still expandable and another one (subsection 7.6) which is a close variant of the `\IsPrime` code above but with the `\xintFor` loop, thus breaking expandability. The `xintiloop` variant does not first evaluate the integer square root, the `xintFor` variant still does. I did not compare their efficiencies.

Let us construct with this expandable primality test a table of the prime numbers up to 1000. We need to count how many we have in order to know how many tab stops one should add in the last row.³⁸ There is some subtlety for this last row. Turns out to be better to insert a `\\` only when we know for sure we are starting a new row; this is how we have designed the `\OneCell` macro. And for the last row, there are many ways, we use again `\xintApplyUnbraced` but with a macro which gobbles its argument and replaces it with a tabulation character. The `\xintFor*` macro would be more elegant here.

```

\newcounter{primecount}
\newcounter{cellcount}
\newcommand{\NbOfColumns}{13}
\newcommand{\OneCell}[1]{%
  \ifnumequal{\IsPrime{#1}}{1}
  {\stepcounter{primecount}
  \ifnumequal{\value{cellcount}}{\NbOfColumns}
  {\setcounter{cellcount}{1}#1}
  {\&\stepcounter{cellcount}#1}%
  } % was prime
  }% not a prime, nothing to do
}
\newcommand{\OneTab}[1]{&}
\begin{tabular}{|*{\NbOfColumns}{r}|}
\hline
2 \setcounter{cellcount}{1}\setcounter{primecount}{1}%
\xintApplyUnbraced \OneCell {\xintSeq [2]{3}{999}}%
\xintApplyUnbraced \OneTab
{\xintSeq [1]{1}{\the\numexpr\nbOfColumns-\value{cellcount}\relax}}%
\\
\hline
\end{tabular}
There are \arabic{primecount} prime numbers up to 1000.

```

The table has been put in `float` which appears on the following page. We had to be careful to use in the last row `\xintSeq` with its optional argument `[1]` so as to not generate a decreasing sequence from 1 to 0, but really an empty sequence in case the row turns out to already have all its cells (which doesn't happen here but would with a number of columns dividing 168).

7.3. Another completely expandable prime test

The `\IsPrime` macro from subsection 7.2 checked expandably if a (short) integer was prime, here is a partial rewrite using `\xintiloop`. We use the `etoolbox` expandable conditionals for convenience, but not everywhere as `\xintiloopindex` can not be evaluated while being braced. This is also the reason why `\xintbreakloopanddo` is delimited, and the next macro `\SmallestFactor` which returns the smallest prime factor exemplifies that. One could write more efficient completely expandable routines, the aim here was only to illustrate use of the general purpose `\xintiloop`. A little table giving the first values of `\SmallestFactor` follows, its coding uses `\xintFor`, which is described later; none of this uses count registers.

³⁸ although a tabular row may have less tabs than in the preamble, there is a problem with the vertical rule, if one does that.

2	3	5	7	11	13	17	19	23	29	31	37	41
43	47	53	59	61	67	71	73	79	83	89	97	101
103	107	109	113	127	131	137	139	149	151	157	163	167
173	179	181	191	193	197	199	211	223	227	229	233	239
241	251	257	263	269	271	277	281	283	293	307	311	313
317	331	337	347	349	353	359	367	373	379	383	389	397
401	409	419	421	431	433	439	443	449	457	461	463	467
479	487	491	499	503	509	521	523	541	547	557	563	569
571	577	587	593	599	601	607	613	617	619	631	641	643
647	653	659	661	673	677	683	691	701	709	719	727	733
739	743	751	757	761	769	773	787	797	809	811	821	823
827	829	839	853	857	859	863	877	881	883	887	907	911
919	929	937	941	947	953	967	971	977	983	991	997	

There are 168 prime numbers up to 1000.

```
% clean up possible left-over mess from previous examples
\let\IsPrime\undefined \let\SmallestFactor\undefined
\newcommand{\IsPrime}[1] % returns 1 if #1 is prime, and 0 if not
  {\ifnumodd {#1}
    {\ifnumless {#1}{8}
      {\ifnumequal{#1}{1}{0}{1}}% 3,5,7 are primes
      {\if
        \xintilooop [3+2]
        \ifnum#1<\numexpr\xintilooopindex*\xintilooopindex\relax
          \expandafter\xintbreakilooopanddo\expandafter1\expandafter.%
        \fi
        \ifnum#1=\numexpr (#1/\xintilooopindex)*\xintilooopindex\relax
        \else
        \repeat 00\expandafter0\else\expandafter1\fi
      }%
    }% END OF THE ODD BRANCH
    {\ifnumequal {#1}{2}{1}{0}}% EVEN BRANCH
  }%
\catcode`_ 11
\newcommand{\SmallestFactor}[1] % returns the smallest prime factor of #1>1
  {\ifnumodd {#1}
    {\ifnumless {#1}{8}
      {#1}% 3,5,7 are primes
      {\xintilooop [3+2]
        \ifnum#1<\numexpr\xintilooopindex*\xintilooopindex\relax
          \xint_afterfi{\xintbreakilooopanddo#1.}%
        \fi
        \ifnum#1=\numexpr (#1/\xintilooopindex)*\xintilooopindex\relax
          \xint_afterfi{\expandafter\xintbreakilooopanddo\xintilooopindex.}%
        \fi
        \iftrue\repeat
      }%
    }% END OF THE ODD BRANCH
    {2}% EVEN BRANCH
  }%
\catcode`_ 8
```

```
{\centering
\begin{tabular}{|c|*{10}c|}
\hline
\intFor #1 in {0,1,2,3,4,5,6,7,8,9}\do {\&\bfseries #1}\
\hline
\bfseries 0&--&--&2&3&2&5&2&7&2&3\
\intFor #1 in {1,2,3,4,5,6,7,8,9}\do
{\bfseries #1%
\intFor #2 in {0,1,2,3,4,5,6,7,8,9}\do
{\&\SmallestFactor{#1#2}}\}%
\hline
\end{tabular}\par
}
```

	0	1	2	3	4	5	6	7	8	9
0	--	--	2	3	2	5	2	7	2	3
1	2	11	2	13	2	3	2	17	2	19
2	2	3	2	23	2	5	2	3	2	29
3	2	31	2	3	2	5	2	37	2	3
4	2	41	2	43	2	3	2	47	2	7
5	2	3	2	53	2	5	2	3	2	59
6	2	61	2	3	2	5	2	67	2	3
7	2	71	2	73	2	3	2	7	2	79
8	2	3	2	83	2	5	2	3	2	89
9	2	7	2	3	2	5	2	97	2	3

7.4. Miller-Rabin Pseudo-Primality expandably

The `isPseudoPrime(n)` is usable in `\xintiiexpr`-essions and establishes if its (positive) argument is a Miller-Rabin PseudoPrime to the bases 2, 3, 5, 7, 11, 13, 17. If this is true and $n < 341550071728321$ (which has 15 digits) then n really is a prime number.

Similarly $n = 3825123056546413051$ (19 digits) is the smallest composite number which is a strong pseudo prime for bases 2, 3, 5, 7, 11, 13, 17, 19 and 23. It is easy to extend the code below to include these additional tests (we could make the list of tested bases an argument too, now that I think about it.)

For more information see

https://en.wikipedia.org/wiki/Miller%E2%80%93Rabin_primality_test#Deterministic_variants_of_the_test
and

http://primes.utm.edu/prove/prove2_3.html

In particular, according to JÄESCHKE *On strong pseudoprimes to several bases*, Math. Comp., 61 (1993) 915-926, if $n < 4,759,123,141$ it is enough to establish Rabin-Miller pseudo-primality to bases $a = 2, 7, 61$ to prove that n is prime. This range is enough for $\text{T}_{\text{E}}\text{X}$ numbers and we could then write a very fast expandable primality test for such numbers using only `\numexpr`. Left as an exercise...

```
% I ----- Modular Exponentiation
% Computes x^m modulo n (with m non negative).
% We will always use it with 1 < x < n
%
% With xint 1.4 we should use ? and ?? (although in the case at hand ifsgn()
% and if() would be ok but I should not say that).
%
\xintdefiifunc powmod_a(x, m, n) :=
  isone(m)?
    % m=1, return x modulo n
```

```

    { x /: n }
    % m > 1 test if odd or even and do recursive call
    { odd(m)? { x*sqr(powmod_a(x, m//2, n)) /: n }
      { sqr(powmod_a(x, m//2, n)) /: n }
    }
  }
;
\xintdefiifunc powmod(x, m, n) := (m)?{powmod_a(x, m, n)}{1};

%% Syntax used before xint 1.4:
% \xintdefiifunc powmod_a(x, m, n) :=
%   ifone(m,
%     % m=1, return x modulo n
%     x /: n,
%     % m > 1 test if odd or even and do recursive call
%     if(odd(m), (x*sqr(powmod_a(x, m//2, n))) /: n,
%       sqr(powmod_a(x, m//2, n)) /: n
%     )
%   );
\xintdefiifunc powmod(x, m, n) := if(m, powmod_a(x, m, n), 1);

% II ----- Miller-Rabin compositeness witness

% n=2^k m + 1 with m odd and k at least 1

% Choose 1<x<n.
% compute y=x^m modulo n
% if equals 1 we can't say anything
% if equals n-1 we can't say anything
% else put j=1, and
% compute repeatedly the square, incrementing j by 1 each time,
% thus always we have y^{2^{j-1}}
% -> if at some point n-1 mod n found, we can't say anything and break out
% -> if however we never find n-1 mod n before reaching
%   z=y^{2^{k-1}} with j=k
%   we then have z^2=x^{n-1}.
% Suppose z is not -1 mod n. If z^2 is 1 mod n, then n can be prime only if
% z is 1 mod n, and we can go back up, until initial y, and we have already
% excluded y=1. Thus if z is not -1 mod n and z^2 is 1 then n is not prime.
% But if z^2 is not 1, then n is not prime by Fermat. Hence (z not -1 mod n)
% implies (n is composite). (Miller test)

% let's use again xintexpr indecipherable (except to author) syntax. Of course
% doing it with macros only would be faster.

% Here \xintdefiifunc is not usable because not compatible with iter, break, ...
% but \xintNewFunction comes to the rescue.

\xintNewFunction{isCompositeWitness}[4]{% x=#1, n=#2, m=#3, k=#4
  subs((y==1)?{0}
    {iter(y;(j==#4)?{break(!(@==#2-1))}
      {(@==#2-1)?{break(0)}{sqr(@)/:#2}},j=1++)}
    ,y=powmod(#1,#3,#2))}

```

```

% added note (2018/03/07) it is possible in the above that m=#3 is never
% zero, so we should rather call powmod_a for a small gain, but I don't
% have time to re-read the code comments and settle this.

% III ----- Strong Pseudo Primes

% cf
% http://oeis.org/A014233
% http://mathworld.wolfram.com/Rabin-MillerStrongPseudoprimeTest.html>
% http://mathworld.wolfram.com/StrongPseudoprime.html>

% check if positive integer <49 si a prime.
% 2,3,5,7,11,13,17,19,23,29,31,37,41,43,47
\def\IsVerySmallPrime #1%
  {\ifnum#1=1 \xintdothis0\fi
  \ifnum#1=2 \xintdothis1\fi
  \ifnum#1=3 \xintdothis1\fi
  \ifnum#1=5 \xintdothis1\fi
  \ifnum#1=\numexpr (#1/2)*2\relax\xintdothis0\fi
  \ifnum#1=\numexpr (#1/3)*3\relax\xintdothis0\fi
  \ifnum#1=\numexpr (#1/5)*5\relax\xintdothis0\fi
  \xintorthat 1}

\xintNewFunction{isPseudoPrime}[1]{% n = #1
  (#1<49)?% use ? syntax to evaluate only what is needed
% prior to 1.4 we had \xintthe#1 here but the actual tokens represented
% by this #1 when isPseudoPrime() function expands have changed and
% the correct way is now \xintiieval{#1} to hand over explicit digits to
% the \IsVerySmallPrime macro.
  {\IsVerySmallPrime{\xintiieval{#1}}}
  {(even(#1))?
  {0}
  {subs(%
  % L expands to two values m, k hence isCompositeWitness does get
  % its four variables x, n, m, k
  isCompositeWitness(2, #1, L)?
  {0}%
  isCompositeWitness(3, #1, L)?
  {0}%
  isCompositeWitness(5, #1, L)?
  {0}%
  isCompositeWitness(7, #1, L)?
  {0}%
% above enough for N<3215031751 hence all TeX numbers
  isCompositeWitness(11, #1, L)?
  {0}%
% above enough for N<2152302898747, hence all 12-digits numbers
  isCompositeWitness(13, #1, L)?
  {0}%
% above enough for N<3474749660383
  isCompositeWitness(17, #1, L)?

```



```

        {0}%
% above enough for N<341550071728321
        {1}%
        }% not needed to comment-out end of lines spaces inside
        }% \xintexpr but this is too much of a habit for me with TeX!
        }% I left some after the ? characters.
        }%
        }%
        }% this computes (m, k) such that n = 2^k m + 1, m odd, k>=1
        , L=iter(#1//2;(even(@))?{@//2}{break(@,k)},k=1++)%
        }%
        }%
}

% if needed:
%\def\IsPseudoPrime #1{\xinttheiiexpr isPseudoPrime(#1)\relax}

\noindent The smallest prime number at least equal to 3141592653589 is
\xintiexpr
    seq(isPseudoPrime(3141592653589+n)?
        {break(3141592653589+n)}{omit}, n=0++)\relax.
% we could not use 3141592653589++ syntax because it works only with TeX numbers
\par

```

The smallest prime number at least equal to 3141592653589 is 3141592653601.

7.5. A table of factorizations

As one more example with `\xintilooop` let us use an alignment to display the factorization of some numbers. The loop will actually only play a minor rôle here, just handling the row index, the row contents being almost entirely produced via a macro `\factorize`. The factorizing macro does not use `\xintilooop` as it didn't appear to be the convenient tool. As `\factorize` will have to be used on `\xintilooopindex`, it has been defined as a delimited macro.

To spare some fractions of a second in the compilation time of this document (which has many many other things to do), 2147483629 and 2147483647, which turn out to be prime numbers, are not given to `factorize` but just typeset directly; this illustrates use of `\xintilooopskiptonext`.

The code next generates a `table` which has been made into a float appearing on page 115. Here is now the code for factorization; the conditionals use the package provided `\xint_firstoftwo` and `\xint_secondoftwo`, one could have employed rather \TeX 's own `\@firstoftwo` and `\@secondoftwo`, or, simpler still in \TeX context, the `\ifnumequal`, `\ifnumless` . . . , utilities from the package `etoolbox` which do exactly that under the hood. Only \TeX acceptable numbers are treated here, but it would be easy to make a translation and use the `xint` macros, thus extending the scope to big numbers; naturally up to a cost in speed.

The reason for some strange looking expressions is to avoid arithmetic overflow.

```

\catcode`_ 11
\def\abortfactorize #1\xint_secondoftwo\fi #2#3{\fi}

\def\factorize #1.{\ifnum#1=1 \abortfactorize\fi
    \ifnum\numexpr #1-2=\numexpr ((#1/2)-1)*2\relax
        \expandafter\xint_firstoftwo
    \else\expandafter\xint_secondoftwo
    \fi
    {2&\expandafter\factorize\the\numexpr#1/2.}%
    {\factorize_b #1.3.}}%

```

```

\def\factorize_b #1.#2.{\ifnum#1=1 \abortfactorize\fi
  \ifnum\numexpr #1-(#2-1)*#2<#2
    #1\abortfactorize
  \fi
  \ifnum \numexpr #1-#2=\numexpr ((#1/#2)-1)*#2\relax
    \expandafter\xint_firstoftwo
  \else\expandafter\xint_secondoftwo
  \fi
  {#2&\expandafter\factorize_b\the\numexpr#1/#2.#2.}%
  {\expandafter\factorize_b\the\numexpr #1\expandafter.%
    \the\numexpr #2+2.}%
}

\catcode`_ 8
\begin{figure*}[ht!]
\centering\phantomsection\label{floatfactorize}\normalcolor
\tabskip1ex
\centeredline{\vbox{\halign {\hfil\strut#\hfil&&\hfil#\hfil\cr\noalign{\hrule}
  \xintiloop ["7FFFFFFE0+1]
  \expandafter\bfseries\xintiloopindex &
  \ifnum\xintiloopindex="7FFFFFFD
    \number"7FFFFFFD\cr\noalign{\hrule}
  \expandafter\xintiloopskiptonext
  \fi
  \expandafter\factorize\xintiloopindex.\cr\noalign{\hrule}
  \ifnum\xintiloopindex<"7FFFFFFE
    \repeat
    \bfseries \number"7FFFFFFF&\number "7FFFFFFF\cr\noalign{\hrule}
  }}}
\centeredline{A table of factorizations}
\end{figure*}

```

7.6. Another table of primes

As a further example, let us dynamically generate a tabular with the first 50 prime numbers after 12345. First we need a macro to test if a (short) number is prime. Such a completely expandable macro was given in [subsection 7.2](#), here we consider a variant which will be slightly more efficient. This new `\IsPrime` has two parameters. The first one is a macro which it redefines to expand to the result of the primality test applied to the second argument. For convenience we use the `etoolbox` wrappers to various `\ifnum` tests, although here there isn't anymore the constraint of complete expandability (but using explicit `\if.. \fi` in tabulars has its quirks); equivalent tests are provided by `xint`, but they have some overhead as they are able to deal with arbitrarily big integers.

```

\def\IsPrime #1#2% #1=\Result, #2=tested number (assumed >0).
{\edef\TheNumber {\the\numexpr #2}% hence #2 may be a count or \numexpr.
  \ifnumodd {\TheNumber}
  {\ifnumgreater {\TheNumber}{1}
  {\edef\ItsSquareRoot{\xintiiSqrt \TheNumber}%
  \xintFor ##1 in {\xintintegers [3+2]}\do
  {\ifnumgreater {##1}{\ItsSquareRoot} % ##1 is a \numexpr.
    {\def#1{1}\xintBreakFor}
  }%
  \ifnumequal {\TheNumber}{(\TheNumber/##1)*##1}
  {\def#1{0}\xintBreakFor }

```

2147483616	2	2	2	2	2	3	2731	8191
2147483617	6733	318949						
2147483618	2	7	367	417961				
2147483619	3	3	23	353	29389			
2147483620	2	2	5	4603	23327			
2147483621	14741	145681						
2147483622	2	3	17	467	45083			
2147483623	79	967	28111					
2147483624	2	2	2	11	13	1877171		
2147483625	3	5	5	5	7	199	4111	
2147483626	2	19	37	1527371				
2147483627	47	53	862097					
2147483628	2	2	3	3	59652323			
2147483629	2147483629							
2147483630	2	5	6553	32771				
2147483631	3	137	263	19867				
2147483632	2	2	2	2	7	73	262657	
2147483633	5843	367531						
2147483634	2	3	12097	29587				
2147483635	5	11	337	115861				
2147483636	2	2	536870909					
2147483637	3	3	3	13	6118187			
2147483638	2	2969	361651					
2147483639	7	17	18046081					
2147483640	2	2	2	3	5	29	43	113 127
2147483641	2699	795659						
2147483642	2	23	46684427					
2147483643	3	715827881						
2147483644	2	2	233	1103	2089			
2147483645	5	19	22605091					
2147483646	2	3	3	7	11	31	151	331
2147483647	2147483647							

A table of factorizations

```

    {}%
  }}
  {\def#1{0}}% 1 is not prime
  {\ifnumequal {\TheNumber}{2}{\def#1{1}}{\def#1{0}}}%
}

```

As we used `\xintFor` inside a macro we had to double the `#` in its `#1` parameter. Here is now the code which creates the prime table (the table has been put in a `float`, which should be found on page 116):

```

\newcounter{primecount}
\newcounter{cellcount}
\begin{figure*}[ht!]
  \centering
  \begin{tabular}{|*{7}c|}
  \hline
  \setcounter{primecount}{0}\setcounter{cellcount}{0}%
  \xintFor #1 in {\xintintegers [12345+2]} \do
  % #1 is a \numexpr.

```

```

{\IsPrime\Result{#1}%
 \ifnumgreater{\Result}{0}
 {\stepcounter{primecount}%
 \stepcounter{cellcount}%
 \ifnumequal {\value{cellcount}}{7}
 {\the#1 \\setcounter{cellcount}{0}}
 {\the#1 &}}
 {}%
 \ifnumequal {\value{primecount}}{50}
 {\xintBreakForAndDo
 {\multicolumn {6}{l}{These are the first 50 primes after 12345.}}
 {}%
 }\hline
 \end{tabular}
 \end{figure*}

```

12347	12373	12377	12379	12391	12401	12409
12413	12421	12433	12437	12451	12457	12473
12479	12487	12491	12497	12503	12511	12517
12527	12539	12541	12547	12553	12569	12577
12583	12589	12601	12611	12613	12619	12637
12641	12647	12653	12659	12671	12689	12697
12703	12713	12721	12739	12743	12757	12763
12781	These are the first 50 primes after 12345.					

7.7. Factorizing again

Here is an *f-expandable* macro which computes the factors of an integer. It uses the `xint` macros only.

```

\catcode\@ 11
\let\factorize\relax
\newcommand\Factorize [1]
  {\romannumeral0\expandafter\factorize\expandafter{\romannumeral-`0#1}}%
\newcommand\factorize [1]{\xintiiifOne{#1}{ 1}{\factors@a #1.#1};}%
\def\factors@a #1.{\xintiiifOdd{#1}
  {\factors@c 3.#1.}%
  {\expandafter\factors@b \expandafter1\expandafter.\romannumeral0\xinthalff{#1}.}}%
\def\factors@b #1.#2.{\xintiiifOne{#2}
  {\factors@end {2, #1}}%
  {\xintiiifOdd{#2}{\factors@c 3.#2.{2, #1}}%
    {\expandafter\factors@b \the\numexpr #1+\@ne\expandafter.%
      \romannumeral0\xinthalff{#2}.}}%
  }%
\def\factors@c #1.#2.%
  \expandafter\factors@d\romannumeral0\xintiidivision {#2}{#1}{#1}{#2}%
}%
\def\factors@d #1#2#3#4{\xintiiifNotZero{#2}
  {\xintiiifGt{#3}{#1}
    {\factors@end {#4, 1}}% ultimate quotient is a prime with power 1
    {\expandafter\factors@c\the\numexpr #3+\tw@.#4.}}%
  {\factors@e 1.#3.#1.}%
}%

```

```

\def\factorse #1.#2.#3.{\xintiifOne{#3}
  {\factors@end {#2, #1}}%
  {\expandafter\factorse\romannumeral0\xintiidivision {#3}{#2}{#1}{#2}{#3}}%
}%
\def\factorse #1#2#3#4#5{\xintiifNotZero{#2}
  {\expandafter\factorse\the\numexpr #4+\tw@.#5.{#4, #3}}%
  {\expandafter\factorse\the\numexpr #3+\@ne.#4.#1.}%
}%
\def\factorse #1;{\xintlistwithsep{, }{\xintRevWithBraces {#1}}}%
\catcode`@ 12

```

The macro will be acceptably efficient only with numbers having somewhat small prime factors.

```
\Factorize{16246355912554185673266068721806243461403654781833}
```

```
16246355912554185673266068721806243461403654781833, 13, 5, 17, 8, 29, 5, 37, 6, 41, 4, 59, 6
```

It puts a little stress on the input save stack in order not be bothered with previously gathered things.³⁹

Its output is a comma separated list with the number first, then its prime factors with multiplicity. Let's produce something prettier:

```

\catcode`_ 11
\def\ShowFactors #1{\expandafter
  \ShowFactors_a\romannumeral-`0\Factorize{#1},\relax,\relax,}
\def\ShowFactors_a #1,{#1=\ShowFactors_b}
\def\ShowFactors_b #1,#2,{\if\relax#1\else#1^{#2}\expandafter\ShowFactors_b\fi}
\catcode`_ 8
$$\ShowFactors{16246355912554185673266068721806243461403654781833}$$

```

```
16246355912554185673266068721806243461403654781833 = 135178295376414596
```

If we only considered small integers, we could write pure `\numexpr` methods which would be very much faster (especially if we had a table of small primes prepared first) but still ridiculously slow compared to any non expandable implementation, not to mention use of programming languages directly accessing the CPU registers. . .

7.8. The Quick Sort algorithm illustrated

First a completely expandable macro which sorts a comma separated list of numbers.⁴⁰

The `\QSx` macro expands its list argument, which may thus be a macro; its comma separated items must expand to integers or decimal numbers or fractions or scientific notation as acceptable to `xintfrac`, but if an item is itself some (expandable) macro, this macro will be expanded each time the item is considered in a comparison test! This is actually good if the macro expands in one step to the digits, and there are many many digits, but bad if the macro needs to do many computations. Thus `\QSx` should be used with either explicit numbers or with items being macros expanding in one step to the numbers (particularly if these numbers are very big).

If the interest is only in \TeX integers, then one should replace the `\xintifCmp` macro with a suitable conditional, possibly helped by tools such as `\ifnumgreater`, `\ifnumequal` and `\ifnumless` from `etoolbox` (\TeX only; I didn't see a direct equivalent to `\xintifCmp`.) Or, if we are dealing with decimal numbers with at most four+four digits, then one should use suitable `\ifdim` tests. Naturally this will boost consequently the speed, from having skipped all the overhead in parsing fractions and scientific numbers as are acceptable by `xintfrac` macros, and subsequent treatment.

³⁹ 2015/11/18 I have not revisited this code for a long time, and perhaps I could improve it now with some new techniques.

⁴⁰ The code in earlier versions of this manual handled inputs composed of braced items. I have switched to comma separated inputs on the occasion of (link removed) The version here is like `code 3` on (link removed) (which is about 3x faster than the earlier code it replaced in this manual) with a modification to make it more efficient if the data has many repeated values. A faster routine (for sorting hundreds of values) is provided as `code 6` at the link mentioned in the footnote, it is based on Merge Sort, but limited to inputs which one can handle as \TeX dimensions. This `code 6` could be extended to handle more general numbers, as acceptable by `xintfrac`. I have also written a non expandable version, which is even faster, but this matters really only when handling hundreds or rather thousands of values.

```

% THE QUICK SORT ALGORITHM EXPANDABLY
% \usepackage{xintfrac} in the preamble (latex)
\makeatletter
% use extra safe delimiters
\catcode`! 3 \catcode`? 3
\def\QSx {\romannumeral0\qsx }%
% first we check if empty list (else \qsx@finish will not find a comma)
\def\qsx #1{\expandafter\qsx@a\romannumeral-`0#1,!,%}
\def\qsx@a #1{\ifx,#1\expandafter\qsx@abort\else
\expandafter\qsx@start\fi #1}%
\def\qsx@abort #1?{ }%
\def\qsx@start {\expandafter\qsx@finish\romannumeral0\qsx@b,%}
\def\qsx@finish ,#1{ #1}%
%
% we check if empty of single and if not pick up the first as Pivot:
\def\qsx@b ,#1#2,#3{\ifx?#3\xintdothis\qsx@empty\fi
\ifx!#3\xintdothis\qsx@single\fi
\xintorthat\qsx@separate {#1#2}{#1#2#3}%
\def\qsx@empty #1#2#3#4#5{ }%
\def\qsx@single #1#2#3#4#5?{, #4}%
\def\qsx@separate #1#2#3#4#5#6,%
{%
\ifx!#5\expandafter\qsx@separate@done\fi
\xintifCmp {#5#6}{#4}%
\qsx@separate@appendtosmaller
\qsx@separate@appendtoequal
\qsx@separate@appendtogreater {#5#6}{#1}{#2}{#3}{#4}%
}%
%
\def\qsx@separate@appendtoequal #1#2{\qsx@separate {#2,#1}}%
\def\qsx@separate@appendtogreater #1#2#3{\qsx@separate {#2}{#3,#1}}%
\def\qsx@separate@appendtosmaller #1#2#3#4{\qsx@separate {#2}{#3}{#4,#1}}%
%
\def\qsx@separate@done\xintifCmp #1%
\qsx@separate@appendtosmaller
\qsx@separate@appendtoequal
\qsx@separate@appendtogreater #2#3#4#5#6#7%
}%
\expandafter\qsx@f\expandafter {\romannumeral0\qsx@b #4,!,%}{\qsx@b #5,!,%}{#3}%
}%
%
\def\qsx@f #1#2#3{#2, #3#1}%
%
\catcode`! 12 \catcode`? 12
\makeatother

% EXAMPLE
\begingroup
\edef\z {\QSx {1.0, 0.5, 0.3, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
1.3, 1.1, 0.7, 1.6, 0.6, 0.9, 0.8, 0.2, 0.1, 1.9}}
\meaning\z

```

```

\def\@ {3.123456789123456789}\def\@b {3.123456789123456788}
\def\@c {3.123456789123456790}\def\@d {3.123456789123456787}
\oodef\@z {\@Sx { \@, \@b, \@c, \@d}}%
% The space before \@a to let it not be expanded during the conversion from CSV
% values to List. The \oodef expands exactly twice (via a bunch of \expandafter's)
\meaning\@z
\endgroup
macro:->0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8,
1.9, 2.0
macro:->\@d , \@b , \@a , \@c (the spaces after \@d, etc... come from the use of the \meaning primitive.)

```

The choice of pivot as first element is bad if the list is already almost sorted. Let's add a variant which will pick up the pivot index randomly. The previous routine worked also internally with comma separated lists, but for a change this one will use internally lists of braced items (the initial conversion via `\xintCSVtoList` handles all potential spurious space problems).

```

% QuickSort expandably on comma separated values with random choice of pivots
% =====> Requires availability of \pdfuniformdeviate <=====
% \usepackage{xintfrac, xinttools} in preamble
\makeatletter
\def\@Sx {\romannumeral0\@sx }% This is a f-expandable macro.
% This converts from comma separated values on input and back on output.
% **** NOTE: these steps (and the other ones too, actually) are costly if input
% has thousands of items.
\def\@sx #1{\xintlistwithsep{, }%
{\expandafter\@sx@sort@a\expandafter{\romannumeral0\xintcsvtolist{#1}}}%
%
% we check if empty or single or double and if not pick up the first as Pivot:
\def\@sx@sort@a #1%
{\expandafter\@sx@sort@b\expandafter{\romannumeral0\xintlength{#1}}{#1}}%
\def\@sx@sort@b #1{\ifcase #1
\expandafter\@sx@sort@empty
\or\expandafter\@sx@sort@single
\or\expandafter\@sx@sort@double
\else\expandafter\@sx@sort@c\fi {#1}}%
\def\@sx@sort@empty #1#2{ }%
\def\@sx@sort@single #1#2{#2}%
\catcode`_ 11
\def\@sx@sort@double #1#2{\xintifGt #2{\xint_exchangetwo_keepbraces}}{#2}%
\catcode`_ 8
\def\@sx@sort@c #1#2{%
\expandafter\@sx@sort@sep@a\expandafter
{\romannumeral0\xintnthelt{\pdfuniformdeviate #1+\@ne}{#2}}{#2?}%
\def\@sx@sort@sep@a #1{\@sx@sort@sep@loop {}{}{#1}}%
\def\@sx@sort@sep@loop #1#2#3#4#5%
{%
\ifx?#5\expandafter\@sx@sort@sep@done\fi
\xintifCmp {#5}{#4}%
\@sx@sort@sep@appendtosmaller
\@sx@sort@sep@appendtoequal
\@sx@sort@sep@appendtogreater {#5}{#1}{#2}{#3}{#4}%
}%
%
%

```

```

\def\qsx@sort@sep@appendtoequal #1#2{\qsx@sort@sep@loop {#2{#1}}}%
\def\qsx@sort@sep@appendtogreater #1#2#3{\qsx@sort@sep@loop {#2}{#3{#1}}}%
\def\qsx@sort@sep@appendtosmaller #1#2#3#4{\qsx@sort@sep@loop {#2}{#3}{#4{#1}}}%
%
\def\qsx@sort@sep@done\xintifCmp #1%
    \qsx@sort@sep@appendtosmaller
    \qsx@sort@sep@appendtoequal
    \qsx@sort@sep@appendtogreater #2#3#4#5#6%
{%
    \expandafter\qsx@sort@recurse\expandafter
        {\romannumeral0\qsx@sort@a {#4}}{\qsx@sort@a {#5}}{#3}%
}%
%
\def\qsx@sort@recurse #1#2#3{#2#3#1}%
%
\makeatother

% EXAMPLES
\begingroup
\edef\z {\QSx {1.0, 0.5, 0.3, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
    1.3, 1.1, 0.7, 1.6, 0.6, 0.9, 0.8, 0.2, 0.1, 1.9}}
\meaning\z

\def\a {3.123456789123456789}\def\b {3.123456789123456788}
\def\c {3.123456789123456790}\def\d {3.123456789123456787}
\oodef\z {\QSx { \a, \b, \c, \d}}%
% The space before \a to let it not be expanded during the conversion from CSV
% values to List. The \oodef expands exactly twice (via a bunch of \expandafter's)
\meaning\z

\def\somenumbers{%
3997.6421, 8809.9358, 1805.4976, 5673.6478, 3179.1328, 1425.4503, 4417.7691,
2166.9040, 9279.7159, 3797.6992, 8057.1926, 2971.9166, 9372.2699, 9128.4052,
1228.0931, 3859.5459, 8561.7670, 2949.6929, 3512.1873, 1698.3952, 5282.9359,
1055.2154, 8760.8428, 7543.6015, 4934.4302, 7526.2729, 6246.0052, 9512.4667,
7423.1124, 5601.8436, 4433.5361, 9970.4849, 1519.3302, 7944.4953, 4910.7662,
3679.1515, 8167.6824, 2644.4325, 8239.4799, 4595.1908, 1560.2458, 6098.9677,
3116.3850, 9130.5298, 3236.2895, 3177.6830, 5373.1193, 5118.4922, 2743.8513,
8008.5975, 4189.2614, 1883.2764, 9090.9641, 2625.5400, 2899.3257, 9157.1094,
8048.4216, 3875.6233, 5684.3375, 8399.4277, 4528.5308, 6926.7729, 6941.6278,
9745.4137, 1875.1205, 2755.0443, 9161.1524, 9491.1593, 8857.3519, 4290.0451,
2382.4218, 3678.2963, 5647.0379, 1528.7301, 2627.8957, 9007.9860, 1988.5417,
2405.1911, 5065.8063, 5856.2141, 8989.8105, 9349.7840, 9970.3013, 8105.4062,
3041.7779, 5058.0480, 8165.0721, 9637.7196, 1795.0894, 7275.3838, 5997.0429,
7562.6481, 8084.0163, 3481.6319, 8078.8512, 2983.7624, 3925.4026, 4931.5812,
1323.1517, 6253.0945}%

\oodef\z {\QSx \somenumbers}% produced as a comma+space separated list
% black magic as workaround to the shrinkability of spaces in last line...
\hsize 87\fontcharwd\font`0
\lccode`~ = 32
\lowercase{\def~}{\discretionary}{~}{\kern\fontcharwd\font`0}}\catcode32 13

```



```

\noindent\phantom{00}\scantokens\expandafter{\meaning\z}\par
\endgroup
macro:->0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8,
1.9, 2.0
macro:->\d , \b , \a , \c
macro:->1055.2154, 1228.0931, 1323.1517, 1425.4503, 1519.3302, 1528.7301, 1560.2458, 1698.3952,
1795.0894, 1805.4976, 1875.1205, 1883.2764, 1988.5417, 2166.9040, 2382.4218, 2405.1911,
2625.5400, 2627.8957, 2644.4325, 2743.8513, 2755.0443, 2899.3257, 2949.6929, 2971.9166,
2983.7624, 3041.7779, 3116.3850, 3177.6830, 3179.1328, 3236.2895, 3481.6319, 3512.1873,
3678.2963, 3679.1515, 3797.6992, 3859.5459, 3875.6233, 3925.4026, 3997.6421, 4189.2614,
4290.0451, 4417.7691, 4433.5361, 4528.5308, 4595.1908, 4910.7662, 4931.5812, 4934.4302,
5058.0480, 5065.8063, 5118.4922, 5282.9359, 5373.1193, 5601.8436, 5647.0379, 5673.6478,
5684.3375, 5856.2141, 5997.0429, 6098.9677, 6246.0052, 6253.0945, 6926.7729, 6941.6278,
7275.3838, 7423.1124, 7526.2729, 7543.6015, 7562.6481, 7944.4953, 8008.5975, 8048.4216,
8057.1926, 8078.8512, 8084.0163, 8105.4062, 8165.0721, 8167.6824, 8239.4799, 8399.4277,
8561.7670, 8760.8428, 8809.9358, 8857.3519, 8989.8105, 9007.9860, 9090.9641, 9128.4052,
9130.5298, 9157.1094, 9161.1524, 9279.7159, 9349.7840, 9372.2699, 9491.1593, 9512.4667,
9637.7196, 9745.4137, 9970.3013, 9970.4849

```

All the previous examples were with numbers which could have been handled via `\ifdim` tests rather than the `\xintifCmp` macro from `xintfrac`; using `\ifdim` tests would naturally be faster. Even faster routine is `code 6` at (link removed) which uses `\pdfescapestring` and a Merge Sort algorithm.

We then turn to a graphical illustration of the algorithm.⁴¹ For simplicity the pivot is always chosen as the first list item. Then we also give a variant which picks up the last item as pivot.

```

% in LaTeX preamble:
% \usepackage{xintfrac, xinttools}
% \usepackage{color}
% or, when using Plain TeX:
% \input xintfrac.sty \input xinttools.sty
% \input color.tex
%
% Color definitions
\definecolor{LEFT}{RGB}{216,195,88}
\definecolor{RIGHT}{RGB}{208,231,153}
\definecolor{INERT}{RGB}{199,200,194}
\definecolor{INERTpivot}{RGB}{237,237,237}
\definecolor{PIVOT}{RGB}{109,8,57}
% Start of macro defintions
\makeatletter
% \catcode`? 3 % a bit too paranoid. Normal ? will do.
%
% argument will never be empty
\def\QS@cmp@a #1{\QS@cmp@b #1??}%
\def\QS@cmp@b #1{\noexpand\QS@sep@A\@ne{#1}\QS@cmp@d {#1}}%
\def\QS@cmp@d #1#2{\ifx ?#2\expandafter\QS@cmp@done\fi
\xintifCmp {#1}{#2}\tw@\@ne\z@{#2}\QS@cmp@d {#1}}%
\def\QS@cmp@done #1{??}%
%
\def\QS@sep@A #1?{\QSLr\QS@sep@L #1\thr@@?#1\thr@@?#1\thr@@?}%
\def\QS@sep@L #1#2{\ifcase #1{#2}\or\or\else

```

⁴¹ I have rewritten (2015/11/21) the routine to do only once (and not thrice) the needed calls to `\xintifCmp`, up to the price of one additional `\edef`, although due to the context execution time on our side is not an issue and moreover is anyhow overwhelmed by the TikZ's activities. Simultaneously I have updated the code. The variant with the choice of pivot on the right has more overhead: the reason is simply that we do not convert the data into an array, but maintain a list of tokens with self-reorganizing delimiters.

```

\expandafter\QS@sep@I@start\fi \QS@sep@L}%
\def\QS@sep@I@start\QS@sep@L {\noexpand\empty?\QSIr\QS@sep@I}%
\def\QS@sep@I #1#2{\ifcase#1\or{#2}\or\else\expandafter\QS@sep@R@start\fi\QS@sep@I}%
\def\QS@sep@R@start\QS@sep@I {\noexpand\empty?\QSRr\QS@sep@R}%
\def\QS@sep@R #1#2{\ifcase#1\or\or{#2}\else\expandafter\QS@sep@done\fi\QS@sep@R}%
\def\QS@sep@done\QS@sep@R {\noexpand\empty?}%
%
\def\QS@loop {%
  \xintloop
  % pivot phase
  \def\QS@pivotcount{0}%
  \let\QSLr\DecoLEFTwithPivot \let\QSIr \DecoINERT
  \let\QSRr\DecoRIGHTwithPivot \let\QSIrr\DecoINERT
  \centerline{\QS@list}%
  % sorting phase
  \ifnum\QS@pivotcount>\z@
    \def\QSLr {\QS@cmp@a}\def\QSRr {\QS@cmp@a}%
    \def\QSIr {\QSIrr}\let\QSIrr\relax
    \edef\QS@list{\QS@list}% compare
    \let\QSLr\relax\let\QSRr\relax\let\QSIr\relax
    \edef\QS@list{\QS@list}% separate
    \def\QSLr ##1##2?{\ifx\empty##1\else\noexpand \QSLr {{{##1}##2}\fi}%
    \def\QSIr ##1##2?{\ifx\empty##1\else\noexpand \QSIr {{{##1}##2}\fi}%
    \def\QSRr ##1##2?{\ifx\empty##1\else\noexpand \QSRr {{{##1}##2}\fi}%
    \edef\QS@list{\QS@list}% gather
    \let\QSLr\DecoLEFT \let\QSRr\DecoRIGHT
    \let\QSIr\DecoINERTwithPivot \let\QSIrr\DecoINERT
    \centerline{\QS@list}%
  \repeat }%
%
% \xintFor* loops handle gracefully empty lists.
\def\DecoLEFT #1{\xintFor* ##1 in {#1} \do {\colorbox{LEFT}{##1}}}%
\def\DecoINERT #1{\xintFor* ##1 in {#1} \do {\colorbox{INERT}{##1}}}%
\def\DecoRIGHT #1{\xintFor* ##1 in {#1} \do {\colorbox{RIGHT}{##1}}}%
\def\DecoPivot #1{\begingroup
  \color{PIVOT}\advance\fbboxsep-\fbboxrule\fbbox{#1}\endgroup}%
%
\def\DecoLEFTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForFirst {\DecoPivot {##1}}{\colorbox{LEFT}{##1}}}%
\def\DecoINERTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForFirst {\colorbox{INERTpiv}{##1}}{\colorbox{INERT}{##1}}}%
\def\DecoRIGHTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForFirst {\DecoPivot {##1}}{\colorbox{RIGHT}{##1}}}%
%
\def\QuickSort #1{% warning: not compatible with empty #1.
  % initialize, doing conversion from comma separated values
  % to a list of braced items
  \edef\QS@list{\noexpand\QSRr{\xintCSVtoList{#1}}}%
  % may \edef's are to follow anyhow

```

```
% earlier I did a first drawing of the list, here with the color of RIGHT elements,
% but the color should have been for example white, anyway I drop this first line
\let\QSRr\DecoRIGHT
%\par\centerline{\QS@list}%
%
% loop as many times as needed
\QS@loop }%
%
% \catcode`? 12 % in case we had used a funny ? as delimiter.
\makeatother
%% End of macro definitions.
%% Start of Example
\begingroup\offinterlineskip
\small
% \QuickSort {1.0, 0.5, 0.3, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
%             1.3, 1.1, 0.7, 1.6, 0.6, 0.9, 0.8, 0.2, 0.1, 1.9}
% \medskip
% with repeated values
\QuickSort {1.0, 0.5, 0.3, 0.8, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
            1.3, 1.1, 0.7, 0.3, 1.6, 0.6, 0.3, 0.8, 0.2, 0.8, 0.7, 1.2}
\endgroup
```

1.0	0.5	0.3	0.8	1.5	1.8	2.0	1.7	0.4	1.2	1.4	1.3	1.1	0.7	0.3	1.6	0.6	0.3	0.8	0.2	0.8	0.7	1.2
0.5	0.3	0.8	0.4	0.7	0.3	0.6	0.3	0.8	0.2	0.8	0.7	1.0	1.5	1.8	2.0	1.7	1.2	1.4	1.3	1.1	1.6	1.2
0.5	0.3	0.8	0.4	0.7	0.3	0.6	0.3	0.8	0.2	0.8	0.7	1.0	1.5	1.8	2.0	1.7	1.2	1.4	1.3	1.1	1.6	1.2
0.3	0.4	0.3	0.3	0.2	0.5	0.8	0.7	0.6	0.8	0.8	0.7	1.0	1.2	1.4	1.3	1.1	1.2	1.5	1.8	2.0	1.7	1.6
0.3	0.4	0.3	0.3	0.2	0.5	0.8	0.7	0.6	0.8	0.8	0.7	1.0	1.2	1.4	1.3	1.1	1.2	1.5	1.8	2.0	1.7	1.6
0.2	0.3	0.3	0.3	0.4	0.5	0.7	0.6	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.4	1.3	1.5	1.7	1.6	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.7	0.6	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.4	1.3	1.5	1.7	1.6	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0

Here is the variant which always picks the pivot as the rightmost element.

```
\makeatletter
%
\def\QS@cmp@a #1{\noexpand\QS@sep@a\expandafter\QS@cmp@d\expandafter
                {\romannumeral0\xintnthelt{-1}{#1}{#1??}}%
%
\def\DecoLEFTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForLast {\DecoPivot {##1}}{\colorbox{LEFT}{##1}}}}
\def\DecoINERTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForLast {\colorbox{INERTpiv}{##1}}{\colorbox{INERT}{##1}}}}
\def\DecoRIGHTwithPivot #1{\xdef\QS@pivotcount{\the\numexpr\QS@pivotcount+\@ne}%
  \xintFor* ##1 in {#1} \do
    {\xintifForLast {\DecoPivot {##1}}{\colorbox{RIGHT}{##1}}}}
\def\QuickSort #1{%
  % initialize, doing conversion from comma separated values
  % to a list of braced items
  \edef\QS@list{\noexpand\QSLr {\xintCSVtoList{#1}}}%
  % many \edef's are to follow anyhow
  %
```

```

% loop as many times as needed
\QS@loop }%
\makeatother
\begingroup\offinterlineskip
\small
% \QuickSort {1.0, 0.5, 0.3, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
%             1.3, 1.1, 0.7, 1.6, 0.6, 0.9, 0.8, 0.2, 0.1, 1.9}
% \medskip
% with repeated values
\QuickSort {1.0, 0.5, 0.3, 0.8, 1.5, 1.8, 2.0, 1.7, 0.4, 1.2, 1.4,
            1.3, 1.1, 0.7, 0.3, 1.6, 0.6, 0.3, 0.8, 0.2, 0.8, 0.7, 1.2}
\endgroup

```

1.0	0.5	0.3	0.8	1.5	1.8	2.0	1.7	0.4	1.2	1.4	1.3	1.1	0.7	0.3	1.6	0.6	0.3	0.8	0.2	0.8	0.7	1.2
1.0	0.5	0.3	0.8	0.4	1.1	0.7	0.3	0.6	0.3	0.8	0.2	0.8	0.7	1.2	1.2	1.5	1.8	2.0	1.7	1.4	1.3	1.6
1.0	0.5	0.3	0.8	0.4	1.1	0.7	0.3	0.6	0.3	0.8	0.2	0.8	0.7	1.2	1.2	1.5	1.8	2.0	1.7	1.4	1.3	1.6
0.5	0.3	0.4	0.3	0.6	0.3	0.2	0.7	0.7	1.0	0.8	1.1	0.8	0.8	1.2	1.2	1.5	1.4	1.3	1.6	1.8	2.0	1.7
0.5	0.3	0.4	0.3	0.6	0.3	0.2	0.7	0.7	1.0	0.8	1.1	0.8	0.8	1.2	1.2	1.5	1.4	1.3	1.6	1.8	2.0	1.7
0.2	0.5	0.3	0.4	0.3	0.6	0.3	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.5	1.4	1.6	1.7	1.8	2.0
0.2	0.5	0.3	0.4	0.3	0.6	0.3	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.5	1.4	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.5	0.4	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.5	0.4	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.5	0.4	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.5	0.4	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0
0.2	0.3	0.3	0.3	0.4	0.5	0.6	0.7	0.7	0.8	0.8	0.8	1.0	1.1	1.2	1.2	1.3	1.4	1.5	1.6	1.7	1.8	2.0

The choice of the first or last item as pivot is not a good one as nearly ordered lists will take quadratic time. But for explaining the algorithm via a graphical interpretation, it is not that bad. If one wanted to pick up the pivot randomly, the routine would have to be substantially rewritten: in particular the `\Deco.withPivot` macros need to know where the pivot is, and currently this is implemented by using either `\xintifForFirst` or `\xintifForLast`.

Part II.

The macro layer for expandable computations: [xintcore](#), [xint](#), [xintfrac](#), and some extras

WARNING !
 The documentation is getting old, and is in need of rewrites for many sections, particularly for examples.
 We do try to keep updated the description of macros provided by the packages.

8	The xint bundle	125
9	Macros of the xintkernel package	140
10	Macros of the xintcore package	146
11	Macros of the xint package	151
12	Macros of the xintfrac package	163
13	Macros of the xintbinhex package	193
14	Macros of the xintgcd package	196
15	Macros of the xintseries package	198
16	Macros of the xintcfrac package	214

8. The [xint](#) bundle

.1	Characteristics	125		.8	\ifcase, \ifnum, ... constructs	134
.2	Floating point evaluations	127	.9	No variable declarations are needed	135	
.3	Expansion matters	128	.10	Possible syntax errors to avoid	135	
.4	Input formats for macros	130	.11	Error messages	136	
.5	Output formats of macros	132	.12	Package namespace, catcodes	137	
.6	Count registers and variables	132	.13	Origins of the package	138	
.7	Dimension registers and variables	133				

8.1. Characteristics

The main characteristics are:

1. exact algebra on ``big numbers'', integers as well as fractions,
2. floating point variants with user-chosen precision,
3. the computational macros are compatible with expansion-only context,
4. the bundle comes with parsers (integer-only, or handling fractions, or doing floating point computations) of infix operations implementing beyond infix operations extra features such as dummy variables.

Since 1.2 ``big numbers'' must have less than about 19950 digits: the maximal number of digits for addition is at 19968 digits, and it is 19959 for multiplication. The reasonable range of use of the package is with numbers of up to a few hundred digits.⁴²

\TeX does not know off-hand how to print on the page such very long numbers, see [subsection 1.6](#).

Integers with only 10 digits and starting with a 3 already exceed the \TeX bound; and \TeX does not have a native processing of floating point numbers (multiplication by a decimal number of a dimension register is allowed --- this is used for example by the [pgf](#) basic math engine.)

\TeX elementary operations on numbers are done via the non-expandable `\advance`, `\multiply`, and `\divide` assignments. This was changed with ϵ - \TeX 's `\numexpr` which does expandable computations using standard infix notations with \TeX integers. But ϵ - \TeX did not modify the \TeX bound on acceptable integers, and did not add floating point support.

The [bigintcalc](#) package by HEIKO OBERDIEK provided expandable macros (using some of `\numexpr` possibilities, when available) on arbitrarily big integers, beyond the \TeX bound. It does not provide an expression parser.⁴³ `xint` did it again using more of `\numexpr` for higher speed, and in a later evolution added handling of exact fractions, of scientific numbers, and an expression parser. Arbitrary precision floating points operations were added as a derivative, and not part of the initial design goal.

The concept of signed infinities, signed zeroes, NaN's, error traps...⁴⁴ have not been implemented, only the notion of `scientific notation with a given number of significant figures'.⁴⁵

The \mathbb{X}_3 project has implemented expandably floating-point computations with 16 significant figures ([l3fp](#)), including functions such as exp, log, sine and cosine.⁴⁶

More directly related to the `xint` bundle there is the [l3bigint](#) package, also devoted to big integers and in development a.t.t.o.w (2015/10/09, no division yet). It is part of the experimental trunk of the [\$\mathbb{X}_3\$ Project](#) and provides an expression parser for expandable arithmetic with big integers. Its author Bruno LE FLOCH succeeded brilliantly into implementing expandably the Karatsuba multiplication algorithm and he achieves *sub-quadratic growth for the computation time*. This shows up very clearly with numbers having thousands of digits, up to the maximum which a.t.t.o.w is at 8192 digits.

The [l3bigint](#) multiplication from late 2015 is observed to be roughly $3x$ -- $4x$ faster than the one from `\xintiexpr` in the range of 4000 to 5000 digits integers, and isn't far from being $9x$ faster

⁴² For example multiplication of integers having from 50 to 100 digits takes roughly of the order of the millisecond on a 2012 desktop computer. I compared this to using Python3: using `timeit` module on a wrapper defined as `return w*z` with random integers of 100 digits, I observe on the same computer a computation time of roughly $4 \cdot 10^{-7}$ s per call. And with `return str(w*z)` then this becomes more like $16 \cdot 10^{-7}$ s per call. And with `return str(int(W)*int(Z))` where W and Z are strings, this becomes about $26 \cdot 10^{-7}$ s (I am deliberately ignoring Python's Decimal module here...) Anyway, my sentence from earlier version of this documentation: *this is, I guess, at least about 1000 times slower than what can be expected with any reasonable programming language*, is about right. I then added: *nevertheless as compilation of a typical \LaTeX document already takes of the order of seconds and even dozens of seconds for long ones, this leaves room for reasonably many computations via `xintexpr` or via direct use of the macros of `xint/xintfrac`.*

⁴³ One can currently use package `bnumexpr` to associate the `bigintcalc` macros with an expression parser. This may be unavailable in future if `bnumexpr` becomes more tightly associated with future evolutions or variants of `xintcore`.
⁴⁴ The latter exist as work-in-progress for some time in the source code.
⁴⁵ Multiplication of two floats with `P=\xinttheDigits` digits is first done exactly then rounded to P digits, rather than using a specially tailored multiplication for floating point numbers which would be more efficient (it is a waste to evaluate fully the multiplication result with $2P$ or $2P-1$ digits.)
⁴⁶ at the time of writing (2014/10/28) the [l3fp](#) (exactly represented) floating point numbers have their exponents limited to ± 9999 .

at 8000 digits. On the other hand `\xintiexpr`'s multiplication is found to be on average roughly 2.5x faster than `l3bigint`'s for numbers up to 100 digits and the two packages achieve about the same speed at 900 digits: but each such multiplication of numbers of 900 digits costs about one or two tenths of a second on a 2012 desktop computer, whereas the order of magnitude is rather the ms for numbers with 50--100 digits.⁴⁷

Even with the superior `l3bigint` Karatsuba multiplication it takes about 3.5s on this 2012 desktop computer for a single multiplication of two 5000-digits numbers. Hence it is not possible to do routinely such computations in a document. I have long been thinking that without the expandability constraint much higher speeds could be achieved, but perhaps I have not given enough thought to sustain that optimistic stance.⁴⁸

I remain of the opinion that if one really wants to do computations with *thousands* of digits, one should drop the expandability requirement. Indeed, as clearly demonstrated long ago by the [pi computing file](#) by D. ROEGEL one can program \TeX to compute with many digits at a much higher speed than what `xint` achieves: but, direct access to memory storage in one form or another seems a necessity for this kind of speed and one has to renounce at the complete expandability.⁴⁹

8.2. Floating point evaluations

Floating point macros are provided by package `xintfrac` to work with a given arbitrary precision `P`. The default value is `P = 16` meaning that the significands of the produced (non-zero) numbers have 16 decimal digits. The syntax to set the precision to `P` is

```
\xintDigits:=P\relax
```

The value is local to the group or environment (if using \LaTeX). To query the current value use `\xinttheDigits`.

Most floating point macros accept an optional first argument `[P]` which then sets the target precision and replaces the `\xintDigits` assigned value (the `[P]` must be repeated if the arguments are themselves `xintfrac` macros with arguments of their own.) In this section `P` refers to the prevailing `\xinttheDigits` float precision or to the target precision set in this way as an optional argument.

`\xintfloatexpr[Q]...\relax` also admits an optional argument `[Q]` but it has an altogether different meaning: the computations are always done with the prevailing `\xinttheDigits` precision and the optional argument `Q` is used for the final rounding. This makes sense only if `Q < \xinttheDigits` and is intended to clean up the result from dubious last digits (when `Q < 0` it indicates rather by how many digits one should reduce the mantissa lengths via a final rounding).

The IEEE 754⁵⁰ requirement of *correct rounding* for addition, subtraction, multiplication, division and square root is achieved (in arbitrary precision) by the macros of `xintfrac` hence also by the infix operators `+`, `-`, `*`, `/`.

This means that for operands given with at most `P` significant digits (and arbitrary exponents) the output coincides exactly with the rounding of the exact theoretical result (barring overflow or underflow).

Due to a typographical oversight, this documentation (up to 1.2j) adjoined `^` and `**` to the above list of infix operators. But as is explained in [subsection 12.97](#), what is guaranteed regarding integer powers is an error of at most 0.52ulp, not the correct rounding. Half-integer powers are computed as square roots of integer powers.

The rounding mode is ``round to nearest, ties away from zero''. It is not customizable.

Currently `xintfrac` has no notion of NaNs or signed infinities or signed zeroes, but this is intended for the future.

⁴⁷ I have tested this again on 2016/12/19, but the macros have not changed on the `l3bigint` side and barely on the `xintcore` side, hence I got again the same results. . . ⁴⁸ The `apnum` package implements (non-expandably) arbitrary precision fixed point algebra and (v1.6) functions `exp`, `log`, `sqrt`, the trigonometrical direct and inverse functions. ⁴⁹ The \LaTeX project possibly makes endeavours such as `xint` appear even more insane that they are, in truth: `xint` is able to handle fast enough computations involving numbers with less than one hundred digits and brings this to all engines.

Since release 1.2f, square root extraction achieves correct rounding in arbitrary precision.

See [xintlog](#) for fractional powers and [xinttrig](#) for trigonometrical functions.

The maximal floating point decimal exponent is currently 2147483647 which is the maximal number handled by \TeX . The minimal exponent is its opposite. But this means that overflow or underflow are detected only via low-level `\numexpr` arithmetic overflows which are basically un-recoverable. Besides there are some border effects as the routines need to add or subtract lengths of numbers from exponents, possibly triggering the low-level overflows. In the future not only the Precision but also the maximal and minimal exponents `Emin` and `Emax` will be specifiable by the user.

Since 1.2f, the float macros round their inputs to the target precision `P` before further processing. Formerly, the initial rounding was done to `P+2` digits (and at least `P+3` for the power operation.)

The more ambitious model would be for the computing macros to obey the intrinsic precision of their inputs, i.e. to compute the correct rounding to `P` digits of the exact mathematical result corresponding to inputs allowed to have their own higher precision.⁵¹ This would be feasible by [xintfrac](#) which after all knows how to compute exactly, but I have for the time being decided that for reasons of efficiency, the chosen model is the one of rounding inputs to the target precision first.

The float macros of [xintfrac](#) have to handle inputs which not only may have much more digits than the target float precision, but may even be fractions: in a way this means infinite precision.

From releases 1.08a to 1.2j a fraction input `AeM/BeN` had its numerator and denominator `A` and `B` truncated to `Q+2` digits of precision, then the substituted fraction was correctly rounded to `Q` digits of precision (usually with `Q` set to `P+2`) and then the operation was implemented on such rounded inputs. But this meant that two fractions representing the same rational number could end up being rounded differently (with a difference of one unit in the last place), if it had numerators and denominators with at least `Q+3` digits.

Starting with release 1.2k a fractional input `AeM/BeN` is handled intrinsically: the fraction, independently of its representation `AeM/BeN`, is *correctly rounded* to `P` digits during the input parsing. Hence the output depends only on its arguments as mathematical fractions and not on their representatives as quotients.

Notice that in float expressions, the `/` is treated as operator, and is applied to arguments which are generally already `P`-floats, hence the above discussion becomes relevant in this context only for the special input form `qfloat(A/B)` or when using a sub-expression `\xintexpr A/B\relax` embedded in the float expression with `A` or `B` having more digits than the prevailing float precision `P`.

8.3. Expansion matters

8.3.1. Full expansion of the first token

The whole business of `xint` is to build upon `\numexpr` and handle arbitrarily large numbers. Each basic operation is thus done via a macro: `\xintiiAdd`, `\xintiiSub`, `\xintiiMul`, `\xintiiDivision`. In order to handle more complex operations, it must be possible to nest these macros. An expandable macro can not execute a `\def` or an `\edef`. But the macro must expand its arguments to find the digits it is supposed to manipulate. \TeX provides a tool to do the job of (expandable !) repeated expansion of the first token found until hitting something non expandable, such as a digit, a `\def` token, a brace, a `\count` token, etc... is found. A space token also will stop the expansion (and be swallowed, contrarily to the non-expandable tokens).

By convention in this manual *f-expansion* ('`full expansion' or '`full first expansion') will be this \TeX process of expanding repeatedly the first token seen. For those familiar with \TeX 3

⁵⁰ The IEEE 754-1985 standard was for hardware implementations of binary floating-point arithmetic with a specific value for the precision (24 bits for single precision, 53 bits for double precision). The newer IEEE 754-2008 (https://en.wikipedia.org/wiki/IEEE_floating_point) normalizes five basic formats, three binaries and two decimals (16 and 34 decimal digits) and discusses extended formats with higher precision. These standards are only indirectly relevant to libraries like `xint` dealing with arbitrary precision. ⁵¹ The MPFR library <http://www.mpfr.org/> implements this but it does not know fractions!

(which is not used by `xint`) this is what is called in its documentation full expansion, whereas expansion inside `\edef` would be described I think as ``exhaustive'' expansion and will be referred too in this manual as *x-expansion*.

Most of the package macros, and all those dealing with computations⁵², are expandable in the strong sense that they expand to their final result via this *f-expansion*. This will be signaled in their descriptions via a star in the margin.

★ These macros not only have this property of *f-expandability*, they all begin by first applying *f-expansion* to their arguments. Again from \TeX 's conventions this will be signaled by a margin annotation next to the description of the arguments.

8.3.2. Summary of important expandability aspects

1. the macros *f-expand* their arguments, this means that they expand the first token seen (for each argument), then expand, etc..., until something un-expandable such as a digit or a brace is hit against. This example

```
\def\x{98765}\def\y{43210} \xintiiAdd {\x}{\x\y}
```

is *not* a legal construct, as the `\y` will remain untouched by expansion and not get converted into the digits which are expected by the sub-routines of `\xintiiAdd`. It is a `\numexpr` which will expand it and an arithmetic overflow will arise as `9876543210` exceeds the \TeX bounds. The same would hold for `\xintAdd`.

To the contrary `\xinttheiexpr` and others have no issues with things such as `\xinttheiexpr \x+\x\y\relax`.

2. using `\if...\fi` constructs *inside* the package macro arguments requires suitably mastering \TeX niques (`\expandafter`'s and/or swapping techniques) to ensure that the *f-expansion* will indeed absorb the `\else` or closing `\fi`, else some error will arise in further processing. Therefore it is highly recommended to use the package provided conditionals such as `\xintifEq`, `\xintifGt`, `\xintifSgn`,... or, for \TeX users and when dealing with short integers the `etoolbox`⁵³ expandable conditionals (for small integers only) such as `\ifnumequal`, `\ifnumgreater`,... Use of *non-expandable* things such as `\ifthenelse` is impossible inside the arguments of `xint` macros.

One can use naive `\if...\fi` things inside an `\xinttheexpr`-ession and cousins, as long as the test is expandable, for example

```
\xinttheiexpr\ifnum3>2 143\else 33\fi 0^2\relax→2044900=1430^2
```

3. after the definition `\def\x {12}`, one can not use `-\x` as input to one of the package macros: the *f-expansion* will act only on the minus sign, hence do nothing. The only way is to use the `\xintOpp` macro (or `\xintiiOpp` which is integer only) which obtains the opposite of a given number.

Again, this is otherwise inside an `\xinttheexpr`-ession or `\xintthefloatexpr`-ession. There, the minus sign may prefix macros which will expand to numbers (or parentheses etc...)

4. With the definition

```
\def\AplusBC #1#2#3{\xintAdd {#1}{\xintMul {#2}{#3}}}
```

one obtains an expandable macro producing the expected result, not in two, but rather in three steps: a first expansion is consumed by the macro expanding to its definition. As the package macros expand their arguments until no more is possible (regarding what comes first), this `\AplusBC` may be used inside them: `\xintAdd {\AplusBC {1}{2}{3}}{4}` does work and returns `11/1[0]`.

If, for some reason, it is important to create a macro expanding in two steps to its final value, one may either do:

⁵² except `\xintXTrunc`. ⁵³ <https://ctan.org/pkg/etoolbox>

```
\def\AplusBC #1#2#3{\romannumeral-`0\xintAdd {#1}{\xintMul {#2}{#3}}}
```

or use the *lowercase* form of `\xintAdd`:

```
\def\AplusBC #1#2#3{\romannumeral0\xintadd {#1}{\xintMul {#2}{#3}}}
```

and then `\AplusBC` will share the same properties as do the other `xint` 'primitive' macros.

5. The `\romannumeral0` and `\romannumeral-`0` things above look like an invitation to hacker's territory; if it is not important that the macro expands in two steps only, there is no reason to follow these guidelines. Just chain arbitrarily the package macros, and the new ones will be completely expandable and usable one within the other.

Since release 1.07 the `\xintNewExpr` macro automatizes the creation of such expandable macros:

```
\xintNewExpr\AplusBC[3]{#1+#2*#3}
```

creates the `\AplusBC` macro doing the above and expanding in two expansion steps.

6. In the expression parsers of `xintexpr` such as `\xintexpr.\relax`, `\xintfloatexpr.\relax` the contents are expanded completely from left to right until the ending `\relax` is found and swallowed, and spaces and even (to some extent) catcodes do not matter.
7. For all variants, prefixing with `\xintthe` allows to print the result or use it in other contexts. Shortcuts `\xinttheexpr`, `\xintthefloatexpr`, `\xinttheiexpr`, ... are available.

8.4. Input formats for macros

Macros can have different types of arguments (we do not consider here the `\xintexpr`-parsers but only the macros of `xintcore/xint/xintfrac`). In a macro description, a margin annotation signals what is the argument type.

`num`
`x`

1. `TeX` integers are handled inside a `\numexpr.\relax` hence may be count registers or variables. Beware that `-(1+1)` is not legal and raises an error, but `0-(1+1)` is. Also `2\cnta` with `\cnta a \count` isn't legal. Integers must be kept less than 2147483647 in absolute value, although the scaling operation `(a*b)/c` computes the intermediate product with twice as many bits.

The slash `/` does a `rounded` division which is a fact of life of `\numexpr` which I have found very annoying in at least nine cases out of ten, not to say ninety-nine cases out of one hundred. Besides, it is at odds with `TeX`'s `\divide` which does a truncated division (non-expandably).

But to follow-suit `/` also does rounded integer division in `\xintiexpr.\relax`, and the operator `//` does there the truncated division.

`f`

2. the strict format applies to macros handling big integers but only `f`-expanding their arguments. After this `f`-expansion the input should be a string of digits, optionally preceded by a unique minus sign. The first digit can be zero only if it is the only digit. A plus sign is not accepted. `-0` is not legal in the strict format. Macros of `xint` with a double `ii` require this 'strict' format for the inputs.

`Num`
`f`

3. the extended integer format applies when the macro parses its arguments via `\xintNum`. The input may then have arbitrarily many leading minus and plus signs, followed by leading zeroes, and further digits. With `xintfrac` loaded, `\xintNum` is extended to accept fractions and its action is to truncate them to integers.

`Frac`
`f`

4. the fraction input format applies to the arguments of `xintfrac` macros handling genuine fractions. It allows two types of inputs: general and restricted. The restricted type is parsed faster, but... is restricted.

general: inputs of the shape `A.BeC/D.EeF`. Example:

```
\noindent\xintRaw{+--0367.8920280e17/--+278.289287e-15}\newline
\xintRaw{+---1253.2782e+++3/---0087.123e---5}\par
```

TOC, `\xint bundle`, `\xintkernel`, `\xintcore`, `\xint`, `\xintfrac`, `\xintbinhex`, `\xintgcd`, `\xintseries`, `\xintcfrac`

`-3678920280/278289287[31]`

`-12532782/87123[7]`

The input parser does not reduce fractions to smallest terms. Here are the rules of this general fraction format:

- everything is optional, absent numbers are treated as zero, here are some extreme cases:

```
\xintRaw{,} \xintRaw{.} \xintRaw{./1.e} \xintRaw{-.e} \xintRaw{e/-1}
0/1[0], 0/1[0], 0/1[0], 0/1[0], 0/1[0]
```

- **AB** and **DE** may start with pluses and minuses, then leading zeroes, then digits.
- **C** and **F** will be given to `\numexpr` and can be anything recognized as such and not provoking arithmetic overflow (the lengths of **B** and **E** will also intervene to build the final exponent naturally which must obey the \TeX bound).
- the `/`, `.` (numerator and/or denominator) and `e` (numerator and/or denominator) are all optional components.
- each of **A**, **B**, **C**, **D**, **E** and **F** may arise from *f-expansion* of a macro.
- the whole thing may arise from *f-expansion*, however the `/`, `.`, and `e` should all come from this initial expansion. The `e` of scientific notation is mandatorily lowercased.

restricted: inputs either of the shape `A[N]` or `A/B[N]`, which represents the fraction `A/B` times 10^N . The whole thing or each of **A**, **B**, **N** (but then not `/` or `[]`) may arise from *f-expansion*, **A** (after expansion) *must* have a unique optional minus sign and no leading zeroes, **B** (after expansion) if present *must* be a positive integer with no signs and no leading zeroes, `[N]` if present will be given to `\numexpr`. Any deviation from the rules above will result in errors.

Frac
f

Notice that `*`, `+` and `-` contrarily to the `/` (which is treated simply as a kind of delimiter) are not acceptable within arguments of this type (see [subsection 8.6](#) for some exceptions to this.)

Generally speaking, there should be no spaces among the digits in the inputs (in arguments to the package macros). Although most would be harmless in most macros, there are some cases where spaces could break havoc.⁵⁴ So the best is to avoid them entirely.

This is entirely otherwise inside an `\xintexpr`-ession, where spaces are ignored (except when they occur inside arguments to some macros, thus escaping the `\xintexpr` parser). See the [section 2](#).

There are also some slightly more obscure expansion types: in particular, the `\xintApplyInline` and `\xintFor*` macros from `\xinttools` apply a special iterated *f-expansion*, which gobbles spaces, to the non-braced items (braced items are submitted to no expansion because the opening brace stops it) coming from their list argument; this is denoted by a special symbol in the margin. Some other macros such as `\xintSum` from `\xintfrac` first do an *f-expansion*, then treat each found (braced or not) item (skipping spaces between such items) via the general fraction input parsing, this is signaled as here in the margin where the signification of the `*` is thus a bit different from the previous case.

*f

Frac
f → * f

`n`, resp. `o` A few macros from `\xinttools` do not expand, or expand only once their argument. This is also signaled in the margin with notations à la \LaTeX 3.

⁵⁴ The `\xintNum` macro does not remove spaces between digits beyond the first non zero ones; however this should not really alter the subsequent functioning of the arithmetic macros, and besides, since `\xintcore` 1.2 there is an initial parsing of the entire number, during which spaces will be gobbled. However I have not done a complete review of the legacy code to be certain of all possibilities after 1.2 release. One thing to be aware of is that `\numexpr` stops on spaces between digits (although it provokes an expansion to see if an infix operator follows); the exponent for `\xintiiPow` or the argument of the factorial `\xintiiFac` are only subjected to such a `\numexpr` (there are a few other macros with such input types in `\xint`). If the input is given as, say `1 2\x` where `\x` is a macro, the macro `\x` will not be expanded by the `\numexpr`, and this will surely cause problems afterwards. Perhaps a later `\xint` will force `\numexpr` to expand beyond spaces, but I decided that that was not really worth the effort. Another immediate cause of problems is an input of the type `\xintiiAdd {<space>\x }{\y }`, because the space will stop the initial expansion; this will most certainly cause an arithmetic overflow later when the `\x` will be expanded in a `\numexpr`. Thus in conclusion, damages due to spaces are unlikely if only explicit digits are involved in the inputs, or arguments are single macros with no preceding space.

8.5. Output formats of macros

We do not consider here the `\xintexpr`-parsers but only the macros from `xintcore`, `xint` and `xintfrac`. Macros of other components of the bundle may have their own output formats, for example for continuous fractions with `xintfrac`. There are mainly three types of outputs:

- arithmetic macros from `xintcore/xint` deliver integers in the strict format as described in the previous section.
- arithmetic macros from `xintfrac` produce on output the strict fraction format $A/B[N]$, which stands for $(A/B) \times 10^N$, where A and B are integers, B is positive, and N is a 'short' integer. The output is not reduced to smallest terms. The A and B may end with zeroes (i.e. N does not represent all powers of ten). The denominator B is always strictly positive. There is no $+$ sign. The $-$ is always first if present (i.e. the denominator on output is always positive.) The output will be expressed as such a fraction even if the inputs are both integers and the mathematical result is an integer. The $B=1$ is not removed.⁵⁵
- macros from `xintfrac` having `Float` in their names deliver a number in the scientific notation as described in the documentation of `\xintFloat`.

The exception is `\xintPFloat` which does some customizable pretty printing of the result.

8.6. Count registers and variables

Inside `\xintexpr..\relax` and its variants, a count register or count control sequence is automatically unpacked using `\number`, with tacit multiplication: `1.23\counta` is like `1.23*\number\counta`. There is a subtle difference between count *registers* and count *variables*. In `1.23*\counta` the unpacked `\counta` variable defines a complete operand thus `1.23*\counta 7` is a syntax error. But `1.23*\count0` just replaces `\count0` by `\number\count0` hence `1.23*\count0 7` is like `1.23*57` if `\count0` contains the integer value 5.

Regarding now the package macros, there is first the case of arguments having to be short integers: this means that they are fed to a `\numexpr..\relax`, hence submitted to a *complete expansion* which must deliver an integer, and count registers and even algebraic expressions with them like `\mycountA+\mycountB*17-\mycountC/12+\mycountD` are admissible arguments (the slash stands here for the rounded integer division done by `\numexpr`). This applies in particular to the number of digits to truncate or round with, to the indices of a series partial sum, ...

The macros allowing the extended format for long numbers or dealing with fractions will to *some extent* allow the direct use of count registers and even infix algebra inside their arguments: a count register `\mycountA` or `\count 255` is admissible as numerator or also as denominator, with no need to be prefixed by `\the` or `\number`. It is possible to have as argument an algebraic expression as would be acceptable by a `\numexpr..\relax`, under this condition: *each of the numerator and denominator is expressed with at most nine tokens*.⁵⁶ ⁵⁷ Important: a slash for rounded division in a `\numexpr` should be written with braces `{/}` to not be confused with the `xintfrac` delimiter between numerator and denominator (braces will be removed internally and the slash will count for one token). Example: `\mycountA+\mycountB{/}17/1+\mycountA*\mycountB`, or `\count 0+\count 2{/}17/1+\count 0*\count 2`.

```
\cnta 10 \cntb 35 \xintRaw {\cnta+\cntb{/}17/1+\cnta*\cntb}->12/351[0]
```

For longer algebraic expressions using count registers, there are two possibilities:

1. let the numerator and the denominator be presented as `\the\numexpr..\relax`,
2. or as `\numexpr {...}\relax` (the braces are removed during processing; they are not legal for `\numexpr..\relax` syntax.)

```
\cnta 100 \cntb 10 \cntc 1
```

⁵⁵ refer to the documentation of `\xintPRaw` for an alternative. ⁵⁶ The 1.2k and earlier versions manual claimed up to 8 tokens, but low-level TeX error arose if the `\numexpr ..\relax` occupied exactly 8 tokens and evaluated to zero. With 1.21 and later, up to 9 tokens are always safe and one may even drop the ending `\relax`. But well, all these explanations are somewhat silly because prefixing by `\the` or `\number` is always working with arbitrarily many tokens. ⁵⁷ Attention! in the \LaTeX context a `\value{countname}` will behave ok only if it is first in the input, if not it will not get expanded, and braces around the name will be removed and chaos will ensue inside a `\numexpr`. One should enclose the whole input in `\the\numexpr..\relax` in such cases.



```
\xintPRaw {\numexpr {\cnta*\cnta+\cntb*\cntb+\cntc*\cntc+
                    2*\cnta*\cntb+2*\cnta*\cntc+2*\cntb*\cntc}\relax/%
          \numexpr {\cnta*\cnta+\cntb*\cntb+\cntc*\cntc}\relax }
12321/10101
```

8.7. Dimension registers and variables

$\langle dimen \rangle$ variables can be converted into (short) integers suitable for the `xint` macros by prefixing them with `\number`. This transforms a dimension into an explicit short integer which is its value in terms of the `sp` unit (1/65536 pt). When `\number` is applied to a $\langle glue \rangle$ variable, the stretch and shrink components are lost.

For \TeX users: a length is a $\langle glue \rangle$ variable, prefixing a length macro defined by `\newlength` with `\number` will thus discard the `plus` and `minus` glue components and return the dimension component as described above, and usable in the `xint` bundle macros.

This conversion is done automatically inside an `\xintexpr`-essions, with tacit multiplication implied if prefixed by some (integral or decimal) number.

One may thus compute areas or volumes with no limitations, in units of `sp^2` respectively `sp^3`, do arithmetic with them, compare them, etc. . . . , and possibly express some final result back in another unit, with the suitable conversion factor and a rounding to a given number of decimal places.

A [table of dimensions](#) illustrates that the internal values used by \TeX do not correspond always to the closest rounding. For example a millimeter exact value in terms of `sp` units is $72.27/10/2.54*65536=186467.981\dots$ and \TeX uses internally `186467sp` (\TeX truncates to get an integral multiple of the `sp` unit; see at the end of this section the exact rules applied internally by \TeX).

Unit	definition	Exact value in sp units	\TeX 's value in sp units	Relative error
cm	0.01 m	$236814336/127 = 1864679.811\dots$	1864679	-0.0000%
mm	0.001 m	$118407168/635 = 186467.981\dots$	186467	-0.0005%
in	2.54 cm	$118407168/25 = 4736286.720\dots$	4736286	-0.0000%
pc	12 pt	$786432 = 786432.000\dots$	786432	0%
pt	1/72.27 in	$65536 = 65536.000\dots$	65536	0%
bp	1/72 in	$1644544/25 = 65781.760\dots$	65781	-0.0012%
3bp	1/24 in	$4933632/25 = 197345.280\dots$	197345	-0.0001%
12bp	1/6 in	$19734528/25 = 789381.120\dots$	789381	-0.0000%
72bp	1 in	$118407168/25 = 4736286.720\dots$	4736286	-0.0000%
dd	1238/1157 pt	$81133568/1157 = 70124.086\dots$	70124	-0.0001%
11dd	$11*1238/1157$ pt	$892469248/1157 = 771364.950\dots$	771364	-0.0001%
12dd	$12*1238/1157$ pt	$973602816/1157 = 841489.037\dots$	841489	-0.0000%
sp	1/65536 pt	$1 = 1.000\dots$	1	0%

\TeX dimensions

There is something quite amusing with the Didot point. According to the \TeX Book, $1157\text{ dd}=1238\text{ p}_2\text{ t}$. The actual internal value of `1 dd` in \TeX is `70124 sp`. We can use `xintcfrac` to display the list of centered convergents of the fraction $70124/65536$:

```
\xintListWithSep{, }{\xintFtoCCv{70124/65536}}
```

1/1, 15/14, 61/57, 107/100, 1452/1357, 17531/16384, and we don't find 1238/1157 therein, but another approximant 1452/1357!

And indeed multiplying $70124/65536$ by 1157, and respectively 1357, we find the approximations (wait for more, later):

```
`1157 dd' '=1237.998474121093. . . pt
`1357 dd' '=1451.999938964843. . . pt
```

and we seemingly discover that `1357 dd=1452 pt` is *far more accurate* than the \TeX Book formula `1157 dd=1238 pt` ! The formula to compute `N dd` was

```
\xinttheexpr trunc(N\dimexpr 1dd\relax/\dimexpr 1pt\relax,12)\relax}
```

What's the catch? The catch is that \TeX *does not* compute `1157 dd` like we just did:

```
1157 dd=\number\dimexpr 1157dd\relax/65536=1238.000000000000...pt
```

```
1357 dd=\number\dimexpr 1357dd\relax/65536=1452.001724243164...pt
```

We thus discover that \TeX (or rather here, $e\text{-}\TeX$, but one can check that this works the same in $\TeX 82$), uses $1238/1157$ as a conversion factor (and necessarily intermediate computations simulate higher precision than a priori available with integers less than 2^{31} or rather 2^{30} for dimensions). Hence the $1452/1357$ ratio is irrelevant, an artefact of the rounding (or rather, as we see, truncating) for one `dd` to be expressed as an integral number of `sp`'s.

Let us now use `\xintexpr` to compute the value of the Didot point in millimeters, if the above rule is exactly verified:

```
\xinttheexpr trunc(1238/1157*25.4/72.27,12)\relax=0.376065027442...mm
```

This fits very well with the possible values of the Didot point as listed in the [Wikipedia Article](#). The value `0.376065 mm` is said to be *the traditional value in European printers' offices*. So the `1157 dd=1238 pt` rule refers to this Didot point, or more precisely to the *conversion factor* to be used between this Didot and \TeX points.

The actual value in millimeters of exactly one Didot point as implemented in \TeX is

```
\xinttheexpr trunc(\dimexpr 1dd\relax/65536/72.27*25.4,12)\relax
=0.376064563929...mm
```

The difference of circa 5\AA is arguably tiny!

By the way the *European printers' offices* (dixit Wikipedia) *Didot* is thus exactly

```
\xinttheexpr reduce(.376065/(25.4/72.27))\relax=543564351/508000000 pt
```

and the centered convergents of this fraction are `1/1`, `15/14`, `61/57`, `107/100`, `1238/1157`, `11249/10513`, `23736/22183`, `296081/276709`, `615898/575601`, `11382245/10637527`, `22148592/20699453`, `188570981/176233151`, `543564351/508000000`. We do recover the `1238/1157` therein!

Here is how \TeX converts `abc.xyz...<unit>`. First the decimal is *rounded* to the nearest integral multiple of $1/65536$, say `X/65536`. The `<unit>` is associated to a ratio `N/D`, which represents `<unit>/pt`. For the Didot point the ratio is indeed `1238/1157`. \TeX *truncates* the fraction `XN/D` to an integer `M`. The dimension is represented by `M sp`.

8.8. `\ifcase`, `\ifnum`, ... constructs

When using things such as `\ifcase \xintSgn{\A}` one has to make sure to leave a space after the closing brace for \TeX to stop its scanning for a number: once \TeX has finished expanding `\xintSgn{\A}` and has so far obtained either `1`, `0`, or `-1`, a space (or something `\unexpandable`) must stop it looking for more digits. Using `\ifcase\xintSgn\A` without the braces is very dangerous, because the blanks (including the end of line) following `\A` will be skipped and not serve to stop the number which `\ifcase` is looking for.

```
\begin{enumerate}[nosep]\def\A{1}
\item \ifcase \xintSgn\A 0\or OK\else ERROR\fi
\item \ifcase \xintSgn\A\space 0\or OK\else ERROR\fi
\item \ifcase \xintSgn{\A} 0\or OK\else ERROR\fi
\end{enumerate}
```

1. ERROR
2. OK
3. OK

In order to use successfully `\if...\fi` constructions either as arguments to the `xint` bundle expandable macros, or when building up a completely expandable macro of one's own, one needs some \TeX nical expertise (see also [item 2](#) on page 129).

It is thus much to be recommended to use the expandable branching macros, provided by `xintfrac` such as `\xintifSgn`, `\xintifZero`, `\xintifOne`, `\xintifNotZero`, `\xintifTrueAelseB`, `\xintifCmp`, `\xintifGt`, `\xintifLt`, `\xintifEq`, `\xintifInt`... See their respective documentations. All these conditionals always have either two or three branches, and empty brace pairs `{}` for unused branches should not be forgotten.

If these tests are to be applied to standard \TeX short integers, it is more efficient to use (under \LaTeX) the equivalent conditional tests from the [etoolbox](#)⁵⁸ package.

8.9. No variable declarations are needed

There is no notion of a *declaration of a variable*.

To do a computation and assign its result to some macro `\z`, the user will employ the `\def`, `\edef`, or `\newcommand` (in \LaTeX) as usual, keeping in mind that two expansion steps are needed, thus `\edef` is initially the main tool:

```
\def\x{1729728} \def\y{352827927} \edef\z{\xintiiMul {\x}{\y}}
\meaning\z
```

macro:->610296344513856

As an alternative to `\edef` the package provides `\oodef` which expands exactly twice the replacement text, and `\fdef` which applies *f-expansion* to the replacement text during the definition.

```
\def\x{1729728}\def\y{352827927}
\oodef\w {\xintiiMul\x\y} \fdef\z{\xintiiMul {\x}{\y}}
\meaning\w, \meaning\z
```

macro:->610296344513856, macro:->610296344513856



In practice `\oodef` is slower than `\edef`, except for computations ending in very big final replacement texts (thousands of digits). On the other hand `\fdef` appears to be slightly faster than `\edef` already in the case of expansions leading to only a few dozen digits.

`xintexpr` does provide an interface to declare and assign values to identifiers which can then be used in expressions: [subsection 2.9](#).

8.10. Possible syntax errors to avoid

Here is a list of imaginable input errors. Some will cause compilation errors, others are more annoying as they may pass through unsignaled.

- using `-` to prefix some macro: `-\xintiiSqr{35}/271`.⁵⁹
- using one pair of braces too many `\xintIrr{{\xintiiPow {3}{13}}/243}` (the computation goes through with no error signaled, but the result is completely wrong).
- things like `\xintiiAdd { \x}{\y}` as the space will cause `\x` to be expanded later, most probably within a `\numexpr` thus provoking possibly an arithmetic overflow.
- using `[]` and decimal points at the same time `1.5/3.5[2]`, or with a sign in the denominator `3/-5[7]`. The scientific notation has no such restriction, the two inputs `1.5/-3.5e-2` and `-1.5e2/3.5` are equivalent: `\xintRaw{1.5/-3.5e-2}=-15/35[2]`, `\xintRaw{-1.5e2/3.5}=-15/35[2]`.
- generally speaking, using in a context expecting an integer (possibly restricted to the \TeX bound) a macro or expression which returns a fraction: `\xinttheexpr 4/2\relax` outputs `4/2`, not `2`. Use `\xintNum {\xinttheexpr 4/2\relax}` or `\xinttheiexpr 4/2\relax` (which rounds the result to the nearest integer, here, the result is already an integer) or `\xinttheiiexpr 4/2\relax`. Or, divide in your head `4` by `2` and insert the result directly in the \TeX source.

⁵⁸ <https://ctan.org/pkg/etoolbox> ⁵⁹ to the contrary, this *is* allowed inside an `\xintexpr`-ession.

8.11. Error messages

In situations such as division by zero, the \TeX run will be interrupted with some error message. It conveys some short information on the cause of the problem,⁶⁰ then an optimistic statement about a possible recovery if the user (in interactive mode) simply hits the `<return>` key. In non-interactive (`nonstopmode`) the \TeX run goes on uninterrupted and the error data will be found in the compilation log. Often, `xint` will fall-back to using a zero value. This is still an experimental feature.^{61 62}

The encouragements will be slightly better formatted if the run is with \LaTeX compared to Plain $\varepsilon\text{-}\TeX$: Plain by default does not set the `\newlinechar` which allows to issue linebreaks in messages at chosen locations. In the examples here, `xintsession` is used, and it loads `xint` in a way activating the nicer `\newlinechar` formatted messages, even though it runs (a priori, but not necessarily) under Plain $\varepsilon\text{-}\TeX$.

```
>>> 1/0;
Runaway argument?
! xint error: Division by zero: 1/0.
! Paragraph ended before \xint<...> is done, but will resume:
  hit <return> at the ? prompt to try fixing the error above
  which has been encountered before expansion was complete.
<to be read again>
          \par
...
1.602 \xintsession
          \endinput%^M
?
@_1    0
>>> (-1)^3.2;
Runaway argument?
! xint error: Fractional power 32/1[-1] of negative -1[0].
! Paragraph ended before \xint<...> is done, but will resume:
  hit <return> at the ? prompt to try fixing the error above
  which has been encountered before expansion was complete.
<to be read again>
          \par
...
1.602 \xintsession
          \endinput%^M
?
@_2    0
>>> cos 1);
Runaway argument?
! xint error: `cos1' unknown, say `Isome_var' or I use 0.
! Paragraph ended before \xint<...> is done, but will resume:
  hit <return> at the ? prompt to try fixing the error above
  which has been encountered before expansion was complete.
<to be read again>
          \par
...
1.602 \xintsession
          \endinput%^M
```

**Changed
at 1.4m!**

⁶⁰ The wording of these messages has been last modified at 1.4m. ⁶¹ Customizable handlers, error traps, error flags are implemented in embryonic form but without user interface since 1.21 release. This is not ready yet. ⁶² The 1.4g new formatting implementation benefited from a May 2021 thread at the \LaTeX 3 site where expandable error messages were discussed, with in particular contributions of [@blefloch](#) and [@Skillmon](#).


```

?
Runaway argument?
! xint error: Extra ) removed. Hit <return>, fingers crossed.
! Paragraph ended before \xint<...> is done, but will resume:
  hit <return> at the ? prompt to try fixing the error above
  which has been encountered before expansion was complete.
<to be read again>
      \par
...
1.602 \xintsession
      \endinput%^M
?
@_3      0
>>> 3=4;
Runaway argument?
! xint error: Expected an operator but got `='. Ignoring.
! Paragraph ended before \xint<...> is done, but will resume:
  hit <return> at the ? prompt to try fixing the error above
  which has been encountered before expansion was complete.
<to be read again>
      \par
...
1.602 \xintsession
      \endinput%^M
?
@_4      12
>>> &bye

```

In the last example, tacit multiplication was applied as `xintexpr` was looking for an operator, got some invalid input and then a number.

Some constructs in `xintexpr`-essions use delimited macros and there is thus possibility in case of an ill-formed expression to end up beyond the `\relax` end-marker. Such a situation can also occur from `\relax` being swallowed by a non-terminated `\numexpr`:

```
\xintexpr 3 + \numexpr 5+4\relax followed by some LaTeX code...
```

The correct input is

```
\xintexpr 3 + \numexpr 5+4\relax\relax
```

But people in their right mind will have done

```
\xintexpr 3 + 5 + 4\relax
```

A few will have done the computation in their heads.

In such cases low-level errors will arise and may lead to very cryptic messages; but nothing unusual or especially traumatizing for the daring experienced \TeX/\LaTeX user, whose has seen zillions of un-helpful error messages already in her daily practice of \TeX/\LaTeX .

8.12. Package namespace, catcodes

This section reviews (probably with some omissions) important miscellany regarding control sequence names and catcode matters and is basically in its entirety a

\TeX -hackers note:

- The bundle packages force the `\space` and `\empty` control sequences into having their default meanings as in Plain \TeX or \LaTeX 2e formats.
- Private macros (or internally used `\count` registers, and one private `\toks`) have names starting with `\xint_` or `\XINT_`. Some, for legacy or technical reasons, have `\xint` or `\XINT` prefix with no underscore.

- All public macros have their names starting with `\xint` except for: the `xintkernel` provided `\odef`, `\oodef`, `\f`, `\fdef`. If macros with these names already exist `xinttools` will not overwrite them. Their meanings are also available under the names `\xintodef`, `\xintoodef`, etc...
- Whether in arguments of `xintfrac` macros or inside `\xintexpressions` active characters will cause naturally problems, one thinks in particular of characters made active by `babel` languages such as `babel-fr` `ench` for `!?:;`, and for example the `!` if used as factorial operator will cause breakage. Prefixing with `\string` is one option, and there is also `\detokenize`.
- Precautions are not needed (but see next item) for `\xintdefvar` and `\xintdeffunc` as they use automatically `\xintexprSafeCatcodes` and `\xintexprRestoreCatcodes` to temporarily set catcodes to safe values.
- `\xintexprSafeCatcodes` and `\xintexprRestoreCatcodes` can be employed at user level too. They should be used in case of need before and after macro definitions themselves using `\xintexpr`, `\xintdefvar` or `\xintdeffunc`. Indeed the latter two won't be able to 'fix' problematic characters already present at definition time when expansion is later done. In the \TeX preamble, `babel` has not yet activated (normally) characters, so these precautions should be unneeded there.
- For the `xintfrac` macros to be able to parse their inputs, standard catcodes in the argument are assumed for the digits (of course), the plus and minus signs, the dot, the letter `e`, the forward slash, the square brackets. Spaces should be avoided although they may go unnoticed sometimes.
- For the `xintexpr` expressions there is more leeway: the digit tokens must have their standard catcodes, the letters must have their standard catcodes for variable and function names to be recognized, but other characters may mostly have unusual catcodes, if non-letter and not extremes. A few syntax elements are implemented via delimited macros and normal catcodes must be in effect for the comma, the equal sign and the closing parenthesis for those to work. Spaces are gobbled. The `e` of scientific notation may be `E` on input, `xintfrac` macros on the other hand will not recognize the `E`.
- At loading time the catcode configuration may be arbitrary as long as it satisfies the following requirements:
 - `%` has its normal category code,
 - `\` has its normal category code,
 - Latin letters have their normal category code "letter",
 - Digits have their normal category code "other".
 Nothing more is assumed, for example `{` and `}` may have unusual catcodes at package loading time. This will be admittedly unusual especially in \TeX as `\usepackage{xintexpr}` would then have had to be replaced by something such as `\usepackage<xintexpr>...`
- Loading the packages causes no insertion of space tokens.
- The previous two items also apply to usage of `\xintreloadxintlog` and of `\xintreloadxinttrig`.

8.13. Origins of the package

2013/03/28. Package `bigintcalc` by HEIKO OBERDIEK already provides expandable arithmetic operations on "big integers", i.e. integers beyond the \TeX bound $2^{31} - 1$, so why another⁶³ one?

I got started on this in early March 2013, via a thread on the `c.t.tex` usenet group, where ULRICH DIEZ used the previously cited package together with a macro (`\ReverseOrder`) which I had contributed to another thread.⁶⁴ What I had learned in this other thread thanks to interaction with ULRICH DIEZ and GL on expandable manipulations of tokens motivated me to try my hands at addition and multiplication.

I wrote macros `\bigMul` and `\bigAdd` which I posted to the newsgroup; they appeared to work comparatively fast. These first versions did not use the ε - \TeX `\numexpr` primitive, they worked one digit at a time, having previously stored carry-arithmetic in 1200 macros.

I noticed that the `bigintcalc` package used `\numexpr` if available, but (as far as I could tell) not to do computations many digits at a time. Using `\numexpr` for one digit at a time for `\bigAdd` and

⁶³ this section was written before the `xintfrac` package; the author is not aware of another package allowing expandable computations with arbitrarily big fractions. ⁶⁴ the `\ReverseOrder` could be avoided in that circumstance, but it does play a crucial rôle here.

`\bigMul` slowed them a tiny bit but avoided cluttering \TeX memory with the 1200 macros storing pre-computed digit arithmetic. I wondered if some speed could be gained by using `\numexpr` to do four digits at a time for elementary multiplications (as the maximal admissible number for `\numexpr` has ten digits).

2013/04/14. This initial `xint` was followed by `xintfrac` which handled exactly fractions and decimal numbers.

2013/05/25. Later came `xintexpr` and at the same time `xintfrac` got extended to handle floating point numbers.

2013/11/22. Later, `xinttools` was detached.

2014/10/28. Release 1.1 significantly extended the `xintexpr` parsers.

2015/10/10. Release 1.2 rewrote the core integer routines which had remained essentially unmodified, apart from a slight improvement of division early 2014.

This 1.2 release also got its impulse from a fast ```reversing''` macro, which I wrote after my interest got awakened again as a result of correspondence with Bruno LE FLOCH during September 2015: this new reverse uses a \TeX technique which *requires* the tokens to be digits. I wrote a routine which works (expandably) in quasi-linear time, but a less fancy $O(N^2)$ variant which I developed concurrently proved to be faster all the way up to perhaps 7000 digits, thus I dropped the quasi-linear one. The less fancy variant has the advantage that `xint` can handle numbers with more than 19900 digits (but not much more than 19950). This is with the current common values of the input save stack and maximal expansion depth: 5000 and 10000 respectively.

9. Macros of the **xintkernel** package

The **xintkernel** package contains mainly the common code base for handling the load-order of the bundle packages, the management of catcodes at loading time, definition of common constants and macro utilities which are used throughout the code etc ... it is automatically loaded by all packages of the bundle.

It provides a few macros possibly useful in other contexts.

<p>.1 \odef, \oodef, \fdef 140</p> <p>.2 \xintReverseOrder 140</p> <p>.3 \xintLength 140</p> <p>.4 \xintFirstItem 141</p> <p>.5 \xintLastItem 141</p>	<p>.6 \xintFirstOne 141</p> <p>.7 \xintLastOne 141</p> <p>.8 \xintReplicate, \xintreplicate 141</p> <p>.9 \xintGobble, \xintgobble 141</p> <p>.10 (WIP) \xintUniformDeviate 142</p>
---	---

9.1. \odef, \oodef, \fdef

\oodef\controlsequence {<stuff>} does

```
\expandafter\expandafter\expandafter\def
\expandafter\expandafter\expandafter\controlsequence
\expandafter\expandafter\expandafter{<stuff>}
```

This works only for a single \controlsequence, with no parameter text, even without parameters. An alternative would be:

```
\def\oodef #1#\def\oodefparametertext{#1}%
\expandafter\expandafter\expandafter\expandafter
\expandafter\expandafter\expandafter\def
\expandafter\expandafter\expandafter\oodefparametertext
\expandafter\expandafter\expandafter }
```

but it does not allow \global as prefix, and, besides, would have anyhow its use (almost) limited to parameter texts without macro parameter tokens (except if the expanded thing does not see them, or is designed to deal with them).

There is a similar macro \odef with only one expansion of the replacement text <stuff>, and \fdef which expands fully <stuff> using \romannumeral-`0.

They can be prefixed with \global. It appears that \fdef is generally a bit faster than \edef when expanding macros from the **xint** bundle, when the result has a few dozens of digits. \oodef needs thousands of digits it seems to become competitive.

xintkernel will not define these macros if the control sequence names already exist. It provides them always under the names \xintodef, \xintoodef and \xintfdef respectively.

9.2. \xintReverseOrder

n ★ \xintReverseOrder{<list>} does not do any expansion of its argument and just reverses the order of the tokens in the <list>. Braces are removed once and the enclosed material, now unbraced, does not get reversed. Unprotected spaces (of any character code) are gobbled.

```
\xintReverseOrder{\xintDigitsOf\xintiiPow {2}{100}\to\Stuff}
gives: \Stuff\to1002\xintiiPow\xintDigitsOf
```

xinttools provides a variant \xintRevWithBraces which keeps brace pairs in the output, and *f*-expands its input first.

For inputs consisting only digit tokens, see \xintReverseDigits from **xint**.

9.3. \xintLength

n ★ \xintLength{<list>} counts how many tokens (or braced items) there are (possibly none). It does no expansion of its argument, so to use it to count things in the replacement text of a macro \x

one should do `\expandafter\xintLength\expandafter{\x}`. Blanks between items are not counted. See also `\xintNthElt{0}` (from [xinttools](#)) which first *f*-expands its argument and then applies the same code.

```
\xintLength {\xintiiPow {2}{100}}=3
≠ \xintLen {\xintiiPow {2}{100}}=31
```

9.4. `\xintFirstItem`

n ★ `\xintFirstItem{⟨list⟩}` returns the first item of its argument, one pair of braces removed. If the list has no items the output is empty.

It does no expansion. For this and the next similar ones, see [sourcexint.pdf](#) for comments on limitations.

9.5. `\xintLastItem`

Added at 1.2i.

n ★ `\xintLastItem{⟨list⟩}` returns the last item of its argument, one pair of braces removed. If the list has no items the output is empty.

It does no expansion, which should be obtained via suitable `\expandafter`'s. See also `\xintNthElt{-1}` from [xinttools](#) which obtains the same result (but with another code) after having however *f*-expanded its argument first.

9.6. `\xintFirstOne`

n ★ `\xintFirstOne{⟨list⟩}` returns the first item as a braced item. i.e. if it was braced the braces are kept, else the braces are added. It looks like using `\xintFirstItem` within braces, but the difference is when the input was empty. Then the output is empty.

It does no expansion, which should be obtained via suitable `\expandafter`'s.

9.7. `\xintLastOne`

n ★ `\xintLastOne{⟨list⟩}` returns the last item as a braced item. i.e. if it was braced the braces are kept, else the braces are added. It looks like using `\xintLastItem` within braces, but the difference is when the input was empty. Then the output is empty.

It does no expansion, which should be obtained via suitable `\expandafter`'s.

9.8. `\xintReplicate`, `\xintreplicate`

^{num}_X *n* ★ `\romannumeral\xintreplicate{x}{⟨stuff⟩}` is simply copied over from \TeX 's `\prg_replicate:nn` with some minor changes.⁶⁵

And `\xintReplicate{x}` integrates the `\romannumeral` prefix.

It does not do any expansion of its second argument but inserts it in the upcoming token stream precisely *x* times. Using it with a negative *x* raises no error and does nothing.⁶⁶

9.9. `\xintGobble`, `\xintgobble`

^{num}_X ★ `\romannumeral\xintgobble{x}` is a Gobbling macro written in the spirit of \TeX 's `\prg_replicate:nn` (which I cloned as `\xintreplicate`.) It gobbles *x* tokens upstream, with *x* allowed to be as large as 531440. Don't use it with *x*<0.

And `\xintGobble{x}` integrates the `\romannumeral`.

`\xintgobble` looks as if it must be related to `\xintTrim` from [xinttools](#), but the latter uses different code (using directly `\xintgobble` is not possible because one must make sure not to gobble

⁶⁵ I started with the code from Joseph WRIGHT available on an online site. ⁶⁶ This behaviour may change in future.

more than the number of available items; and counting available items first is an overhead which `\xintTrim` avoids.) It is rather `\xintKeep` with a negative first argument which hands over to `\xintgobble` (because in that case it is needed to count anyhow beforehand the number of items, hence `\xintgobble` can then be used safely.)

I wrote an `\xintcount` in the same spirit as `\xintreplicate` and `\xintgobble`. But it needs to be counting hundreds of tokens to be worth its salt compared to `\xintLength`.

9.10. (WIP) `\xintUniformDeviate`

$\frac{\text{num}}{x}$ ★ `\xintUniformDeviate{x}` is a wrapper of engine `\pdfuniformdeviate` (or `\uniformdeviate`).⁶⁷ The implementation is to be considered experimental for the time being.

The argument is expanded in `\numexpr` and the macro itself needs two expansion steps. It produces like the engine primitive an integer (digit tokens) with minimal value 0 and maximal one $x-1$ if x is positive, or minimal value $x+1$ and maximal value 0 if x is negative. For the discussion next, x is supposed positive as this avoids having to insert absolute values in formulas.

The underlying engine Random Number Generator works with an array of 55 28bits integers. To produce a “uniform” random integer in a given range $0..x-1$ it produces next pseudo-random y (supposedly uniformly distributed, i.e. non-uniformity can be neglected) such that $0 \leq y < 2^{28}$ and the output is the rounding of $x * (y/2^{28})$, with upper bound x remapped to 0. This has following corollaries:

1. with $x=2^{29}$ or $x=2^{30}$ the engine primitive produces only even numbers,
2. with $x=3*2^{26}$ the integers produced by the RNG when taken modulo three obey the proportion 1:1:2, not 1:1:1,
3. with $x=3*2^{14}$ there is analogous although weaker non-uniformity of the random integers when taken modulo 3,
4. generally speaking pure powers of two should generate uniform random integers, but when the range is divisible by large powers of two, the non-uniformity may be amplified in surprising ways by modulo operations.

These observations are not to be construed as criticism of the engine primitive itself, which comes from MetaPost, as the code comments and more generally the whole of *The Art of Computer Programming, Vol. 2* stresses that it should rather be seen as producing random fractions (the unit fraction being 2^{28}). Using it as a generator for *integers* is a bit of an abuse.

The first goal of `\xintUniformDeviate` is to guarantee a better uniformity for the distribution of random integers in any given range x .

If the probability to obtain a given y in $0..x-1$ is $(1+e(y))/x$, the “relative non-uniformity” for that value y is $|e(y)|$.

The engine primitive guarantees only $x/2^{28}$ relative non-uniformity, and `\xintUniformDeviate` (in its current implementation) improves this by a factor $2^{28}=268435456$: the non-uniformity is guaranteed to be bounded by $x/2^{56}$.⁶⁸ With such a small non-uniformity, modulo phenomena as mentioned earlier are not observable in reasonable computing time.

```
%\xintdefiifunc mod3(x):= x 'mod' 3;
\xintNewIIExpr\ModThree[1]{#1 'mod' 3}

\pdfsetrandomseed 87654321
```

⁶⁷ The `\uniformdeviate` primitive has been added to XeTeX and will be available with TeXLive 2019 release. ⁶⁸ These estimates assume that the engine RNG underlying stream of 28-bits integers can be considered uniform; it is known that the parity bits of these 28-bits integers have a period of $55(2^{55}-1)$ and that after that many draws the count of 1s has only an excess of 55 compared to the count of 0s, so the scale seems to be an intrinsic non-uniformity of 2^{-55} but it is not obvious if it applies to much shorter ranges. At any rate we assumed that the non-uniformity for x a power of two less than 2^{28} is negligible in comparison to 2^{-28} . Bigger powers of 2 produce only even integers because the output is rescaled by factor $x/2^{28}$!

```
\xintdefiivar BadDigits:=qraw(%
  \romannumeral\xintreplicate{504}{\ModThree{\pdfuniformdeviate "C000000}}}%
);%

\pdfsetrandomseed 87654321
\xintdefiivar GoodDigits:=qraw(%
  \romannumeral\xintreplicate{504}{\ModThree{\xintUniformDeviate{"C000000}}}%
);%
```

These 504 digits generated from `\string\pdfuniformdeviate:`
`\xinttheiexpr BadDigits\relax\hfill\break`
 contain these respective amounts of 0, 1, and 2:
 % (this is definitely not the fastest way to count, but it is fun - and expandable)
 % !(i) (logical not) is short-cut for the vanishing test `i==0`, but it would be
 % simpler to use `(i)?{i is not zero}{i is zero}` with permuted branches
`\xinttheiexpr iter(0,0,0;(!i))?{[0][0]+1,[0][1],[0][2]}`
 `{(isone(i))?[0][0],[0][1]+1,[0][2]}`
 `{[0][0],[0][1],[0][2]+1}},`
 `i=BadDigits)\relax\par`

These 504 digits generated from `\string\xintUniformDeviate:`
`\xinttheiexpr GoodDigits\relax\hfill\break`
 contain these respective amounts of 0, 1, and 2:
`\xinttheiexpr iter(0,0,0;(!i))?{[0][0]+1,[0][1],[0][2]}`
 `{(isone(i))?[0][0],[0][1]+1,[0][2]}`
 `{[0][0],[0][1],[0][2]+1}},`
 `i=GoodDigits)\relax\par`

```
% % output to data file for double-check with python
% \newwrite\out
% \immediate\openout\out=\jobname.data
% \immediate\write\out{Lbad=[\xinttheiexpr BadDigits\relax]}
% \immediate\write\out{Lgood=[\xinttheiexpr GoodDigits\relax]}
% \immediate\closeout\out
```

These 504 digits generated from `\pdfuniformdeviate`: 1, 2, 1, 2, 1, 0, 0, 1, 2, 1, 1, 0, 1, 0, 0,
 2, 2, 0, 0, 2, 2, 2, 1, 2, 2, 2, 2, 0, 0, 2, 1, 1, 0, 0, 1, 1, 1, 2, 2, 2, 2, 0, 1, 1, 2, 1, 1, 1, 2,
 1, 2, 2, 1, 0, 0, 2, 0, 0, 2, 0, 2, 1, 1, 2, 2, 0, 0, 2, 2, 2, 2, 2, 0, 2, 1, 2, 2, 1, 1, 0, 1, 2, 2, 0,
 1, 0, 2, 2, 2, 0, 1, 1, 1, 0, 1, 1, 0, 0, 1, 2, 2, 2, 1, 2, 2, 0, 1, 0, 2, 1, 2, 2, 1, 1, 2, 2, 1, 2, 2,
 1, 2, 2, 1, 2, 2, 2, 1, 0, 2, 2, 2, 1, 1, 2, 1, 2, 2, 1, 2, 1, 0, 2, 2, 2, 2, 1, 2, 0, 2, 2, 2, 0, 0, 2,
 0, 1, 0, 2, 2, 2, 2, 2, 2, 2, 2, 0, 0, 2, 2, 1, 0, 2, 1, 2, 0, 1, 0, 1, 0, 2, 2, 0, 2, 2, 2, 2, 2, 0, 2,
 2, 2, 2, 2, 1, 2, 2, 1, 1, 1, 1, 2, 0, 2, 1, 0, 1, 2, 1, 1, 1, 2, 1, 1, 2, 0, 0, 2, 1, 1, 2, 2, 0, 1, 2,
 2, 1, 2, 0, 1, 1, 2, 2, 1, 2, 2, 1, 2, 1, 1, 2, 1, 2, 2, 1, 1, 0, 1, 1, 2, 2, 0, 1, 1, 2, 2, 2, 0, 1, 0,
 2, 1, 1, 1, 0, 1, 0, 2, 1, 1, 2, 2, 1, 2, 0, 2, 2, 1, 0, 1, 2, 2, 0, 2, 0, 2, 2, 2, 2, 0, 2, 0, 2, 2, 2,
 2, 2, 1, 1, 2, 1, 2, 1, 1, 0, 2, 2, 2, 1, 2, 0, 0, 2, 2, 0, 1, 0, 2, 2, 2, 2, 2, 1, 2, 0, 1, 0, 0, 2, 0,
 2, 2, 1, 0, 1, 0, 1, 1, 2, 0, 1, 0, 0, 2, 1, 0, 2, 2, 0, 1, 2, 1, 1, 0, 0, 0, 0, 0, 0, 2, 2, 2, 0, 0, 2,
 2, 2, 0, 2, 2, 1, 1, 2, 1, 1, 2, 0, 1, 2, 1, 2, 1, 1, 2, 2, 0, 2, 2, 2, 1, 2, 2, 1, 0, 1, 2, 2, 2, 2, 2,
 0, 2, 2, 1, 0, 2, 0, 1, 2, 2, 2, 2, 1, 0, 1, 2, 0, 2, 2, 2, 2, 2, 2, 2, 0, 1, 0, 2, 1, 2, 1, 0, 2, 1, 1,
 2, 2, 0, 1, 0, 0, 0, 2, 2, 1, 2, 2, 1, 1, 1, 0, 1, 2, 2, 2, 0, 2, 1, 0, 1, 2, 2, 1, 2, 0, 2, 0, 1, 0, 1,
 0, 0, 2, 0, 2, 1, 2, 0, 1, 0, 0, 1, 1, 2, 0, 2, 0, 0, 2, 0, 2, 0, 0, 0, 0, 2, 1, 0, 2, 2, 0, 2, 1, 2

contain these respective amounts of 0, 1, and 2: 124, 147, 233

These 504 digits generated from `\xintUniformDeviate`: 2, 1, 0, 1, 0, 2, 1, 1, 2, 2, 2, 1, 2, 2, 0,
 1, 0, 2, 2, 0, 1, 2, 0, 2, 0, 0, 2, 2, 1, 1, 1, 2, 1, 1, 1, 1, 2, 1, 1, 0, 2, 1, 1, 0, 0, 2, 2, 1, 0, 0,
 1, 1, 0, 2, 0, 2, 0, 0, 1, 0, 2, 0, 0, 1, 1, 0, 0, 1, 2, 0, 0, 2, 2, 0, 1, 1, 0, 2, 1, 2, 0, 1, 0, 0, 1,

2, 1, 1, 0, 1, 0, 2, 2, 1, 0, 1, 1, 2, 1, 0, 2, 0, 1, 1, 0, 1, 0, 0, 1, 0, 1, 1, 2, 0, 2, 1, 2, 2, 1, 0, 0, 1, 2, 2, 2, 1, 1, 1, 0, 2, 1, 0, 2, 2, 2, 2, 0, 0, 1, 2, 2, 2, 2, 2, 0, 2, 2, 1, 0, 1, 2, 0, 0, 2, 0, 0, 2, 1, 1, 0, 1, 1, 1, 1, 2, 1, 1, 1, 2, 0, 1, 1, 1, 1, 0, 1, 0, 0, 0, 2, 0, 2, 1, 1, 0, 0, 2, 1, 1, 1, 0, 2, 2, 0, 1, 1, 1, 2, 0, 1, 0, 2, 1, 2, 1, 0, 2, 0, 1, 1, 2, 2, 0, 0, 2, 1, 0, 2, 0, 2, 2, 2, 1, 0, 0, 1, 2, 2, 2, 1, 1, 0, 1, 1, 2, 2, 2, 2, 2, 2, 2, 2, 1, 1, 1, 1, 1, 1, 1, 2, 2, 0, 0, 0, 1, 0, 0, 0, 1, 0, 2, 1, 0, 2, 2, 1, 1, 1, 0, 1, 0, 0, 1, 1, 1, 0, 2, 1, 1, 0, 0, 0, 0, 0, 0, 1, 2, 0, 0, 0, 0, 0, 2, 1, 1, 0, 2, 0, 1, 1, 1, 1, 2, 1, 1, 2, 2, 1, 2, 2, 1, 1, 0, 1, 2, 2, 2, 0, 1, 1, 2, 0, 2, 1, 0, 0, 2, 1, 0, 1, 1, 2, 1, 1, 0, 1, 0, 2, 0, 1, 1, 1, 0, 2, 2, 0, 1, 2, 2, 0, 1, 0, 2, 2, 0, 1, 1, 2, 2, 0, 2, 1, 0, 1, 0, 1, 2, 0, 1, 1, 2, 1, 2, 2, 0, 2, 2, 0, 1, 0, 0, 0, 2, 2, 2, 2, 0, 0, 2, 1, 0, 0, 1, 2, 2, 2, 2, 0, 2, 1, 0, 0, 2, 1, 0, 1, 2, 1, 1, 2, 1, 0, 0, 0, 2, 1, 0, 1, 0, 1, 1, 2, 2, 0, 2, 1, 1, 2, 0, 0, 1, 0, 2, 0, 1, 0, 2, 0, 1, 2, 1, 0, 2, 0, 1, 1, 2, 0, 2, 2, 2, 1, 1, 2, 2, 2, 1, 2, 2, 2, 0, 0, 0, 2, 0, 2, 2, 1, 2, 2, 1, 1, 2, 1, 1, 0, 0, 0, 0, 0, 0

contain these respective amounts of 0, 1, and 2: 161, 174, 169

There is a second peculiarity of the engine RNG: two seeds sharing the same low k bits generate sequences of 28-bits integers which are identical modulo 2^k ! In particular after setting the seed, there are only 2 distinct sequences of parity bits for the integers generated by `\pdfuniformdeviate` (2 to the power 28)...

In order to mitigate, `\xintUniformDeviate` currently only uses the seven high bits from the underlying random stream, using multiple calls to `\pdfuniformdeviate 128`. From the Birthday Effect, after about 2^{11} seeds one will likely pick a new one sharing its 22 low bits with an earlier one.

1. but as the final random integer is obtained by additional operations involving the range x (currently a modulo operation), for odd ranges it is more difficult for bit correlations to be seen,
2. anyway as they are only 2^{28} seeds in total, after only 2^{14} seeds it is likely to encounter one already explored, and then random integers are identical, however complicated the RNG's raw output is malaxed, and whatever the target range x . And 2^{14} is only eight times as large as 2^{11} .

It would be nice if the engine provided some user interface for letting its RNG execute a given number of iterations without the overhead of replicated executions of `\pdfuniformdeviate`. This could help gain entropy and would reduce correlations across series from distinct seeds.

The description above summarizes parts of discussions held with Bruno Le Floch in May 2018 on occasion of his LaTeX3 contributions related to this.

TeX-hackers note: currently the implementation of `\xintUniformDeviate` consumes exactly 5 calls to the engine primitive at each execution; the improved $x/2^{56}$ non-uniformity could be obtained with only 2 calls, but paranoia about the phenomenon of seeds with common bits has led me to accept the overhead of using the 7 high bits of 4 random 28bits integers, rather than one single 28bits integer, or two, or three.

Timings indicate that one `\xintUniformDeviate` has a time cost about 13 times the one for one call to the engine primitive (and not only 5, as the extra arithmetic expressions add overhead which is more costly than the primitive itself). Except if the code using the pseudo-random number is very short, this time penalty will prove in practice much less severe (and this is one important reason why we opted for obtaining 28bits via the 7 high bits of 4 successive pseudo random numbers from the engine primitive). For example let's raise 100 times a random integer to the tenth power:⁶⁹

```
\pdfsetrandomseed 12345678
\pdfresettimer
\romannumeral\xintreplicate{100}
{\fdef\foo{\xintiiPow{\xintUniformDeviate{100000000}}{10}}}%
\the\dimexpr\pdfelapsedtime sp\relax\space (with \string\xintUniformDeviate)\newline
\pdfsetrandomseed 12345678
\pdfresettimer
```

⁶⁹ Timings done during dvi build on a computer with a 2GHz cpu, old but still kicking, and with a locally compiled latex binary.


```
\romannumeral\xintreplicate{100}  
      {\fdef\foo{\xintiiPow{\pdfuniformdeviate 100000000}{10}}}%  
\the\dimexpr\pdfelapsedtime sp\relax\space (with \string\pdfuniformdeviate)\par
```

0.04562pt (with `\xintUniformDeviate`)

0.04767pt (with `\pdfuniformdeviate`)

The macros `\xintRandomDigits` or `\xintiiRandRange`, and their variants, as well as the supporting macros for `random()` generate random decimal digits eight by eight as if using `\xintUniformDeviate{100000000}`, but via a direct optimized call made possibly by the range being a power of 10.

10. Macros of the **xintcore** package

Package **xintcore** is automatically loaded by **xint**.

xintcore provides for big integers the four basic arithmetic operations (addition, subtraction, multiplication, division), as well as powers and factorials.

In the descriptions of the macros $\{N\}$ and $\{M\}$ stand for (big) integers or macros *f-expanding* to such big integers in strict format as described in [subsection 8.4](#).

All macros require strict integer format on input and produce strict integer format on output, except:



- `\xintiNum` which converts to strict integer format an input in *extended* integer format, i.e. admitting multiple leading plus or minus signs, then possibly leading zeroes, then digits,
- and `\xintNum` which is an alias for the former, which gets redefined by **xintfrac** to accept more generally also decimal numbers or fractions as input and which truncates them to integers.

The **ii** in the names of the macros such as `\xintiiAdd` serves to stress that they accept only strict integers as input (this is signaled by the margin annotation *f*), or macros *f-expanding* to such strict format (big) integers and that they produce strict integers as output.

Other macros, such as `\xintDouble`, lack the **ii**, but this is only a legacy of the history of the package and they have the same requirements for input and format of output as the **ii**-macros.

The letter **x** (with margin annotation ^{num}*x*) stands for an argument which will be handled embedded in `\numexpr..\relax`. It will thus be completely expanded and must give an integer obeying the TeX bounds. See also [subsection 8.6](#). This is the case for the argument of `\xintiiFac` or the exponent argument of `\xintiiPow`.

The *****'s in the margin are there to remind of the complete expandability, even *f-expandability* of the macros, as discussed in [subsubsection 8.3.1](#).

<p>.1 <code>\xintiNum</code> 146</p> <p>.2 <code>\xintDouble</code> 146</p> <p>.3 <code>\xintHalf</code> 146</p> <p>.4 <code>\xintInc</code> 147</p> <p>.5 <code>\xintDec</code> 147</p> <p>.6 <code>\xintDSL</code> 147</p> <p>.7 <code>\xintDSR</code> 147</p> <p>.8 <code>\xintDSRr</code> 147</p> <p>.9 <code>\xintFDg</code> 147</p> <p>.10 <code>\xintLDg</code> 147</p> <p>.11 <code>\xintiiSgn</code> 147</p> <p>.12 <code>\xintiiOpp</code> 147</p> <p>.13 <code>\xintiiAbs</code> 147</p> <p>.14 <code>\xintiiAdd</code> 147</p>	<p>.15 <code>\xintiiCmp</code> 147</p> <p>.16 <code>\xintiiSub</code> 148</p> <p>.17 <code>\xintiiMul</code> 148</p> <p>.18 <code>\xintiiSqr</code> 148</p> <p>.19 <code>\xintiiPow</code> 148</p> <p>.20 <code>\xintiiFac</code> 148</p> <p>.21 <code>\xintiiDivision</code> 148</p> <p>.22 <code>\xintiiQuo</code> 149</p> <p>.23 <code>\xintiiRem</code> 149</p> <p>.24 <code>\xintiiDivRound</code> 149</p> <p>.25 <code>\xintiiDivTrunc</code> 149</p> <p>.26 <code>\xintiiDivFloor</code> 149</p> <p>.27 <code>\xintiiMod</code> 149</p> <p>.28 <code>\xintNum</code> 150</p>
--	---

10.1. `\xintiNum`

f ★ `\xintiNum{N}` removes chains of plus or minus signs, followed by zeroes.

```
\xintiNum{+-----+-----00000000367941789479}
-367941789479
```

10.2. `\xintDouble`

f ★ `\xintDouble{N}` computes $2N$.

10.3. `\xintHalf`

f ★ `\xintHalf{N}` computes $N/2$ truncated towards zero.

10.4. `\xintInc`

f ★ `\xintInc{N}` evaluates $N+1$.

10.5. `\xintDec`

f ★ `\xintDec{N}` evaluates $N-1$.

10.6. `\xintDSL`

f ★ `\xintDSL{N}` is decimal shift left, *i.e.* multiplication by ten.

10.7. `\xintDSR`

f ★ `\xintDSR{N}` is truncated decimal shift right, *i.e.* it is the truncation of $N/10$ towards zero.

10.8. `\xintDSRr`

f ★ `\xintDSRr{N}` is rounded decimal shift right, *i.e.* it is the rounding of $N/10$ away from zero. It is needed in `xintcore` for use by `\xintiiDivRound`.

10.9. `\xintFDg`

f ★ `\xintFDg{N}` outputs the first digit (most significant) of the number.

10.10. `\xintLDg`

f ★ `\xintLDg{N}` outputs the least significant digit. When the number is positive, this is the same as the remainder in the Euclidean division by ten.

10.11. `\xintiiSgn`

f ★ `\xintiiSgn{N}` returns 1 if the number is positive, 0 if it is zero and -1 if it is negative.

10.12. `\xintiiOpp`

f ★ `\xintiiOpp{N}` outputs the opposite $-N$ of the number N .

Important note: an input such as `-\foo` is not legal, generally speaking, as argument to the macros of the `xint` bundle (except, naturally in `\xintexpr`-essions). The reason is that the minus sign stops the *f*-expansion done during parsing of the inputs. One must use the syntax `\xintiiOpp{ } \foo` if one wants to pass `-\foo` as argument to other macros.

10.13. `\xintiiAbs`

f ★ `\xintiiAbs{N}` outputs the absolute value of the number.

10.14. `\xintiiAdd`

ff ★ `\xintiiAdd{N}{M}` computes the sum of the two (big) integers.

10.15. `\xintiiCmp`

ff ★ `\xintiiCmp{N}{M}` produces 1 if $N>M$, 0 if $N=M$, and -1 if $N<M$.

At 1.21 this macro was moved from package `xint` to `xintcore`.

10.16. `\xintiiSub`

`ff` ★ `\xintiiSub{N}{M}` computes the difference $N-M$.

10.17. `\xintiiMul`

`ff` ★ `\xintiiMul{N}{M}` computes the product of two (big) integers.

10.18. `\xintiiSqr`

`f` ★ `\xintiiSqr{N}` produces the square.

10.19. `\xintiiPow`

$f \frac{\text{num}}{X}$ ★ `\xintiiPow{N}{x}` computes N^x . For $x=0$, this is 1. For $N=0$ and $x<0$, or if $|N|>1$ and $x<0$, an error is raised. There will also be an error if x exceeds the maximal ε -TeX number 2147483647, but the real limit for exponents comes from either the computation time or the settings of some TeX memory parameters.

Indeed, the maximal power of 2 which `xint` is able to compute explicitly is $2^{2^{17}}=2^{131072}$ which has 39457 digits. This exceeds the maximal size on input for the `xintcore` multiplication, hence any 2^N with a higher N will fail. On the other hand $2^{2^{16}}$ has 19729 digits, thus it can be squared once to obtain $2^{2^{17}}$ or multiplied by anything smaller, thus all exponents up to and including 2^{17} are allowed (because the power operation works by squaring things and making products).

10.20. `\xintiiFac`

$f \frac{\text{num}}{X}$ ★ `\xintiiFac{x}` computes the factorial.

The (theoretically) allowable range is $0 \leq x \leq 10000$.
However the maximal possible computation depends on the values of some memory parameters of the `tex` executable: with the current default settings of TeXLive 2015, the maximal computable factorial (a.t.t.o.w. 2015/10/06) turns out to be 5971! which has 19956 digits.

The `factorial` function, or equivalently `!` as post-fix operator is available in `\xintiexpr`, `\xintexpr`:

```
\printnumber{\xinttheiexpr 200!\relax}\par
78865786736479050355236321393218506229513597768717326329474253324435944996340334292030428401
19846239041772121389196388302576427902426371050619266249528299311134628572707633172373969889
43922445621451664240254033291864131227428294853277524242407573903240321257405579568660226031
9041703240623517008587961789222227896237038973747200000000000000000000000000000000000000
0000000
```

See also `\xintFloatFac` from package `xintfrac` for the float variant, used in `\xintfloatexpr`.

10.21. `\xintiiDivision`

`ff` ★ `\xintiiDivision{M}{N}` produces `{quotient}{remainder}`, in the sense of (mathematical) Euclidean division: $M = QN + R$, $0 \leq R < |N|$. So the remainder is always non-negative and the formula $M = QN + R$ always holds independently of the signs of N or M . Division by zero is an error (even if M vanishes) and returns `{0}{0}`.

10.22. `\xintiiQuo`

ff★ `\xintiiQuo{M}{N}` computes the quotient from the Euclidean division.

10.23. `\xintiiRem`

ff★ `\xintiiRem{M}{N}` computes the remainder from the Euclidean division.

10.24. `\xintiiDivRound`

ff★ `\xintiiDivRound{M}{N}` returns the rounded value of the algebraic quotient M/N of two big integers. The rounding is ``away from zero.''

```
\xintiiDivRound {100}{3}, \xintiiDivRound {101}{3}
33, 34
```

10.25. `\xintiiDivTrunc`

ff★ `\xintiiDivTrunc{M}{N}` computes $\text{trunc}(M/N)$. For positive arguments $M, N > 0$ it is the same as the Euclidean quotient `\xintiiQuo`.

```
\xintiiQuo{1000}{57} (Euclidean), \xintiiDivTrunc{1000}{57} (truncated),
\xintiiDivRound{1000}{57} (rounded)\newline
\xintiiQuo{-1000}{57}, \xintiiDivTrunc{-1000}{57} (t), \xintiiDivRound{-1000}{57} (r)
\newline
\xintiiQuo{1000}{-57}, \xintiiDivTrunc{1000}{-57} (t), \xintiiDivRound{1000}{-57} (r)
\newline
\xintiiQuo{-1000}{-57}, \xintiiDivTrunc{-1000}{-57} (t), \xintiiDivRound{-1000}{-57} (r)
\par
17 (Euclidean), 17 (truncated), 18 (rounded)
-18, -17 (t), -18 (r)
-17, -17 (t), -18 (r)
18, 17 (t), 18 (r)
```

10.26. `\xintiiDivFloor`

ff★ `\xintiiDivFloor{M}{N}` computes $\text{floor}(M/N)$. For positive divisor $N > 0$ and arbitrary dividend M it is the same as the Euclidean quotient `\xintiiQuo`.

```
\xintiiQuo{1000}{57} (Euclidean), \xintiiDivFloor{1000}{57} (floored)\newline
\xintiiQuo{-1000}{57}, \xintiiDivFloor{-1000}{57}\newline
\xintiiQuo{1000}{-57}, \xintiiDivFloor{1000}{-57}\newline
\xintiiQuo{-1000}{-57}, \xintiiDivFloor{-1000}{-57}\par
17 (Euclidean), 17 (floored)
-18, -18
-17, -18
18, 17
```

10.27. `\xintiiMod`

ff★ `\xintiiMod{M}{N}` computes $M - N * \text{floor}(M/N)$. For positive divisor $N > 0$ and arbitrary dividend M it is the same as the Euclidean remainder `\xintiiRem`.

Formerly, this macro computed $M - N * \text{trunc}(M/N)$. The former meaning is retained as `\xintiiMod-Trunc`.

```
\xintiiRem {1000}{57} (Euclidean), \xintiiMod {1000}{57} (floored),
\xintiiModTrunc {1000}{57} (truncated)\newline
\xintiiRem {-1000}{57}, \xintiiMod {-1000}{57}, \xintiiModTrunc {-1000}{57}\newline
```

```
\xintiiRem {1000}{-57}, \xintiiMod {1000}{-57}, \xintiiModTrunc {1000}{-57}\newline
\xintiiRem {-1000}{-57}, \xintiiMod {-1000}{-57}, \xintiiModTrunc {-1000}{-57}\par
31 (Euclidean), 31 (floored), 31 (truncated)
26, 26, -31
31, -26, 31
26, -31, -31
```

10.28. `\xintNum`

f ★ `\xintNum` is originally an alias for `\xintiNum`. But with `xintfrac` loaded its meaning is modified to accept more general inputs. It then becomes an alias to `\xintTTrunc` which truncates the general input to an integer in strict format.

11. Macros of the `xint` package

This package loads automatically `xintcore` (and `xintkernel`) hence all macros described in [section 10](#) are still available.

This is 1.4m of 2022/06/10.

Version 1.0 was released 2013/03/28. Since 1.1 2014/10/28 the core arithmetic macros have been moved to a separate package `xintcore`, which is automatically loaded by `xint`. Only the `\xintiiSum`, `\xintiiPrd`, `\xintiiSquareRoot`, `\xintiiSqrt`, `\xintiiSqrtR`, `\xintiiPFactorial`, `\xintiiBinomial` genuinely add to the arithmetic macros from `xintcore`. (`\xintiiFac` which computes factorials is already in `xintcore`.)

With the exception of `\xintLen`, of the “Boolean logic macros” (see next paragraphs) all macros require inputs being integers in strict format, see [subsection 8.4](#).⁷⁰ The `ii` in the macro names is here as a reminder of that fact. The output is an integer in strict format, or a pair of two braced such integers for `\xintiiSquareRoot`, with the exception of `\xintiiE` which may produce strings of zero's if its first argument is zero.

Macros `\xintDecSplit` and `\xintReverseDigits` are non-arithmetic and have their own specific rules.

For all macros described here for which it makes sense, package `xintfrac` defines a similar one without `ii` in its name. This will handle more general inputs: decimal, scientific numbers, fractions. The `ii` macros provided here by `xint` can be nested inside macros of `xintfrac` but the opposite does not apply, because the output format of the `xintfrac` macros, even for representing integers, is not understood by the `ii` macros. The “Boolean macros” `\xintAND` etc... are exceptions though, they work fine if served as inputs some `xintfrac` output, despite doing only *f-expansion*. Prior to 1.2o, these macros did apply the `\xintNum` or the more general `xintfrac` general parsing, but this overhead was deemed superfluous as it serves only to handle hand-written input and is not needed if the input is obtained as a nested chain of `xintfrac` macros for example.

Prior to release 1.2o, `xint` defined additional macros which applied `\xintNum` to their input arguments. All these macros were deprecated at 1.2o and have been removed at 1.3.

At 1.3d macros `\xintiiGCD` and `\xintiiLCM` from package `xintgcd` are also available from loading `xint` only. They are support macros for the (multi-arguments) functions `gcd()` and `lcm()` in `\xint-iiexpr`.

See [subsubsection 8.3.1](#) for the significance of the $\overset{\text{Num}}{f}$, f , $\overset{\text{num}}{x}$ and \star margin annotations.

.1	<code>\xintLen</code>	152	.19	<code>\xintiiEven</code>	154
.2	<code>\xintReverseDigits</code>	152	.20	<code>\xintiiMON</code>	154
.3	<code>\xintDecSplit</code>	152	.21	<code>\xintiiMMON</code>	154
.4	<code>\xintDecSplitL</code> , <code>\xintDecSplitR</code>	153	.22	<code>\xintiiifSgn</code>	154
.5	<code>\xintiiE</code>	153	.23	<code>\xintiiifZero</code>	154
.6	<code>\xintDSH</code>	153	.24	<code>\xintiiifNotZero</code>	155
.7	<code>\xintDSHr</code> , <code>\xintDSx</code>	153	.25	<code>\xintiiifOne</code>	155
.8	<code>\xintiiEq</code>	153	.26	<code>\xintiiifCmp</code>	155
.9	<code>\xintiiNotEq</code>	153	.27	<code>\xintiiifEq</code>	155
.10	<code>\xintiiGeq</code>	153	.28	<code>\xintiiifGt</code>	155
.11	<code>\xintiiGt</code>	153	.29	<code>\xintiiifLt</code>	155
.12	<code>\xintiiLt</code>	154	.30	<code>\xintiiifOdd</code>	155
.13	<code>\xintiiGtorEq</code>	154	.31	<code>\xintiiSum</code>	155
.14	<code>\xintiiLtorEq</code>	154	.32	<code>\xintiiPrd</code>	156
.15	<code>\xintiiIsZero</code>	154	.33	<code>\xintiiSquareRoot</code>	156
.16	<code>\xintiiIsNotZero</code>	154	.34	<code>\xintiiSqrt</code> , <code>\xintiiSqrtR</code>	156
.17	<code>\xintiiIsOne</code>	154	.35	<code>\xintiiBinomial</code>	156
.18	<code>\xintiiOdd</code>	154	.36	<code>\xintiiPFactorial</code>	157

⁷⁰ of course for conditionals such as `\xintiiifCmp` this constraint applies only to the first two arguments.

<p>.37 \xintiiMax..... 158</p> <p>.38 \xintiiMin..... 158</p> <p>.39 \xintiiMaxof 158</p> <p>.40 \xintiiMinof 158</p> <p>.41 \xintifTrueAelseB..... 158</p> <p>.42 \xintifFalseAelseB..... 158</p> <p>.43 \xintNOT 158</p> <p>.44 \xintAND 158</p> <p>.45 \xintOR 159</p> <p>.46 \xintXOR 159</p> <p>.47 \xintANDof..... 159</p>	<p>.48 \xintORof..... 159</p> <p>.49 \xintXORof..... 159</p> <p>.50 \xintiiGCD..... 159</p> <p>.51 \xintiiLCM..... 159</p> <p>.52 \xintiiGCDof 159</p> <p>.53 \xintiiLCMof 160</p> <p>.54 \xintLen 160</p> <p>.55 (WIP) \xintRandomDigits 160</p> <p>.56 (WIP) \xintXRandomDigits 160</p> <p>.57 (WIP) \xintiiRandRange 161</p> <p>.58 (WIP) \xintiiRandRangeAtoB..... 161</p>
--	--

11.1. \xintiLen

^{Num}
f ★ `\xintiLen{N}` returns the length of the number, after its parsing via `\xintiNum`. The count does not include the sign.

```
\xintiLen{-12345678901234567890123456789}
```

29

Prior to 1.2o, the package defined only `\xintLen`, which is extended by `xintfrac` to fractions or decimal numbers, hence acquires a bit more overhead then.

11.2. \xintReverseDigits

3.60004pt, 8.39996pt, 12.0pt

f ★ `\xintReverseDigits{N}` will reverse the order of the digits of the number. `\xintRev` is the former denomination and is kept as an alias. Leading zeroes resulting from the operation are not removed. Contrarily to `\xintReverseOrder` this macro *f-expands* its argument; it is only usable with digit tokens. It does *not* apply `\xintiNum` to its argument (so this must be done explicitly if the argument is an integer produced from some `xintfrac` macros). It does accept a leading minus sign which will be left upfront in the output.

```
\oodef\x{\xintReverseDigits
  {9876543210987654321098765432109876543210}}\meaning\x\par
\noindent\oodef\x{\xintReverseDigits {\xintReverseDigits
  {9876543210987654321098765432109876543210}}}\meaning\x\par
```

macro:->01234567890123456789012345678901234567890123456789

macro:->9876543210987654321098765432109876543210

11.3. \xintDecSplit

^{Num}
x *f* ★ `\xintDecSplit{x}{N}` cuts the *N* (a list of digits) into two pieces *L* and *R*: it outputs `{L}{R}` where the original *N* is the concatenation *LR*. These two pieces are decided according to *x*:

- for *x*>0, *R* coincides with the *x* least significant digits. If *x* equals or exceeds the length of *N* the first piece *L* will thus be *empty*,
- for *x*=0, *R* is *empty*, and *L* is all of *N*,
- for *x*<0, the first piece *L* consists of the *|x|* most significant digits and the second piece *R* gets the remaining ones. If *x* equals or exceeds the length of *N* the second piece *R* will thus be *empty*.

This macro provides public interface to some functionality which is primarily of internal interest. It operates only (after *f-expansion*) on ``strings'' of digits tokens: leading zeroes are allowed but a leading sign (even a minus sign) will provoke an error.

Breaking change with 1.2i: formerly *N*<0 was replaced by its absolute value. Now, a sign (positive or negative) will create an error.

11.12. `\xintiilt`

ff ★ `\xintiilt{N}{M}` returns 1 if $N < M$, 0 otherwise.

11.13. `\xintiiGtorEq`

ff ★ `\xintiiGtorEq{N}{M}` returns 1 if $N \geq M$, 0 otherwise. Extended by *xintfrac* to fractions.

11.14. `\xintiiltorEq`

ff ★ `\xintiiltorEq{N}{M}` returns 1 if $N \leq M$, 0 otherwise.

11.15. `\xintiiIsZero`

f ★ `\xintiiIsZero{N}` returns 1 if $N = 0$, 0 otherwise.

11.16. `\xintiiIsNotZero`

f ★ `\xintiiIsNotZero{N}` returns 1 if $N \neq 0$, 0 otherwise.

11.17. `\xintiiIsOne`

f ★ `\xintiiIsOne{N}` returns 1 if $N = 1$, 0 otherwise.

11.18. `\xintiiOdd`

f ★ `\xintiiOdd{N}` is 1 if the number is odd and 0 otherwise.

11.19. `\xintiiEven`

f ★ `\xintiiEven{N}` is 1 if the number is even and 0 otherwise.

11.20. `\xintiiMON`

f ★ `\xintiiMON{N}` computes $(-1)^N$.

```
\xintiiMON {-280914019374101929}
-1
```

11.21. `\xintiiMMON`

f ★ `\xintiiMMON{N}` computes $(-1)^{N-1}$.

```
\xintiiMMON {280914019374101929}
1
```

11.22. `\xintiiifSgn`

fnnn ★ `\xintiiifSgn{⟨N⟩}{⟨A⟩}{⟨B⟩}{⟨C⟩}` executes either the $\langle A \rangle$, $\langle B \rangle$ or $\langle C \rangle$ code, depending on its first argument being respectively negative, zero, or positive.

11.23. `\xintiiifZero`

fnn ★ `\xintiiifZero{⟨N⟩}{⟨IsZero⟩}{⟨IsNotZero⟩}` expandably checks if the first mandatory argument N (a number, possibly a fraction if *xintfrac* is loaded, or a macro expanding to one such) is zero or not. It then either executes the first or the second branch.

Beware that both branches must be present.

11.24. `\xintiifNotZero`

fnn ★ `\xintiifNotZero{⟨N⟩}{⟨IsNotZero⟩}{⟨IsZero⟩}` expandably checks if the first mandatory argument **N** is not zero or is zero. It then either executes the first or the second branch.
Beware that both branches must be present.

11.25. `\xintiifOne`

fnn ★ `\xintiifOne{⟨N⟩}{⟨IsOne⟩}{⟨IsNotOne⟩}` expandably checks if the first mandatory argument **N** is one or not one. It then either executes the first or the second branch. Beware that both branches must be present.

11.26. `\xintiifCmp`

ffnn ★ `\xintiifCmp{⟨A⟩}{⟨B⟩}{⟨A<B⟩}{⟨A=B⟩}{⟨A>B⟩}` compares its first two arguments and chooses accordingly the correct branch.

11.27. `\xintiifEq`

ffnn ★ `\xintiifEq{⟨A⟩}{⟨B⟩}{⟨A=B⟩}{⟨not(A=B)⟩}` checks equality of its two first arguments and executes the corresponding branch.

11.28. `\xintiifGt`

ffnn ★ `\xintiifGt{⟨A⟩}{⟨B⟩}{⟨A>B⟩}{⟨not(A>B)⟩}` checks if $A > B$ and executes the corresponding branch.

11.29. `\xintiifLt`

ffnn ★ `\xintiifLt{⟨A⟩}{⟨B⟩}{⟨A<B⟩}{⟨not(A<B)⟩}` checks if $A < B$ and executes the corresponding branch.

11.30. `\xintiifOdd`

fnn ★ `\xintiifOdd{⟨A⟩}{⟨A odd⟩}{⟨A even⟩}` checks if **A** is an odd integer and executes the corresponding branch.

11.31. `\xintiiSum`

**f* ★ `\xintiiSum{⟨braced things⟩}` after expanding its argument expects to find a sequence of tokens (or braced material). Each is *f*-expanded, and the sum of all these numbers is returned.

```
\xintiiSum{{123}{-98763450}{\xintiiFac{7}}{\xintiiMul{3347}{591}}}\newline
\xintiiSum{1234567890}\newline
\xintiiSum{1234}\newline
\xintiiSum{}
-96780210
45
10
0
```

A sum with only one term returns that number: `\xintiiSum {{-1234}}=-1234`. Attention that `\xintiiSum {-1234}` is not legal input and would make the $\text{T}_{\text{E}}\text{X}$ run fail.

11.32. `\xintiiPrd`

f ★ `\xintiiPrd{braced things}` after expanding its argument expects to find a sequence of (of braced items or unbraced single tokens). Each is expanded (with the usual meaning), and the product of all these numbers is returned.

```
\xintiiPrd{{-9876}}{\xintiiFac{7}}{\xintiiMul{3347}{591}}\newline
\xintiiPrd{123456789123456789}\newline
\xintiiPrd {1234}\newline
\xintiiPrd{}
```

-98458861798080
131681894400
24
1

Attention that `\xintiiPrd {-1234}` is not legal input and would make the TeX compilation fail.

```
$2^{200}3^{100}7^{100}=\printnumber
{\xintiiPrd {{\xintiiPow {2}{200}}{\xintiiPow {3}{100}}{\xintiiPow {7}{100}}}}$
```

$2^{200}3^{100}7^{100} = 2678727931661577575766279517007548402324740266374015348974459614815426412965499$
49000044400724076572713000016531207640654562118014357199401590334353924402821243896682224892
7862988084382716133376

With `xintexpr`, the syntax is the natural one:

```
$2^{200}3^{100}7^{100}=\printnumber{\xinttheiexpr 2^200 * 3^100 * 7^100\relax}$
```

$2^{200}3^{100}7^{100} = 2678727931661577575766279517007548402324740266374015348974459614815426412965499$
49000044400724076572713000016531207640654562118014357199401590334353924402821243896682224892
7862988084382716133376

11.33. `\xintiiSquareRoot`

f ★ `\xintiiSquareRoot{N}` returns two braced integers `{M}{d}` which satisfy $d > 0$ and $M^2 - d = N$ with M the smallest (hence if $N = k^2$ is a perfect square then $M = k + 1$, $d = 2k + 1$).

```
\xintAssign\xintiiSquareRoot {17000000000000000000000000}\to\A\B
\xintiiSub{\xintiiSqr\A}\B=\A\string^2-\B
```

17000000000000000000000000=4123105625618^2-2799177881924

A rational approximation to \sqrt{N} is $M - \frac{d}{2M}$ which is a majorant and the error is at most $1/2M$ (if N is a perfect square k^2 this gives $k+1/(2k+2)$, not k .)

Package `xintfrac` has `\xintFloatSqrt` for square roots of floating point numbers.

11.34. `\xintiiSqrt`, `\xintiiSqrtR`

f ★ `\xintiiSqrt{N}` computes the largest integer whose square is at most equal to N . `\xintiiSqrtR` produces the rounded, not truncated, square root.

```
\begin{itemize}[nosep]
\item \xintiiSqrt {30000000000000000000000000000000}
\item \xintiiSqrtR {30000000000000000000000000000000}
\item \xintiiSqrt {\xintiiE {3}{100}}
\end{itemize}
```

- 1732050807568877293
- 1732050807568877294
- 173205080756887729352744634150587236694280525381038

11.35. `\xintiiBinomial`

$\frac{\text{num}}{X} \frac{\text{num}}{X}$ ★ `\xintiiBinomial{x}{y}` computes binomial coefficients.

If $x < 0$ an out-of-range error is raised. Else, if $y < 0$ or if $x < y$ the macro evaluates to 0.

The allowable range is $0 \leq x \leq 99999999$. But this theoretical range includes binomial coefficients with more than the roughly 19950 digits that the arithmetics of *xint* can handle. In such cases, the computation will end up in a low-level TeX error after a long time.

It turns out that $\binom{65000}{32500}$ has 19565 digits and $\binom{64000}{32000}$ has 19264 digits. The latter can be evaluated (this takes a long long time) but presumably not the former (I didn't try). Reasonable feasible evaluations are with binomial coefficients not exceeding about one thousand digits.

The `binomial` function is available in the *xintexpr* parsers.

```
\xinttheiexpr seq(binomial(100,i), i=47..53)\relax
84413487283064039501507937600, 93206558875049876949581681100, 98913082887808032681188722800,
100891344545564193334812497256, 98913082887808032681188722800, 93206558875049876949581681100,
84413487283064039501507937600
```

See `\xintFloatBinomial` from package *xintfrac* for the float variant, used in `\xintfloatexpr`.

In order to evaluate binomial coefficients $\binom{x}{y}$ with $x > 99999999$, or even $x \geq 2^{31}$, but y is not too large, one may use an ad hoc function definition such as:

```
\xintdeffunc mybigbinomial(x,y):=`*(x-y+1..[1]..x)//y!;%
%
% without [1], x would have been limited to < 2^31
\printnumber{\xinttheexpr mybigbinomial(98765432109876543210,10)\relax}
24338098741940755592729533173058146177070669479669793038510211146784065843698581878582323710
27360575372715482389633359878460739973726786576925067784100587971261422326652270975592667517
4871960261
```

To get this functionality in macro form, one can do:

```
\xintNewIExpr\MyBigBinomial [2]{`*(#1-#2+1..[1]..#1)//#2!}
\printnumber{\MyBigBinomial {98765432109876543210}{10}}
24338098741940755592729533173058146177070669479669793038510211146784065843698581878582323710
27360575372715482389633359878460739973726786576925067784100587971261422326652270975592667517
4871960261
```

As we used `\xintNewIExpr`, this macro will only accept strict integers. Had we used `\xintNewExpr` the `\MyBigBinomial` would have accepted general fractions or decimal numbers, and computed the product at the numerator without truncating them to integers; but the factorial at the denominator would truncate its argument.

11.36. `\xintiiPFactorial`

$\frac{\text{num}}{x} \frac{\text{num}}{x}$ ★ `\xintiiPFactorial{a}{b}` computes the partial factorial $(a+1)(a+2)\dots b$. For $a=b$ the product is considered empty hence returns 1.

The allowed range is $-100000000 \leq a, b \leq 99999999$. The rule is to interpret the formula as the product of the j 's such that $a < j \leq b$, hence in particular if $a \geq b$ the product is empty and the macro evaluates to 1.

Only for $0 \leq a \leq b$ is the behaviour to be considered stable. For $a > b$ or negative arguments, the definitive rules have not yet been fixed.

```
\xintiiPFactorial {100}{130}
69293021885203871012298422845822803287591970060789350400000000
```

This theoretical range allows computations whose result values would have more than the roughly 19950 digits that the arithmetics of *xint* can handle. In such cases, the computation will end up in a low-level TeX error after a long time.

The `pfactorial` function is available in the *xintexpr* parsers.

```
\xinttheiexpr pfactorial(100,130)\relax
69293021885203871012298422845822803287591970060789350400000000
```

See `\xintFloatPFactorial` from package *xintfrac* for the float variant, used in `\xintfloatexpr`.

In case values are needed with $b > 99999999$, or even $b \geq 2^{31}$, but $b - a$ is not too large, one may use an ad hoc function definition such as:

```
\xintdeffunc mybigpfac(a,b):=`*(a+1..[1]..b);%
%
% without [1], b would have been limited to < 2^31
```

```
\printnumber{\xinttheexpr mybigpfac(98765432100,98765432120)\relax}
78000855017567528067298107313023778438653002029049647467208196028116499434050587656870489322
99630604482236853566403912561449912587404607844104078121472675461815442734098676283450069933
322948600573016997034009566576640000
```

11.37. `\xintiiMax`

ff ★ `\xintiiMax{N}{M}` returns the largest of the two in the sense of the order structure on the relative integers (*i.e.* the right-most number if they are put on a line with positive numbers on the right):
`\xintiiMax {-5}{-6}=-5.`

11.38. `\xintiiMin`

ff ★ `\xintiiMin{N}{M}` returns the smallest of the two in the sense of the order structure on the relative integers (*i.e.* the left-most number if they are put on a line with positive numbers on the right): `\xintiiMin {-5}{-6}=-6.`

11.39. `\xintiiMaxof`

f → * *f* ★ `\xintiiMaxof{{a}{b}{c}...}` returns the maximum. The list argument may be a macro, it is *f*-expanded first.

11.40. `\xintiiMinof`

f → * *f* ★ `\xintiiMinof{{a}{b}{c}...}` returns the minimum. The list argument may be a macro, it is *f*-expanded first.

11.41. `\xintifTrueAelseB`

fnn ★ `\xintifTrueAelseB{<f>}{<true branch>}{<false branch>}` is a synonym for `\xintiiifNotZero`.
`\xintiiifnotzero` is lowercase companion macro.

Note 1: as it does only *f*-expansion on its argument it fails with inputs such as `--0`. But with `xintfrac` loaded, it does work fine if nested with other `xintfrac` macros, because the output format of such macros is fine as input to `\xintiiifNotZero`. This remark applies to all other “Boolean logic” macros next.

Note 2: prior to 1.20 this macro was using `\xintifNotZero` which applies `\xintNum` to its argument (or gets redefined by `xintfrac` to handle general decimal numbers or fractions). Hence it would have worked with input such as `--0`. But it was decided at 1.20 that the overhead was not worth it. The same remark applies to the other “Boolean logic” type macros next.

11.42. `\xintifFalseAelseB`

fnn ★ `\xintifFalseAelseB{<f>}{<false branch>}{<true branch>}` is a synonym for `\xintiiifZero`.
`\xintiiifzero` is lowercase companion macro.

11.43. `\xintNOT`

f ★ `\xintNOT` is a synonym for `\xintiiIsZero`.
`\xintiiiszero` serves as lowercase companion macro.

11.44. `\xintAND`

ff ★ `\xintAND{f}{g}` returns 1 if `f!=0` and `g!=0` and 0 otherwise.

11.45. `\xintOR`

ff ★ `\xintOR{f}{g}` returns 1 if $f \neq 0$ or $g \neq 0$ and 0 otherwise.

11.46. `\xintXOR`

ff ★ `\xintXOR{f}{g}` returns 1 if exactly one of f or g is true (i.e. non-zero), else 0.

11.47. `\xintANDof`

f → **f* ★ `\xintANDof{{a}{b}{c}...}` returns 1 if all are true (i.e. non zero) and 0 otherwise. The list argument may be a macro, it (or rather its first token) is *f-expanded* first to deliver its items.

11.48. `\xintORof`

f → **f* ★ `\xintORof{{a}{b}{c}...}` returns 1 if at least one is true (i.e. does not vanish), else it produces 0. The list argument may be a macro, it is *f-expanded* first.

11.49. `\xintXORof`

f → **f* ★ `\xintXORof{{a}{b}{c}...}` returns 1 if an odd number of them are true (i.e. do not vanish), else it produces 0. The list argument may be a macro, it is *f-expanded* first.

11.50. `\xintiiGCD`

ff ★ `\xintiiGCD{N}{M}` computes the greatest common divisor. It is positive, except when both N and M vanish, in which case the macro returns zero.

```
\xintiiGCD{10000}{1113}=1
\xintiiGCD{123456789012345}{9876543210321}=3
```

At 1.3d, this macro (which is used by the `gcd()` function in `\xintiiexpr`) was copied over to `xint`, thus removing a partial dependency of `xintexpr` on `xintgcd`.

At 1.4 `xintgcd` requires `xint` and the latter is thus the one providing the macro.

11.51. `\xintiiLCM`

ff ★ `\xintiiLCM{N}{M}` computes the least common multiple. It is positive, except if one of N or M vanish, in which case the macro returns zero.

```
\xintiiLCM{10000}{1113}=11130000
\xintiiLCM{123456789012345}{9876543210321}=406442103762636081733470915
```

At 1.3d, this macro (which is used by the `lcm()` function in `\xintiiexpr`) was copied over to `xint`, thus removing a partial dependency of `xintexpr` on `xintgcd`.

At 1.4 `xintgcd` requires `xint` and the latter is thus the one providing the macro.

11.52. `\xintiiGCDof`

f → **f* ★ `\xintiiGCDof{{a}{b}{c}...}` computes the greatest common divisor of the integers a , b , \dots . It is a support macro for the `gcd()` function of the `\xintiiexpr` parser.

It replaces the `\xintGCDof` which was formerly provided by `xintgcd` and is now available via `xintfrac` in a version handling also fractions.

11.53. `\xintiiLCMof`

$f \rightarrow *f$ ★ `\xintiiLCMof{{a}{b}{c}...}` computes the least common multiple of the integers `a`, `b`, ... It is a support macro for the `lcm()` function of the `\xintiexpr` parser.

It replaces the `\xintLCMof` which was formerly provided by `xintgcd` and is now available via `xintfrac` in a version handling also fractions.

11.54. `\xintLen`

$\frac{\text{num}}{f}$ ★ `\xintLen` is originally an alias for `\xintiLen`. But with `xintfrac` loaded its meaning is modified to accept more general inputs.

11.55. (WIP) `\xintRandomDigits`

All randomness related macros are Work-In-Progress: implementation and user interface may change. They work only if the TeX engine provides the `\uniformdeviate` or `\pdfuniformdeviate` primitive. See `\xintUniformDeviate` for additional information.

$\frac{\text{num}}{X}$ ★ `\xintRandomDigits{N}` expands in two steps to `N` random decimal digits. The argument must be non-negative and is limited by TeX memory parameters. On TeXLive 2018 with input save stack size at 5000 the maximal allowed `N` is at most 19984 (tested within a `\write` to an auxiliary file, the macro context may cause a reduced maximum).

```
\pdfsetrandomseed 271828182
```

```
\xintRandomDigits{92}
```

```
60033782389146151207277993539344280578090871919638745398735577686436165769394958639376355806
```

TeX-hackers note: the digits are produced eight by eight by the same method which would result from `\xintUniformDeviate{10000000}` but with less overhead.

11.56. (WIP) `\xintXRandomDigits`

$\frac{\text{num}}{X}$ ☆ `\xintXRandomDigits{N}` expands under exhaustive expansion (`\edef`, `\write`, `\csname` ...) to `N` random decimal digits. The argument must be non-negative. For example:

```
\newwrite\out
\immediate\openout\out=\jobname-out.txt
\immediate\write\out{\xintXRandomDigits{4500000}}
\immediate\closeout\out
```

creates a 4500001 bytes file (it ends with a line feed character). Trying with 5000000 raises this error:

```
Runaway text?
588875947168511582764514135070217555354479805240439407753451354223283\ETC.
! TeX capacity exceeded, sorry [main memory size=5000000].
<inserted text> 666515098
```

```
1.15 ...ate\write\out{\xintXRandomDigits{5000000}}
```

```
No pages of output.
Transcript written on temp.log.
```

This can be lifted by increasing the TeX memory settings (installation dependent).

TeX-hackers note: the digits are produced eight by eight by the same method which would result from `\xintUniformDeviate{10000000}` but with less overhead.

11.57. (WIP) `\xintiiRandRange`

f ★ `\xintiiRandRange{A}` expands to a random (big) integer N such that $0 \leq N < A$. It is a supporting macro for `randrange()`. As with Python's function of the same name, it is an error if $A \leq 0$.

```
\pdfsetrandomseed 271828314
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx\newline
\xintiiRandRange{\xintNum{1e40}}\newline
\pdfsetrandomseed 271828314
\xinttheiexpr randrange(num(1e40))\relax\newline
% bare 1e40 not understood by \xintiiexpr
\pdfsetrandomseed 271828314
\xinttheexpr randrange(1e40)\relax
```

```
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
1408107837990425263001878034077495278697
1408107837990425263001878034077495278697
1408107837990425263001878034077495278697
```

Of course, keeping in mind that the set of seeds is of cardinality 2^{28} , randomness is a bit illusory here say with $A=10^N$, $N>8$, if we proceed immediately after having set the seed. If we add some entropy in any way, then it is slightly more credible; but I think that for each seed the period is something like $2^{27}(2^{55}-1)55$,⁷¹ so we expect at most about $2^{110}55$ 'points in time', and this is already small compared to the 10^{40} from example above. Thus already we are very far from being intrinsically able to generate all numbers with forty digits as random numbers, and this makes the previous section about usage of `\xintXRandDigits` to generate millions of digits a bit comical...

T_EX-hackers note: the digits are produced eight by eight by the same method which would result from `\xintUniformDeviate{10000000}` but with less overhead.

11.58. (WIP) `\xintiiRandRangeAtoB`

ff ★ `\xintiiRandRangeAtoB{A}{B}` expands to a random (big) integer N such that $A \leq N < B$. It is a supporting macro for `randrange()`. As with Python's function of the same name, it is an error if $B \leq A$.

```
\pdfsetrandomseed 271828314
12345678911111111111111111\newline
\xintiiRandRangeAtoB{12345678911111111111111111}{123456789222222222222222222222}%
\newline
\pdfsetrandomseed 271828314
\def\test{%
\xinttheiexpr
    randrange(1234567891111111111111111111,123456789222222222222222222222)
\relax}%
\romannumerals\xintreplicate{10}{\test\newline}%
1234567892222222222222222222222222222
```

```
12345678911111111111111111111111111
12345678916037426188606389808
12345678916037426188606389808
12345678916060337223949101536
12345678912190033095886250034
12345678917323740152668511995
12345678915424847208552293485
12345678921595726610650510660
```

⁷¹ Compare the result of exercise 3.2.2-30 in TAOCP, vol II.

TOC, xint bundle, xintkernel, xintcore, xint, xintfrac, xintbinhex, xintgcd, xintseries, xintfrac

12345678911673261982088192858
12345678911339325803675947159
12345678917791540296982027151
12345678913602899909728811895
123456789222222222222222222222

T_EX-hackers note: the digits are produced eight by eight by the same method which would result from `\xint-UniformDeviate{10000000}` but with less overhead.

12. Macros of the **xintfrac** package

First version of this package was in release 1.03 (2013/04/14) of the **xint** bundle.

At release 1.3 (2018/02/28) the behaviour of `\xintAdd` (and of `\xintSub`) was modified: when adding a/b and c/d they will use always the least common multiple of the denominators. This helps limit the build-up of denominators, but the author still hesitates if the fraction should be reduced to smallest terms. The current method allows (for example when multiplying two polynomials) to keep a well-predictable denominator among various terms, even though some may be reducible.

xintfrac loads automatically **xintcore** and **xint** and inherits their macro definitions. Only these two are redefined: `\xintNum` and `\xintLen`. As explained in subsection 8.4 and subsection 8.5 the interchange format for the **xintfrac** macros, i.e. $A/B[N]$, is not understood by the **ii**-named macros of **xintcore**/**xint** which expect the so-called strict integer format. Hence, to use such an **ii**-macro with an output from an **xintfrac** macro, an extra `\xintNum` wrapper is required. But macros already defined by **xintfrac** cover most use cases hence this should be a rarely needed.

Frac
f

num
x

In the macro descriptions, the variable **f** and the margin indicator stand for the **xintfrac** input format for integers, scientific numbers, and fractions as described in subsection 8.4.

As in the **xint.sty** documentation, **x** stands for something which internally will be handled in a `\numexpr`. It may thus be an expression as understood by `\numexpr` but its evaluation and intermediate steps must obey the \TeX bound.

The output format for most macros is the $A/B[N]$ format but naturally the float macros use the scientific notation on output. And some macros are special, for example `\xintTrunc` produces decimal numbers, `\xintIrr` produces an A/B with no $[N]$, `\xintiTrunc` and `\xintiRound` produce integers without trailing $[N]$ either, etc. . .

At 1.4g, old legacy typesetting macros `\xintFrac`, `\xintSignedFrac`, `\xintFwOver` and `\xintSignedFwOver` were renamed into `\xintTeXFrac`, `\xintTeXsignedFrac`, `\xintTeXOver`, `\xintTeXsignedOver`.

Changed
at 1.4m!

The old names will raise errors and will be removed completely soon.

.1	<code>\xintTeXFromSci</code>	164	.26	<code>\xintFloor</code>	175
.2	<code>\xintTeXFrac</code>	165	.27	<code>\xintCeil</code>	175
.3	<code>\xintTeXsignedFrac</code>	165	.28	<code>\xintiTrunc</code>	175
.4	<code>\xintTeXOver</code>	166	.29	<code>\xintTTrunc</code>	176
.5	<code>\xintTeXsignedOver</code>	166	.30	<code>\xintiRound</code>	176
.6	<code>\xintLen</code>	166	.31	<code>\xintiFloor</code>	176
.7	<code>\xintNum</code>	166	.32	<code>\xintiCeil</code>	176
.8	<code>\xintRaw</code>	166	.33	<code>\xintE</code>	177
.9	<code>\xintRawBraced</code>	167	.34	<code>\xintCmp</code>	177
.10	<code>\xintNumerator</code>	167	.35	<code>\xintEq</code>	177
.11	<code>\xintDenominator</code>	167	.36	<code>\xintNotEq</code>	177
.12	<code>\xintRawWithZeros</code>	167	.37	<code>\xintGeq</code>	177
.13	<code>\xintREZ</code>	167	.38	<code>\xintGt</code>	177
.14	<code>\xintIrr</code>	168	.39	<code>\xintLt</code>	177
.15	<code>\xintPIrr</code>	168	.40	<code>\xintGtorEq</code>	177
.16	<code>\xintJrr</code>	168	.41	<code>\xintLtorEq</code>	177
.17	<code>\xintPRaw</code>	168	.42	<code>\xintIsZero</code>	177
.18	<code>\xintDecToStringREZ</code>	168	.43	<code>\xintIsNotZero</code>	177
.19	<code>\xintDecToString</code>	169	.44	<code>\xintIsOne</code>	177
.20	<code>\xintFracToSci</code>	170	.45	<code>\xintOdd</code>	177
.21	<code>\xintFracToDecimal</code>	171	.46	<code>\xintEven</code>	178
.22	<code>\xintTrunc</code>	171	.47	<code>\xintifSgn</code>	178
.23	<code>\xintXTrunc</code>	172	.48	<code>\xintifZero</code>	178
.24	<code>\xintTFrac</code>	174	.49	<code>\xintifNotZero</code>	178
.25	<code>\xintRound</code>	175	.50	<code>\xintifOne</code>	178

<p>.51 \xintifOdd..... 178</p> <p>.52 \xintifCmp..... 178</p> <p>.53 \xintifEq..... 178</p> <p>.54 \xintifGt..... 178</p> <p>.55 \xintifLt..... 178</p> <p>.56 \xintifInt..... 179</p> <p>.57 \xintSgn..... 179</p> <p>.58 \xintSignBit..... 179</p> <p>.59 \xintOpp..... 179</p> <p>.60 \xintAbs..... 179</p> <p>.61 \xintAdd..... 179</p> <p>.62 \xintSub..... 179</p> <p>.63 \xintMul..... 179</p> <p>.64 \xintDiv..... 179</p> <p>.65 \xintDivFloor..... 179</p> <p>.66 \xintMod..... 180</p> <p>.67 \xintDivMod..... 180</p> <p>.68 \xintDivTrunc..... 180</p> <p>.69 \xintModTrunc..... 180</p> <p>.70 \xintDivRound..... 180</p> <p>.71 \xintSqr..... 180</p> <p>.72 \xintPow..... 180</p> <p>.73 \xintFac..... 181</p> <p>.74 \xintBinomial..... 181</p> <p>.75 \xintPFactorial..... 181</p> <p>.76 \xintMax..... 181</p>	<p>.77 \xintMin..... 181</p> <p>.78 \xintMaxof..... 181</p> <p>.79 \xintMinof..... 182</p> <p>.80 \xintSum..... 182</p> <p>.81 \xintPrd..... 182</p> <p>.82 \xintGCD..... 182</p> <p>.83 \xintLCM..... 182</p> <p>.84 \xintGCDof..... 182</p> <p>.85 \xintLCMof..... 182</p> <p>.86 \xintDigits, \xinttheDigits..... 183</p> <p>.87 \xintSetDigits..... 183</p> <p>.88 \xintFloat..... 184</p> <p>.89 \xintFloatBraced..... 185</p> <p>.90 \xintFloatToDecimal..... 185</p> <p>.91 \xintPFloat..... 186</p> <p>.92 \xintFloatAdd..... 190</p> <p>.93 \xintFloatSub..... 190</p> <p>.94 \xintFloatMul..... 190</p> <p>.95 \xintFloatDiv..... 190</p> <p>.96 \xintFloatPow..... 190</p> <p>.97 \xintFloatPower..... 190</p> <p>.98 \xintFloatSqrt..... 191</p> <p>.99 \xintFloatFac..... 192</p> <p>.100 \xintFloatBinomial..... 192</p> <p>.101 \xintFloatPFactorial..... 192</p>
---	--

12.1. \xintTeXFromSci

☆ Experimental. This typesetting math-mode-only macro expects an input which is already in, or will expand to, decimal or scientific notation. A trailing /B is accepted and will be handled differently according to whether it follows some scientific exponent eN part or not.

It was formerly `\xintTeXfromSci`. Old name deprecated at 1.41. Also it used to be *f-expandable* but is now only *x-expandable*. Use `\expanded` if needed.

This macro can be used as a typesetting wrapper for `\xinteval` or `\xintfloateval` output: it expects its input (after expansion) to have been already “prettified” for example via the removal of trailing zeros, usage of fixed point notation if scientific exponent is small, etc... It simply transforms the `e<exponent>` part, if actually present, into `\cdot 10^{exponent}`. A fractional part /B if found in the expansion of the input must be last and will be transformed into `\cdot B^{1}` if there was a scientific part, else the output will be using `\frac{A}{B}` (or the \TeX equivalent in place of `\frac`) with A the numerator.

\TeX -hackers note:

- I am hesitating whether the `\frac{A}{B}` branch choice should require A to be an integer, or will also, as currently, be done with A being a number in decimal notation. Please advise.
- The package does:

```
\ifdefined\frac
\protected\def\xintTeXFromScifracmacro#1#2{\frac{#2}{#1}}%
\else
\protected\def\xintTeXFromScifracmacro#1#2{{#2\over#1}}%
\fi
```

Customize as desired. Notice the interversion of arguments.

Example:

```
\xintTeXFromSci{\xintfloateval{1.1^10000/5}}$,
\xintTeXFromSci{\xinteval{1.1^10000/5}}$\par
1.689980050240070 · 10413, 8.449900251200348 · 10413 · 5-1
```

The above examples are in the case of a single numerical value. To handle more complex outputs from `\xinteval` or `\xintfloateval` one will need to proceed via a redefinition of `\xintfloatexprPrintOne` and/or `\xintexprPrintOne` like this:

```
[\def\xintfloatexprPrintOne[#1]#2{\xintTeXFromSci{\xintPFloat[#1]{#2}}
\xintfloateval[10]{2^100, 3^100, 13^100}\}
2.1267650600 · 1030, 5.153775207 · 1047, 2.479335111 · 10111
```

```
[\def\xintexprPrintOne#1{\xintTeXFromSci{\xintFracToSci{#1}}
\xinteval{sqrt(2^101,60), 355/113, 6.02e23/1000}\}
```

```
1.59226291813144314115595358963043315049844681269444074447413 · 1015,  $\frac{355}{113}$ , 6.02 · 1023 · 1000-1
```

This will however make then impossible, due to the added \TeX mark-up in the output, the nesting of `\xintfloateval`/`\xinteval` inside one another. The core `\xintexpr... \relax` syntax remains usable and is anyhow the recommended way for such nesting as it is more efficient.

Some similar effect is also possible without `\xintTeXFromSci`, simply by a customization of `\xintPFloatE` like this:

```
\begingroup
\def\xintPFloatE#1.{\cdot10^{#1}.}%
\xintfloateval{1.1^10000/5}$, $\xinteval{1.1^10000/5}$
\endgroup\newline
```

```
1.689980050240070 · 10413, 8.449900251200348 · 10413/5
```

This method is simpler-minded but will leave the trailing `/B`'s "as is", even if the numerator has no scientific exponent part. The presence of extra \TeX mark-up in the output has the consequences on nesting which were mentioned above.

12.2. `\xintTeXFrac`

$\frac{f}{f}$ ★ This is a typesetting \TeX only macro, math mode only as it expands to `\frac{A}{B}10^n` for an input which ends up parsed into raw format `A/B[n]`.

If the denominator `B` is 1, the output is `A\cdot 10^n`. If the exponent `n` is 0, the `[\cdot]10^n` part is omitted.

```
\xintTeXFrac {178.000/25600000}$, $\xintTeXFrac {178.000/1}$,
\xintTeXFrac {3.5/5.7}$\newline
```

```
 $\frac{178000}{25600000}10^{-3}$ ,  $178000 \cdot 10^{-3}$ ,  $\frac{35}{57}$ 
```

The input, if in a fraction form, is not simplified in any way, except for transforming numerator and denominator into integers, separating a power of ten part. Macros such as `\xintIrr`, `\xintPIrr`, `\xintREZ` can be inserted to wrap the input and help simplify it. The minus sign ends up in the numerator.

Changed at 1.4m! It is the new name since 1.4g of `\xintFrac`. The old name now raises a \TeX error.

12.3. `\xintTeXsignedFrac`

$\frac{f}{f}$ ★ This is as `\xintTeXFrac` except that a negative fraction has the sign ending up in front, not in the numerator.

```
\xintTeXFrac {-355/113}=\xintTeXsignedFrac {-355/113}$\newline
```

```
 $\frac{-355}{113} = -\frac{355}{113}$ 
```

Changed at 1.4m! It is the new name since 1.4g of `\xintSignedFrac`. The old name now raises a \TeX error.

12.4. `\xintTeXOver`

$\frac{f}{f}$ ★ This does the same as `\xintTeXFrac` except that the `\over` primitive is used for the fraction (in case the denominator is not one; and a pair of braces contains the `A\over B` part).

```
\xintTeXOver {178.000/25600000}$, \xintTeXOver {178.000/1}$,
\xintTeXOver {3.5/5.7}$\newline
```

$\frac{178000}{25600000} 10^{-3}$, $178000 \cdot 10^{-3}$, $\frac{35}{57}$

Changed at 1.4m! It is the new name since 1.4g of `\xintFwOver`. The old name now raises a TeX error.

12.5. `\xintTeXsignedOver`

$\frac{f}{f}$ ★ This is as `\xintTeXOver` except that a negative fraction has the sign put in front, not in the numerator.

```
\xintTeXOver{-355/113}=\xintTeXsignedOver{-355/113}$\newline
```

$-\frac{355}{113} = -\frac{355}{113}$

Changed at 1.4m! It is the new name since 1.4g of `\xintSignedFwOver`. The old name now raises a TeX error.

12.6. `\xintLen`

$\frac{f}{f}$ ★ The `\xintLen` macro from *xint* is extended to accept a fraction on input: the length of `A/B[n]` is the length of `A` plus the length of `B` plus the absolute value of `n` and minus one (an integer input as `N` is internally represented in a form equivalent to `N/1[0]` so the minus one means that the extended `\xintLen` behaves the same as the original for integers).

```
\xintLen{201710/298219}=\xintLen{201710}+\xintLen{298219}-1\newline
\xintLen{1234/1}=\xintLen{1234}=\xintLen{1234[0]}=\xintiLen{1234}\newline
\xintLen{-1e3/5.425} (\xintRaw {-1e3/5.425})\par
```

11=6+6-1

4=4=4=4

10 (-1/5425[6])

The length is computed on the `A/B[n]` which would have been returned by `\xintRaw`, as illustrated by the last example above.

`\xintLen` is only for use with such (scientific) numbers or fractions. See also `\xintNthElt` from *xinttools*. See also `\xintLength` (which however does not expand its argument) from *xintkernel* for counting more general tokens (or rather braced items).

12.7. `\xintNum`

$\frac{f}{f}$ ★ The `\xintNum` from *xint* is transformed into a synonym to `\xintTTrunc`.

Attention that for example `\xintNum{1e100000}` expands to the needed 100001 digits...

The original `\xintNum` from *xintcore* which does not understand the fraction slash or the scientific notation is still available under the name `\xintiNum`.

12.8. `\xintRaw`

$\frac{f}{f}$ ★ This macro 'prints' the fraction `f` as it is received by the package after its parsing and expansion, in a form `A/B[N]` equivalent to the internal representation: the denominator `B` is always strictly positive and is printed even if it has value 1.

```
\xintRaw{\the\numexpr 571*987\relax.123e-10/\the\numexpr-201+59\relax e-7}
-563577123/142[-6]
```

No simplification is done, not even of common zeroes between numerator and denominator:

```
\xintRaw {178000/25600000}
```

178000/25600000[0]

12.9. `\xintRawBraced`

$\frac{f}{f}$ ★ This macro expands and parses its input `f` as all `xintfrac` macros and produces as output `{N}{A}{B}` (with TeX braces) where `\xintRaw` would have returned `A/B[N]`.

12.10. `\xintNumerator`

$\frac{f}{f}$ ★ The input data is parsed as if by `\xintRaw` into `A/B[N]` and then `A` is returned if `N≤0`, or `A` extended by `N` zeroes if `N>0`.

```
\xintNumerator {178000/25600000[17]}\newline
\xintNumerator {312.289001/20198.27}\newline
\xintNumerator {178000e-3/256e5}\newline
\xintNumerator {178.000/25600000}
```

```
1780000000000000000000000000000000
312289001
178000
178000
```

12.11. `\xintDenominator`

$\frac{f}{f}$ ★ The input data is parsed as if by `\xintRaw` into `A/B[N]` and then `B` is returned if `N>0`, or `B` extended by `|N|` zeroes if `N≤0`.

```
\xintDenominator {178000/25600000[17]}\newline
\xintDenominator {312.289001/20198.27}\newline
\xintDenominator {178000e-3/256e5}\newline
\xintDenominator {178.000/25600000}
```

```
25600000
20198270000
25600000000
25600000000
```

12.12. `\xintRawWithZeros`

$\frac{f}{f}$ ★ This macro parses the input and outputs `A'/B'`, with `A'` as would be returned by `\xintNumerator{f}` and `B'` as would be returned by `\xintDenominator{f}`.

```
\xintRawWithZeros{178000/25600000[17]}\newline
\xintRawWithZeros{312.289001/20198.27}\newline
\xintRawWithZeros{178000e-3/256e5}\newline
\xintRawWithZeros{178.000/25600000}\newline
\xintRawWithZeros{\the\numexpr 571*987\relax.123e-10/\the\numexpr -201+59\relax e-7}
```

```
1780000000000000000000000000000000/25600000
312289001/20198270000
178000/25600000000
178000/25600000000
-563577123/142000000
```

12.13. `\xintREZ`

$\frac{f}{f}$ ★ The input is first parsed into `A/B[N]` as by `\xintRaw`, then trailing zeroes of `A` and `B` are suppressed and `N` is accordingly adjusted.

```
\xintREZ {178000/25600000[17]}
178/256[15]
```

12.14. `\xintIrr`

Frac
f ★ This puts the fraction into its unique irreducible form:

```
\xintIrr {178.256/256.1780}, \xintIrr {178000/2560000[17]}
```

6856/9853, 69531250000000/1

The current implementation does not cleverly first factor powers of 2 and 5, and `\xintIrr {2/3 [100]}` will execute the Euclidean division of $2 \cdot 10^{100}$ by 3, which is a bit stupid as it could have known that the 100 trailing zeros can not bring any divisibility by 3.

Starting with release 1.08, `\xintIrr` does not remove the trailing /1 when the output is an integer. This was deemed better for various (questionable?) reasons, anyway the output format is since always A/B with B>0, even in cases where it turns out that B=1. Use `\xintPraw` on top of `\xintIrr` if it is needed to get rid of such a trailing /1.

12.15. `\xintPIrr`

Frac
f ★ This puts the fraction into irreducible form, *keeping as is the decimal part [N] from raw internal A/B[N] format.* (P stands here for *Partial*)

```
\xintPIrr {178.256/256.1780}, \xintPIrr {178000/2560000[17]}
```

3428/49265[1], 89/12800[17]

Notice that the output always has the ending [N], which is exactly the opposite of `\xintIrr`'s behaviour. The interest of this macro is mainly in handling fractions which somehow acquired a big [N] (perhaps from input in scientific notation) and for which the reduced fraction would have a very large number of digits. This large number of digits can considerably slow-down computations done afterwards.

For example package `polexpr` uses `\xintPIrr` when differentiating a polynomial, or in setting up a Sturm chain for localization of the real roots of a polynomial. This is relevant to polynomials whose coefficients were input in decimal notation, as this automatically creates internally some [N]. Keeping and combining those [N]'s during computations significantly increases their speed.

12.16. `\xintJrr`

Frac
f ★ This also puts the fraction into its unique irreducible form:

```
\xintJrr {178.256/256.178}
```

6856/9853

This is (supposedly, not tested for ages) faster than `\xintIrr` for fractions having some big common factor in the numerator and the denominator.

```
\xintJrr {\xintiiPow{\xintiiFac {15}}{3}/%
\xintiiPrd{\xintiiFac{10}}{\xintiiFac{30}}{\xintiiFac{5}}}
```

1001/51705840

But to notice the difference one would need computations with much bigger numbers than in this example. As `\xintIrr`, `\xintJrr` does not remove the trailing /1 from a fraction reduced to an integer.

12.17. `\xintPraw`

Frac
f ★ `Praw` stands for ``pretty raw''. It does like `\xintRaw` apart from removing the [N] part if N=0 and removing the B if B=1.

```
\xintPraw {123e10/321e10}, \xintPraw {123e9/321e10}, \xintPraw {\xintIrr{861/123}}
```

123/321, 123/321[-1], 7

12.18. `\xintDecToStringREZ`

Frac
f ★ `\xintDecToStringREZ` uses fixed point (decimal) notation for the output. The `REZ` means that it

trims (REmoves) trailing Zeros. The name is a bit strange, because it its not limited to *decimal numbers* but accepts the same kind of inputs as most other `xintfrac` macros. The parsing of this input transforms it first into an internal format having a numerator **A**, a denominator **B** and a power of ten exponent **N**, standing for the fraction **A/B** times 10 to the power **N**. Then the following recipe applies:

- the zero value is printed as 0 (no decimal point).
- trailing zeros of **A** and **B** are removed and **N** is adjusted,
- if the new **B** is not 1, it will appear in the output as **/B**,
- fixed point decimal notation is used for **AeN**:
 - if **N** is non-negative, the output is an integer with **N** trailing zeros (and no decimal mark)
 - if **N** is negative a decimal point is used, and if **AeN** is less than one in absolute value, output will start with 0. (i.e. a decimal mark).

The following should be noted:

1. the fraction **AeN/B** or even **A/B** is not pre-reduced into lowest terms,
2. the macro does not check if **B** contains only powers of 2 and 5, so $1/2$ is printed as $1/2$, not as 0.5.

The definitive behaviour remains to be decided regarding these last two points.

```
\xintDecToStringREZ{0}, \xintDecToStringREZ{1/2}, \xintDecToStringREZ{0.5000}\newline
\xintDecToStringREZ{1.23456789e5}, \xintDecToStringREZ {1.23456789e-3}\newline
\xintDecToStringREZ{12345e-1}, \xintDecToStringREZ {12345e-2},
\xintDecToStringREZ{12345e-3}\newline
\xintDecToStringREZ{12345e-4}, \xintDecToStringREZ {12345e-5},
\xintDecToStringREZ{12345e-6}\newline
\xintDecToStringREZ{1.234567890000e12}, \xintDecToStringREZ{1.23456000e-5/10}\newline
\xintDecToStringREZ{70/14} % is not reduced to lowest terms
```

0, 1/2, 0.5

123456.789, 0.00123456789

1234.5, 123.45, 12.345

1.2345, 0.12345, 0.012345

1234567890000, 0.00000123456

70/14

See `\xintFloatToDecimal` for a variant which first rounds the input to some given number of significant digits.

12.19. `\xintDecToString`

Frac f ★ `\xintDecToString` uses fixed point notation for the output. Its behaviour remains somewhat undecided in so far as whether it should identify inputs which correspond to decimal numbers (i.e. fractions with only powers of two and five in their denominator, once reduced to lowest terms).

As with `\xintDecToStringREZ`, the name is a bit strange as inputs are in no way limited to decimal numbers but are of the most general type accepted by the `xintfrac` macros.

It is the same macro as `\xintDecToStringREZ` except that it does not remove trailing zeros, in fact `\xintDecToStringREZ{f}` is defined as `\xintDecToString{\xintREZ{f}}`.

```
\xintDecToString{0}, \xintDecToString{1/2}, \xintDecToString{0.5000}\newline
\xintDecToString{1.23456789e5}, \xintDecToString {1.23456789e-3}\newline
\xintDecToString{12345e-1}, \xintDecToString {12345e-2}, \xintDecToString{12345e-3}%
\newline
\xintDecToString{12345e-4}, \xintDecToString {12345e-5}, \xintDecToString{12345e-6}%
\newline
\xintDecToString{1.234567890000e12}, \xintDecToString{1.23456000e-5/10}\newline
```

```
\xintDecToString{70/14}
0, 1/2, 0.5000
123456.789, 0.00123456789
1234.5, 123.45, 12.345
1.2345, 0.12345, 0.012345
1234567890000, 0.0000123456000/10
70/14
```

Since 1.4e, `\xintDecToString` is the default for `\xintiexprPrintOne`, which governs the `\xintieval` output format: in this use case there is never a `/B` fractional part and the output is always either an integer (if `\xintieval` was used without optional argument) or a decimal string

```
\def\xintiexprPrintOne{\xintDecToString}
```

Any replacement of `\xintDecToString` as the expansion of `\xintiexprPrintOne` should obey the following blueprint:

- ☆ • to be expandable, but not necessarily *f-expandable*,
- to accept on input `A` or `A[N]`.

12.20. `\xintFracToSci`

Frac
f ☆ `\xintFracToSci` was initially at 1.4 a private macro which served as default customization of `\xintexprPrintOne` and, despite being documented in the user manual, was not supposed to be used at user level (not being *f-expandable* it could not be nested within macros of `xintfrac`, and besides accepted a limited range of inputs).

It has been upgraded at 1.41 to behave like all other `xintfrac` macros.

Here is what it does:

- it first parses the input like any other `xintfrac` macro and convert it into the ``raw' A/B[N]` format,
- it then produces this output: `A/B` if `N=0` (and `/B` is omitted if not `1`), and for `N` not zero, the output numerator will be `AeN` written in scientific notation exactly like it would be by `\xintPFloat` but without of course prior rounding to a given number of digits; the trailing zeros in the significand will be removed always (the `\xintPFloatMinTrimmed` configuration is ignored). Then this value in scientific notation (or in decimal fixed point notation if the scientific exponent is in the `\xintPFloatNoSciEmin` to `\xintPFloatNoSciEmax` range) will be attached to a trailing denominator `/B` (omitted if it is `/1`).

Please note:

- there is no reduction of the fraction `A/B` to lowest terms,
- trailing zeros in the integer denominator `B` are not moved and incorporated into the final scientific exponent,
- no attempt is made to check if `B` is a product of only `2`'s and `5`'s and thus could be integrated into some pure decimal notation for the numerator or at least its significand.

Changes of `\xintPFloat` at 1.4k have an impact here. In particular the zero value will give `0` whether the input was some `0`, `0e-5`, `0/3`, `0.00`, etc. . . , whereas at 1.4e it would have been `0.0` for cases triggering some `\xintPFloat` subroutine.

The general blueprint is still to be considered *unstable*.

The output routine of `\xintieval` is customizable via redefining `\xintexprPrintOne` whose current default is (equivalent to):

```
\def\xintexprPrintOne{\xintFracToSci}
```

12.21. `\xintFracToDecimal`

$\frac{\text{Frac}}{f}$ ★ `\xintFracToDecimal` is a variant of `\xintFracToSci` which differs from it in so far as it outputs a numerator using decimal notation, i.e. with as many zeros as are needed (and no more) and no scientific exponent. The denominator goes through ``as is'' except if it is 1, then it is omitted.

In other terms its behaviour is currently intermediate between `\xintDecToString` and `\xintDecToStringREZ`, as it does not remove trailing zeros of the denominator. Consider its behaviour as *unstable*.

It can be used to customize `\xintexprPrintOne`:

```
\def\xintexprPrintOne{\xintFracToDecimal}
```

It was initially at 1.4k a private macro which served as an alternative to `\xintFracToSci` default customization of `\xintexprPrintOne` and, despite being documented in the user manual, was not supposed to be used at user level (not being *f-expandable* it could not be nested within macros of `xintfrac`, and besides accepted a limited range of inputs).

It has been upgraded at 1.41 to behave like all other `xintfrac` macros.

12.22. `\xintTrunc`

$\frac{\text{num}}{x} \frac{\text{Frac}}{f}$ ★ `\xintTrunc{x}{f}` returns the start of the decimal expansion of the fraction `f`, truncated to:

- if `x>0`, `x` digits after the decimal mark,
- if `x=0`, an integer,
- if `x<0`, an integer multiple of 10^{-x} (in scientific notation).

The output is the sole digit token 0 if and only if the input was exactly zero; else it contains always either a decimal mark (even if `x=0`) or a scientific part and it conserves the sign of `f` (even if the truncated value represents the zero value).

Truncation is done towards zero.

```
\begin{multicols}{2}
  \noindent\xintFor* #1 in {\xintSeq[-1]{7}{-14}}:{\xintTrunc{#1}{-11e12/7}\newline}%
  \xintTrunc{10}{1e-11}\newline \xintTrunc{10}{1/65536}\par
\end{multicols}
```

<code>-1571428571428.5714285</code>	<code>-15714285e5</code>
<code>-1571428571428.571428</code>	<code>-1571428e6</code>
<code>-1571428571428.57142</code>	<code>-157142e7</code>
<code>-1571428571428.5714</code>	<code>-15714e8</code>
<code>-1571428571428.571</code>	<code>-1571e9</code>
<code>-1571428571428.57</code>	<code>-157e10</code>
<code>-1571428571428.5</code>	<code>-15e11</code>
<code>-1571428571428.</code>	<code>-1e12</code>
<code>-157142857142e1</code>	<code>-0e13</code>
<code>-15714285714e2</code>	<code>-0e14</code>
<code>-1571428571e3</code>	<code>0.0000000000</code>
<code>-157142857e4</code>	<code>0.0000152587</code>

Warning: *it is not yet decided is the current behaviour is definitive.*

Currently `xintfrac` has no notion of a positive zero or a negative zero. Hence transitivity of `\xintTrunc` is broken for the case where the first truncation gives on output `0.00...0` or `-0.00...0`: a second truncation to less digits will then output 0, whereas if it had been applied

directly to the initial input it would have produced $0.00\dots 0$ or respectively $-0.00\dots 0$ (with less zeros after decimal mark).

If `xintfrac` distinguished zero, positive zero, and negative zero then it would be possible to maintain transitivity.

The problem would also be fixed, even without distinguishing a negative zero on input, if `\xintTrunc` always produced $0.00\dots 0$ (with no sign) when the mathematical result is zero, discarding the information on original input being positive, zero, or negative.

I have multiple times hesitated about what to do and must postpone again final decision.

12.23. `\xintXTrunc`

$\frac{\text{num}}{\text{x}} \frac{\text{Frac}}{f}$ ☆

`\xintXTrunc{x}{f}` is similar to `\xintTrunc` with the following important differences:

- it is completely expandable but not *f-expandable*, as is indicated by the hollow star in the margin,
- hence it can not be used as argument to the other package macros, but as it *f-expands* its `{f}` argument, it accepts arguments expressed with other `xintfrac` macros,
- it requires $x > 0$,
- contrarily to `\xintTrunc` the number of digits on output is not limited to about 19950 and may go well beyond 100000 (this is mainly useful for outputting a decimal expansion to a file),
- when the mathematical result is zero, it always prints it as $0.00\dots 0$ or $-0.00\dots 0$ with x zeros after the decimal mark.

Warning: transitivity is broken too (see discussion of `\xintTrunc`), due to the sign in the last item. Hence *the definitive policy is yet to be fixed*.

Transitivity is here in the sense of using a first `\edef` and then a second one, because it is not possible to nest `\xintXTrunc` directly as argument to itself. Besides, although the number of digits on output isn't limited, nevertheless x should be less than about 19970 when the number of digits of the input (assuming it is expressed as a decimal number) is even bigger: `\xintXTrunc{30000}{Z}` after `\edef Z{\xintXTrunc{60000}{1/66049}}` raises an error in contrast with a direct `\xintXTrunc{30000}{1/66049}`. But `\xintXTrunc{30000}{123.456789}` works, because here the number of digits originally present is smaller than what is asked for, thus the routine only has to add trailing zeros, and this has no limitation (apart from TeX main memory).

`\xintXTrunc` will expand fully in an `\edef` or a `\write (\message, \wlog, ...)` or in an `\xint-expression`, or as list argument to `\xintFor*`.

Here is an example session where the user checks that the decimal expansion of $1/66049 = 1/257^2$ has the maximal period length $257 * 256 = 65792$ (this period length must be a divisor of $\phi(66049)$ and to check it is the maximal one it is enough to show that neither 32896 nor 256 are periods.)

```
$ rlwrap etex -jobname worksheet-66049
This is pdfTeX, Version 3.14159265-2.6-1.40.17 (TeX Live 2016) (preloaded format=etex)
restricted \writel8 enabled.
**xintfrac.sty
entering extended mode
(/usr/local/texlive/2016/texmf-dist/tex/generic/xint/xintfrac.sty
(/usr/local/texlive/2016/texmf-dist/tex/generic/xint/xint.sty
(/usr/local/texlive/2016/texmf-dist/tex/generic/xint/xintcore.sty
(/usr/local/texlive/2016/texmf-dist/tex/generic/xint/xintkernel.sty)))
% we load xinttools for \xintKeep, etc... \xintXTrunc itself has no more

% any dependency on xinttools.sty since 1.2i

*\input xinttools.sty
(/usr/local/texlive/2016/texmf-dist/tex/generic/xint/xinttools.sty)
*\def\m#1;{\message{#1}}
```

```

*\m \the\numexpr 257*257\relax;
66049
*\m \the\numexpr 257*256\relax;
65792
% Thus 1/66049 will have a period length dividing 65792.

% Let us first check it is indeed periodical.

*\edef\Z{\xintXTrunc{66000}{1/66049}}

% Let's display the first decimal digits.

*\m \xintXTrunc{208}{\Z};

0.00001514027464458205271843631243470756559523989765174340262532362337052794137
6856576178291874214598252812306015231116292449545034746930309315810988811337037
6538630410755651107511090251177156353616254598858423
% let's now fetch the trailing digits

*\m \xintKeep{65792-66000}{\Z};% 208 trailing digits

0000151402746445820527184363124347075655952398976517434026253236233705279413768
5657617829187421459825281230601523111629244954503474693030931581098881133703765
38630410755651107511090251177156353616254598858423
% yes they match! we now check that 65792/2 and 65792/257=256 aren't periods.

*\m \xintXTrunc{256}{\Z};

0.00001514027464458205271843631243470756559523989765174340262532362337052794137
6856576178291874214598252812306015231116292449545034746930309315810988811337037
6538630410755651107511090251177156353616254598858423291798513225029902042423049
554118911717058547442
*\m \xintXTrunc{256+256}{\Z};

0.00001514027464458205271843631243470756559523989765174340262532362337052794137
6856576178291874214598252812306015231116292449545034746930309315810988811337037
6538630410755651107511090251177156353616254598858423291798513225029902042423049
5541189117170585474420505987978621932201850141561567926842192917379521264515738
3154930430438008145467758785144362518736089872670290239064936637950612424109373
3440324607488379839210283274538600130206361943405653378552286938485064119063119
8049932625777831609865402958409665551333
% now with 65792/2=32896. Problem: we can't do \xintXTrunc{32896+100}{\Z}

% but only direct \xintXTrunc{32896+100}{1/66049}. Anyway we want to nest it

% hence let's do it all with (slower) \xintKeep, \xintKeepUnbraced.

*\m \xintKeep {-100}{\xintKeepUnbraced{2+65792/2+100}{\Z}};

9999848597253554179472815636875652924344047601023482565973746763766294720586231
434238217081257854017
% This confirms 32896 isn't a period length.

% To conclude let's write the 66000 digits to the log.

*\wlog{\Z}

```

```

% We want always more digits:

\wlog{\xintXTrunc{150000}{1/66049}}

\bye

```

The acute observer will have noticed that there is something funny when one compares the first digits with those after the middle-period:

```

0000151402746445820527184363124347075655952398976517434026253236233705279413768...
9999848597253554179472815636875652924344047601023482565973746763766294720586231...

```

Mathematical exercise: can you explain why the two indeed add to 9999...9999?

You can try your hands at this simpler one:

```

1/49=\xintTrunc{42+5}{1/49}...\newline
\xintTrim{2}{\xintTrunc{21}{1/49}}\newline
\xintKeep{-21}{\xintTrunc{42}{1/49}}
1/49=0.02040816326530612244897959183673469387755102040...
020408163265306122448
979591836734693877551

```

This was again an example of the type $1/N$ with N the square of a prime. One can also find counter-examples within this class: $1/31^2$ and $1/37^2$ have an odd period length (465 and respectively 111) hence they can not exhibit the symmetry.

Mathematical challenge: prove generally that if the period length of the decimal expansion of $1/p^r$ (with p a prime distinct from 2 and 5 and r a positive exponent) is even, then the previously observed symmetry about the two halves of the period adding to a string of nine's applies.

12.24. `\xintTFrac`

`\xintTFrac{f}` returns the fractional part, $f = \text{trunc}(f) + \text{frac}(f)$. Thus if $f < 0$, then $-1 < \text{frac}(f) \leq 0$ and if $f > 0$ one has $0 \leq \text{frac}(f) < 1$. The T stands for 'Trunc', and there should exist also similar macros associated respectively with 'Round', 'Floor', and 'Ceil', each type of rounding to an integer deserving arguably to be associated with a fractional 'modulo'. By sheer laziness, the package currently implements only the 'modulo' associated with 'Truncation'. Other types of modulo may be obtained more clumsily via a combination of the rounding with a subsequent subtraction from f .

Notice that the result is filtered through `\xintREZ`, and will thus be of the form $A/B[N]$, where neither A nor B has trailing zeros. But the output fraction is not reduced to smallest terms.

The function call in expressions (`\xintexpr`, `\xintfloatexpr`) is `frac`. Inside `\xintexpr..relax`, the function `frac` is mapped to `\xintTFrac`. Inside `\xintfloatexpr..relax`, `frac` first applies `\xintTFrac` to its argument (which may be an exact fraction with more digits than the floating point precision) and only in a second stage makes the conversion to a floating point number with the precision as set by `\xintDigits` (default is 16).

```

\xintTFrac {1235/97}, \xintTFrac {-1235/97}\newline
\xintTFrac {1235.973}, \xintTFrac {-1235.973}\newline
\xintTFrac {1.122435727e5}\par
71/97[0], -71/97[0]
973/1[-3], -973/1[-3]
5727/1[-4]

```

12.25. `\xintRound`

$\frac{\text{num}}{x} \frac{\text{Frac}}{f}$ ★ `\xintRound{x}{f}` returns the start of the decimal expansion of the fraction `f`, rounded to:

- if `x>0`, `x` digits after the decimal mark,
- if `x=0`, an integer,
- if `x<0`, an integer multiple of `10^{-x}` (in scientific notation).

The output is the sole digit token `0` if and only if the input was exactly zero; else it contains always either a decimal mark (even if `x=0`) or a scientific part and it conserves the sign of `f` (even if the rounded value represents the zero value).

```
\begin{multicols}{2}
  \noindent\xintFor* #1 in {\xintSeq[-1]{7}{-14}}:{\xintRound{#1}{-11e12/7}\newline}%
  \xintRound{10}{1e-11}\newline \xintRound{10}{1/65536}\newline
\end{multicols}
```

-1571428571428.5714286	-1571429e6
-1571428571428.571429	-157143e7
-1571428571428.57143	-15714e8
-1571428571428.5714	-1571e9
-1571428571428.571	-157e10
-1571428571428.57	-16e11
-1571428571428.6	-2e12
-1571428571429.	-0e13
-157142857143e1	-0e14
-15714285714e2	0.0000000000
-1571428571e3	0.0000152588
-157142857e4	
-15714286e5	

Rounding is done with half-way numbers going towards infinity of the same sign.

12.26. `\xintFloor`

$\frac{\text{Frac}}{f}$ ★ `\xintFloor {f}` returns the largest relative integer `N` with `N ≤ f`.

```
\xintFloor {-2.13}, \xintFloor {-2}, \xintFloor {2.13}
-3/1[0], -2/1[0], 2/1[0] Note the trailing [0], see \xintiFloor if it is not desired.
```

12.27. `\xintCeil`

$\frac{\text{Frac}}{f}$ ★ `\xintCeil {f}` returns the smallest relative integer `N` with `N > f`.

```
\xintCeil {-2.13}, \xintCeil {-2}, \xintCeil {2.13}
-2/1[0], -2/1[0], 3/1[0]
```

12.28. `\xintiTrunc`

$\frac{\text{num}}{x} \frac{\text{Frac}}{f}$ ★ `\xintiTrunc{x}{f}` returns the integer equal to `10^x` times what `\xintTrunc{x}{f}` would produce. Attention that leading zeros are automatically removed: the output is in strict integer format.

```
\begin{multicols}{2}
  \noindent\xintFor* #1 in {\xintSeq[-1]{7}{-14}}:{\xintiTrunc{#1}{-11e12/7}\newline}%
  \xintiTrunc{10}{1e-11}\newline \xintiTrunc{10}{1/65536}\par
\end{multicols}
```

-15714285714285714285	-15714285
-1571428571428571428	-1571428
-157142857142857142	-157142
-15714285714285714	-15714
-1571428571428571	-1571
-157142857142857	-157
-15714285714285	-15
-1571428571428	-1
-157142857142	0
-15714285714	0
-1571428571	0
-157142857	152587

12.29. `\xintTrunc`

$\frac{\text{Frac}}{f}$ ★ `\xintTrunc{f}` truncates to an integer (truncation towards zero). This is the same as `\xintiTrunc{f}` and also the same as `\xintNum`.

12.30. `\xintiRound`

$\frac{\text{num}}{x} \frac{\text{Frac}}{f}$ ★ `\xintiRound{x}{f}` returns the integer equal to 10^x times what `\xintRound{x}{f}` would return. The output has no leading zeroes, it is always in strict integer format.

```
\begin{multicols}{2}
  \noindent\xintFor* #1 in {\xintSeq[-1]{7}{-14}}:{\xintiRound{#1}{-11e12/7}\newline}%
  \xintiRound{10}{1e-11}\newline \xintiRound{10}{1/65536}\par
\end{multicols}
```

-15714285714285714286	-15714286
-1571428571428571429	-1571429
-157142857142857143	-157143
-15714285714285714	-15714
-1571428571428571	-1571
-157142857142857	-157
-15714285714286	-16
-1571428571429	-2
-157142857143	0
-15714285714	0
-1571428571	0
-157142857	152588

12.31. `\xintiFloor`

$\frac{\text{Frac}}{f}$ ★ `\xintiFloor {f}` does the same as `\xintFloor` but without the trailing `/1[0]`.
`\xintiFloor {-2.13}, \xintiFloor {-2}, \xintiFloor {2.13}`
 -3, -2, 2

12.32. `\xintiCeil`

$\frac{\text{Frac}}{f}$ ★ `\xintiCeil {f}` does the same as `\xintCeil` but its output is without the `/1[0]`.
`\xintiCeil {-2.13}, \xintiCeil {-2}, \xintiCeil {2.13}`
 -2, -2, 3

12.33. `\xintE`

$\frac{\text{Frac num}}{f \ x}$ ★ `\xintE {f}{x}` multiplies the fraction `f` by 10^x . The second argument `x` must obey the TeX bounds.
 Example:
`\count 255 123456789 \xintE {10}{\count 255}`
`10/1[123456789]` Don't feed this example to `\xintNum`!

12.34. `\xintCmp`

$\frac{\text{Frac Frac}}{f \ f}$ ★ This compares two fractions `F` and `G` and produces `-1`, `0`, or `1` according to `F<G`, `F=G`, `F>G`.
 For choosing branches according to the result of comparing `f` and `g`, see `\xintifCmp`.

12.35. `\xintEq`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintEq{f}{g}` returns `1` if `f=g`, `0` otherwise.

12.36. `\xintNotEq`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintNotEq{f}{g}` returns `0` if `f=g`, `1` otherwise.

12.37. `\xintGeq`

$\frac{\text{Frac Frac}}{f \ f}$ ★ This compares the absolute values of two fractions. `\xintGeq{f}{g}` outputs `1` if $|f| \geq |g|$ and `0` if not.
 Important: the macro compares absolute values.

12.38. `\xintGt`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintGt{f}{g}` returns `1` if `f>g`, `0` otherwise.

12.39. `\xintLt`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintLt{f}{g}` returns `1` if `f<g`, `0` otherwise.

12.40. `\xintGtorEq`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintGtorEq{f}{g}` returns `1` if `f≥g`, `0` otherwise. Extended by `xintfrac` to fractions.

12.41. `\xintLtorEq`

$\frac{\text{Frac Frac}}{f \ f}$ ★ `\xintLtorEq{f}{g}` returns `1` if `f≤g`, `0` otherwise.

12.42. `\xintIsZero`

`f` ★ `\xintIsZero{f}` returns `1` if `f=0`, `0` otherwise.

12.43. `\xintIsNotZero`

`f` ★ `\xintIsNotZero{f}` returns `1` if `f!=0`, `0` otherwise.

12.44. `\xintIsOne`

`f` ★ `\xintIsOne{f}` returns `1` if `f=1`, `0` otherwise.

12.45. `\xintOdd`

`f` ★ `\xintOdd{f}` returns `1` if the integer obtained by truncation is odd, and `0` otherwise.

12.46. `\xintEven`

f ★ `\xintEven{f}` returns 1 if the integer obtained by truncation is even, and 0 otherwise.

12.47. `\xintifSgn`

$\frac{Frac}{f} nnn$ ★ `\xintifSgn{f}{A}{B}{C}` executes either the $\langle A \rangle$, $\langle B \rangle$ or $\langle C \rangle$ code, depending on its first argument being respectively negative, zero, or positive.

12.48. `\xintifZero`

$\frac{Frac}{f} nn$ ★ `\xintifZero{f}{IsZero}{IsNotZero}` expandably checks if the first mandatory argument N (a number, possibly a fraction if `xintfrac` is loaded, or a macro expanding to one such) is zero or not. It then either executes the first or the second branch.

Beware that both branches must be present.

12.49. `\xintifNotZero`

$\frac{Frac}{f} nn$ ★ `\xintifNotZero{N}{IsNotZero}{IsZero}` expandably checks if the first mandatory argument f is not zero or is zero. It then either executes the first or the second branch.

Beware that both branches must be present.

12.50. `\xintifOne`

$\frac{Frac}{f} nn$ ★ `\xintifOne{N}{IsOne}{IsNotOne}` expandably checks if the first mandatory argument f is one or not one. It then either executes the first or the second branch. Beware that both branches must be present.

12.51. `\xintifOdd`

$\frac{Frac}{f} nn$ ★ `\xintifOdd{N}{odd}{not odd}` expandably checks if the first mandatory argument f , after truncation to an integer, is odd or even. It then executes accordingly the first or the second branch. Beware that both branches must be present.

12.52. `\xintifCmp`

$\frac{Frac}{f} \frac{Frac}{f} nnn$ ★ `\xintifCmp{f}{g}{if f<g}{if f=g}{if f>g}` compares its first two arguments and chooses accordingly the correct branch.

12.53. `\xintifEq`

$\frac{Frac}{f} \frac{Frac}{f} nn$ ★ `\xintifEq{f}{g}{YES}{NO}` checks equality of its two first arguments and executes accordingly the YES or the NO branch.

12.54. `\xintifGt`

$\frac{Frac}{f} \frac{Frac}{f} nn$ ★ `\xintifGt{f}{g}{YES}{NO}` checks if $f > g$ and in that case executes the YES branch.

12.55. `\xintifLt`

$\frac{Frac}{f} \frac{Frac}{f} nn$ ★ `\xintifLt{f}{g}{YES}{NO}` checks if $f < g$ and in that case executes the YES branch.

12.56. `\xintifInt`

$\frac{\text{Frac}}{f} nn$ ★ `\xintifInt{f}{YES branch}{NO branch}` expandably chooses the YES branch if `f` reveals itself after expansion and simplification to be an integer.

12.57. `\xintSgn`

$\frac{\text{Frac}}{f}$ ★ The sign of a fraction.

12.58. `\xintSignBit`

$\frac{\text{Frac}}{f}$ ★ Expands to 1 for negative input, to 0 else.
Added at 1.41.

12.59. `\xintOpp`

$\frac{\text{Frac}}{f}$ ★ The opposite of a fraction. Note that `\xintOpp {3}` produces `-3/1[0]` whereas `\xintiiOpp {3}` produces `-3`.

12.60. `\xintAbs`

$\frac{\text{Frac}}{f}$ ★ The absolute value. Note that `\xintAbs {-2}=2/1[0]` where `\xintiiAbs {-2}` outputs `=2`.

12.61. `\xintAdd`

$\frac{\text{Frac}}{f} \frac{\text{Frac}}{f}$ ★ Computes the addition of two fractions.
Since 1.3 always uses the least common multiple of the denominators.

12.62. `\xintSub`

$\frac{\text{Frac}}{f} \frac{\text{Frac}}{f}$ ★ Computes the difference of two fractions (`\xintSub{F}{G}` computes `F-G`).
Since 1.3 always uses the least common multiple of the denominators.

12.63. `\xintMul`

$\frac{\text{Frac}}{f} \frac{\text{Frac}}{f}$ ★ Computes the product of two fractions.
Output is not reduced to smallest terms.

12.64. `\xintDiv`

$\frac{\text{Frac}}{f} \frac{\text{Frac}}{f}$ ★ Computes the quotient of two fractions. (`\xintDiv{F}{G}` computes `F/G`).
Output is not reduced to smallest terms.

12.65. `\xintDivFloor`

$\frac{\text{Frac}}{f} \frac{\text{Frac}}{f}$ ★ Computes the quotient of two arguments then apply floor function to get an integer (in strict format). This macro was defined at 1.1 (but was left not documented until 1.3a...) and changed at 1.2p, formerly it appended `/1[0]` to output.

`\xintDivFloor{-170/3}{23/2}`

-5

12.66. `\xintMod`

$\frac{f}{f}$ \star Computes the remainder associated to the floored division `\xintDivFloor`. Prior to 1.2p the meaning was the one of `\xintModTrunc`. Was left undocumented until 1.3a.

`\xintMod{-170/3}{23/2}`
 $5/6[0]$
 Modified at 1.3 to use a l.c.m. for the denominator of the result.

12.67. `\xintDivMod`

$\frac{f}{f}$ \star Computes both the floored division and the remainder `\xintDivFloor`. New at 1.2p and documented at 1.3a.

`\oodef\foo{\xintDivMod{-170/3}{23/2}}\meaning\foo`
`macro:->{-5}{5/6[0]}`

12.68. `\xintDivTrunc`

$\frac{f}{f}$ \star Computes the quotient of two arguments then truncates to an integer (in strict format).

`\xintDivTrunc{-170/3}{23/2}`
 -4

12.69. `\xintModTrunc`

$\frac{f}{f}$ \star Computes the remainder associated with the truncated division of two arguments. Prior to 1.2p it was named `\xintMod`, but the latter then got associated with floored division.

`\xintModTrunc{-170/3}{23/2}`
 $-64/6[0]$
 Modified at 1.3 to use a l.c.m. for the denominator of the result.

12.70. `\xintDivRound`

$\frac{f}{f}$ \star Computes the quotient of the two arguments then rounds to an integer (in strict format).

`\xintDivRound{-170/3}{23/2}`
 -5

12.71. `\xintSqr`

$\frac{f}{f}$ \star Computes the square of one fraction.

12.72. `\xintPow`

$\frac{f}{f}$ \star `\xintPow{f}{x}`: computes f^x with f a fraction and the exponent x possibly also, but if only `xintfrac` is loaded it will be truncated to an integer.

At 1.4e the behaviour of the macro is enhanced if `xintexpr` is loaded, at it then becomes the support macro for powers a^b , $a**b$ (and the `pow()` function) in `\xinteval`: it now handles also non-integer exponents. Also, if the exponent is an integer, it checks a priori if an exact evaluation would produce more than about 10000 digits and then does in its place a floating point evaluation.

The check whether the exponent is integer is not on the mathematical value but on the format (for reasons of efficiency). So $4/2$ will not be recognized as integer and it will thus trigger usage of the floating point evaluations; however 2.0 will be recognized as an integer, as of course 2 .

If the exponent is considered an integer it is then checked if it is less than 10000 (in absolute value) and if the output would contain less than 10000 digits (separately for numerator and denominator) and only then is the power computed exactly. Else it is computed as by `\xintFloatPower`

(but the output uses raw `A[N]` format not scientific notation). Use `\xintiiPow` (on integers only, not fractions) for exact powers with larger exponents.

Also, a check is done whether the exponent is half-integer. Again this check is not on the value but on the format, so `2.5` is an half integer, as is `25e-1`, or `2.50` but `5/2` is not considered an half-integer (for reasons of internal efficiency). If the exponent is half-integer the power is computed by combining suitably `\xintFloatPower` with `\xintFloatSqrt` (but the output uses raw `A[N]` format not scientific notation).

If the exponent is neither an integer nor an half-integer, the power is computed using logarithm and exponential based approach (and uses raw `A[N]` output format). If `Digits` is at most 8 (which triggers `poormanlog` usage, for very fast logarithms but only with about 8 or 9 accurate fractional digits) this will start being inaccurate in the last digit already with fractional exponents `x > 10`. It is recommended to split then the exponent into an integer or half-integer part and a fractional part. Powers with integer or half-integer exponents, even very big, are always computed accurately, for any value of `Digits`.

Within an `\xintiexpr` `\relax` the infix operators `^` and `**` are mapped to `\xintiiPow` and powers are always computed exactly even if they would produce more than 10000 digits and melt your CPU; within an `\xintexpr` `^` and `**` are mapped to `\xintPow` as described here.

12.73. `\xintFac`

$\overset{\text{Num}}{f}$ \star This is a convenience variant of `\xintiiFac` which applies `\xintNum` to its argument. Notice however that the output will have a trailing `[0]` according to the `xintfrac` format for integers.

12.74. `\xintBinomial`

$\overset{\text{Num Num}}{f f}$ \star This is a convenience variant of `\xintiiBinomial` which applies `\xintNum` to its arguments. Notice however that the output will have a trailing `[0]` according to the `xintfrac` format for integers.

12.75. `\xintPFactorial`

$\overset{\text{Num Num}}{f f}$ \star This is a convenience variant of `\xintiiPFactorial` which applies `\xintNum` to its arguments. Notice however that the output will have a trailing `[0]` according to the `xintfrac` format for integers.

12.76. `\xintMax`

$\overset{\text{Frac Frac}}{f f}$ \star The maximum of two fractions. Beware that `\xintMax {2}{3}` produces `3/1[0]`. The original, for use with integers only with no need of normalization, is available as `\xintiiMax: \xintiiMax {2}{3}`
 ff \star =3.
`\xintMax {2.5}{7.2}`
`72/1[-1]`

12.77. `\xintMin`

$\overset{\text{Frac Frac}}{f f}$ \star The minimum of two fractions. Beware that `\xintMin {2}{3}` produces `2/1[0]`. The original, for use with integers only with no need of normalization, is available as `\xintiiMin: \xintiiMin {2}{3}`
 ff \star =2.
`\xintMin {2.5}{7.2}`
`25/1[-1]`

12.78. `\xintMaxof`

$f \rightarrow * \overset{\text{Frac}}{f}$ \star The maximum of any number of fractions, each within braces, and the whole thing within braces.
`\xintMaxof {{1.23}{1.2299}{1.2301}}` and `\xintMaxof {{-1.23}{-1.2299}{-1.2301}}`
`12301/1[-4]` and `-12299/1[-4]`

12.79. `\xintMinof`

$f \rightarrow * \overset{\text{Frac}}{f} \star$ The minimum of any number of fractions, each within braces, and the whole thing within braces.
`\xintMinof {{1.23}}{1.2299}{1.2301}}` and `\xintMinof {{-1.23}}{-1.2299}{-1.2301}}`
 12299/1[-4] and -12301/1[-4]

12.80. `\xintSum`

$f \rightarrow * \overset{\text{Frac}}{f} \star$ This computes the sum of fractions. The output will now always be in the form `A/B[n]`. The original, for big integers only (in strict format), is available as `\xintiiSum`.
`\xintSum {{1282/2196921}}{-281710/291927}{4028/28612}}`
 -5037928302100692/6116678670072468[0]
 No simplification attempted.

12.81. `\xintPrd`

$f \rightarrow * \overset{\text{Frac}}{f} \star$ This computes the product of fractions. The output will now always be in the form `A/B[n]`. The original, for big integers only (in strict format), is available as `\xintiiPrd`.
`\xintPrd {{1282/2196921}}{-281710/291927}{4028/28612}}`
 -1454721142160/18350036010217404[0]
 No simplification attempted.
`\xintIsOne {21921379213/21921379213}\neq`
`\xintIsOne {1.00000000000000000000000000000001}$`
`1 ≠ 0`

12.82. `\xintGCD`

$\overset{\text{Frac}}{f} \overset{\text{Frac}}{f} \star$ The greatest common divisor of its two arguments, which are possibly *fractions*.
 Prior to 1.4 a macro of the same name existed in `xintgcd`. But it truncated its two arguments to integers via `\xintNum`.
 See `\xintiiGCD` for the integer only variant.

12.83. `\xintLCM`

$\overset{\text{Frac}}{f} \overset{\text{Frac}}{f} \star$ The least common multiple of its two arguments, which are possibly *fractions*.
 Prior to 1.4 a macro of the same name existed in `xintgcd`. But it truncated its two arguments to integers via `\xintNum`.
 See `\xintiiLCM` for the integer only variant.

12.84. `\xintGCDoF`

$f \rightarrow * \overset{\text{Frac}}{f} \star$ `\xintGCDoF{{a}}{b}{c}...` computes the greatest common divisor of `a`, `b`, `c`, `...`. The arguments are allowed to be *fractions*: the macro produces the non-negative generator of the fractional ideal they generate. The list argument may be a macro as it is *f-expanded* first. If all arguments vanish, then also the output.
 Prior to 1.4 a macro of the same name existed in `xintgcd`. But it truncated all its arguments to integers via `\xintNum` and then proceeded with integer only computations.
 See `\xintiiGCDoF` for the integer only variant.

12.85. `\xintLCMoF`

$f \rightarrow * \overset{\text{Frac}}{f} \star$ `\xintLCMoF{{a}}{b}{c}...` computes the least common multiple of `a`, `b`, `c`, `...`. The arguments are allowed to be *fractions*: the macro produces the non-negative generator of the intersection of the

corresponding fractional ideals. The list argument may be a macro, it is *f-expanded* first. If one of the item vanishes, then also the output.

Prior to 1.4 a macro of the same name existed in *xintgcd*. But it truncated all its arguments to integers via `\xintNum`.

See `\xintiLCMof` for the integer only variant.

12.86. `\xintDigits`, `\xinttheDigits`

The syntax `\xintDigits := N`; or (recommended) `\xintDigits := N\relax` assigns the value of *N* to the number of digits to be used by floating point operations (this uses internally a `\mathchardef` *f* assignment, and *N* stands for (or expands to) a legal \TeX number). The default is 16. The maximal value is 32767.

Accepted syntax includes also `\xintDigits = N`; or `\xintDigits = N\relax`, i. e. the colon before the equality sign is optional.



`xintexpr` adds the variant `\xintDigits*` which executes `\xintreloadxinttrig` and `\xint-reloadxintlog`.

A priori, you want `\xintDigits*:=N\relax`. Use `\xintDigits:=N\relax` only if not needing trigonometric or logarithm/exponential functions and wanting to avoid the overhead of reloading their libraries. Perhaps for a local temporary configuration.

Spaces do not matter as long as they do not occur in-between digits:

```
\xintDigits := 24\relax\xinttheDigits, %
\xintDigits:=36 \relax\xinttheDigits, %
\xintDigits:= 16 \relax and \xinttheDigits.
```

★ 24, 36, and 16. As shown above `\xinttheDigits` expands to the stored value.

An ending active semi-colon ; is *not* compatible: it can and will cause low-level \TeX errors. This is why the alternative syntax

```
\xintDigits:= N\relax
```

is recommended (with or without the semi-colon). This is hopefully the syntax now in use in most examples from the documentation.

Actually, any non-expanding token can be used in place of the `\relax`. This non-expanding ending token (for example a full stop) will get removed from the token stream.

```
\xintDigits = 24\def\xinttheDigits, % only for showing it works! don't do that!
\xintDigits := 36.\xinttheDigits, % one can use a dot in place of semi-colon
\xintDigits = 16\relax and \xinttheDigits.\par % with \relax, even better
```

24, 36, and 16.

12.87. `\xintSetDigits`

^{num}
x

To be used as `\xintSetDigits{<expression>}` where the expression will be fed to `\numexpr`. It is a shortcut for doing `\xintDigits := \numexpr<expression>\relax \relax`.

```
\xintSetDigits{1+2+3+4+5}The value is now \xinttheDigits.
\xintSetDigits{2*8}The value is now \xinttheDigits.\par
```

The value is now 15. The value is now 16.

The `xintexpr`-added variant `\xintSetDigits*` is the preferred usage as it does the extra work to update the math functions from `xinttrig` and `xintlog`.


```
\do{#1: \xintFloat[#1]{\x}\newline}}%
\xintFloat{5/999999999999999}\newline
\xintFloat[32]{5/999999999999999}\newline
\xintFloat[48]{5/999999999999999}\par
```

```
4: 9.927e6
5: 9.9275e6
6: 9.92746e6
7: 9.927460e6
8: 9.9274600e6
9: 9.92745997e6
10: 9.927459975e6
11: 9.9274599746e6
12: 9.92745997457e6
13: 9.927459974572e6
14: 9.9274599745717e6
15: 9.92745997457166e6
16: 9.927459974571665e6
17: 9.9274599745716647e6
18: 9.92745997457166465e6
19: 9.927459974571664655e6
20: 9.9274599745716646545e6
5.0000000000000001e-16
5.000000000000000500000000000001e-16
5.000000000000000500000000000005000000000001e-16
```

12.89. `\xintFloatBraced`

[_{num}/_x] Frac f ★

The experimental macro `\xintFloatBraced[P]{f}` does like `\xintFloat` but its output consists of three \TeX -braced groups

$\langle\textit{sign bit}\rangle\langle\textit{scientific exponent}\rangle\langle\textit{full width mantissa with decimal point}\rangle$

It is provided for users knowing how to pick one or the other of these constituents from usage of auxiliary macros. Or one can use `\xintAssign`:

```
\begingroup
\xintAssign\xintFloatBraced[7]{-3.1234e-14}\to\A\B\C
\string\A\ has meaning \meaning\A\newline
\string\B\ has meaning \meaning\B\newline
\string\C\ has meaning \meaning\C\par
\endgroup
```

```
\A has meaning macro:->1
\B has meaning macro:->-14
\C has meaning macro:->3.123400
```

Some aspects are undecided:

- should the first item be rather `-1`, `0`, or `1`? or `-`, nothing, nothing?
- in case of zero value the output ignores `\xintFloatZero`, it uses a zero exponent and full width fractional mantissa `0.000000000000000` (here no `[P]` and `\xinttheDigits` has value `16`), should it be otherwise?
- should the mantissa be without the decimal separator? should it incorporate the sign?
- in case the mantissa is without separator, should the exponent be biased to match it?

12.90. `\xintFloatToDecimal`

[_{num}/_x] Frac f ★

`\xintFloatToDecimal [P]{f}` does float rounding on input like `\xintFloat` then outputs the number

using decimal notation, i.e. with as many zeros as are needed (and no more) and no scientific exponent.

In other terms it behaves (and is essentially defined) as:

```
\xintDecToStringREZ{\xintFloat[optional P]{<input>}}
```

Examples:

```
\xintFloatToDecimal{6.02e23}\newline
\xintFloatToDecimal{6.020000000000000e23}\newline
\xintFloatToDecimal[20]{1/7e10}\newline
\xintFloatToDecimal[30]{1/7e10}
```

```
60200000000000000000000000000000000000000000
60200000000000000000000000000000000000000000
0.000000000014285714285714285714285714
0.0000000000142857142857142857142857142857
```

See `\xintDecToString`.

12.91. `\xintPFloat`

`\xintPFloat [P]{f}` is like `\xintFloat` but "pretty-prints" the output.

This macro was initially added at 1.1 as a (very primitive) "prettifying printer" for floating point number, and was then somewhat influenced by Maple, for example the zero value was printed as "0.". Then at 1.4e there was breaking change and the rules became somewhat similar to observed Python behaviour: mantissas trimmed of trailing zeros (whether or not scientific notation was used in the output) and integers printed with a trailing ".0", in particular the zero value was printed as "0.0".

1.4k brought some breaking changes, which are reversible via customizing macros:

- Integers (when scientific notation is dropped according to criteria mentioned next) without a ".0" suffix.
- Same for the zero value, now "0".
- Significant digits are trimmed of trailing zeros only if that removes at least 4 zeros. The rationale is that automatic removal of trailing zeros (which was influenced at 1.4e from practice with Python in interactive mode) proves annoying visually with aligned values in tables, as this creates voids, so we want to do this only when really the presence of trailing zeros is not some kind of numerical fluke.

These changes impact the `\xintfloateval` output as `\xintfloatexprPrintOne` defaults to using `\xintPFloat`.

In this documentation "trailing zeros" refers not to how the input looked like, but to the corresponding mantissa of width `P` or `\xinttheDigits`.

The default rules are thus now:

1. The input is float-rounded to either `Digits` or the optional argument.
2. zero is printed as 0.
3. `x.yz...eN` is printed in decimal fixed point if $-4 \leq N \leq +5$ else it is printed in scientific notation.
4. Trailing zeros of the mantissa are trimmed if, and only if there are at least 4 of them.
5. In case of fixed point output format, and the value is an integer, the integer is printed with no decimal mark.
6. In case of scientific notation output format, and the mantissa has only one digit, no decimal mark is used.

<p><code>\xintDigits at 16</code></p> <ul style="list-style-type: none"> • <code>0</code> → <code>0</code> • <code>1.2340000e-7</code> → <code>1.234e-7</code> • <code>1.2340000e-6</code> → <code>1.234e-6</code> • <code>1.2340000e-5</code> → <code>1.234e-5</code> • <code>1.2340000e-4</code> → <code>0.0001234</code> • <code>1.2340000e-3</code> → <code>0.001234</code> • <code>1.2340000e-2</code> → <code>0.01234</code> • <code>1.2340000e-1</code> → <code>0.1234</code> • <code>1.2340000e0</code> → <code>1.234</code> • <code>1.2340000e1</code> → <code>12.34</code> • <code>1.2340000e2</code> → <code>123.4</code> • <code>1.2340000e3</code> → <code>1234</code> • <code>1.2340000e4</code> → <code>12340</code> • <code>1.2340000e5</code> → <code>123400</code> • <code>1.2340000e6</code> → <code>1.234e6</code> 	<ul style="list-style-type: none"> • <code>1.2340000e7</code> → <code>1.234e7</code> • <code>1e-7/7</code> → <code>1.428571428571429e-8</code> • <code>1e-6/7</code> → <code>1.428571428571429e-7</code> • <code>1e-5/7</code> → <code>1.428571428571429e-6</code> • <code>1e-4/7</code> → <code>1.428571428571429e-5</code> • <code>1e-3/7</code> → <code>0.0001428571428571429</code> • <code>1e-2/7</code> → <code>0.001428571428571429</code> • <code>1e-1/7</code> → <code>0.01428571428571429</code> • <code>1e0/7</code> → <code>0.1428571428571429</code> • <code>1e1/7</code> → <code>1.428571428571429</code> • <code>1e2/7</code> → <code>14.28571428571429</code> • <code>1e3/7</code> → <code>142.8571428571429</code> • <code>1e4/7</code> → <code>1428.571428571429</code> • <code>1e5/7</code> → <code>14285.71428571429</code> • <code>1e6/7</code> → <code>142857.1428571429</code> • <code>1e7/7</code> → <code>1.428571428571429e6</code>
---	---

12.91.1. Customizing macros of `\xintPFloat`

A number of macros allow to customize the behaviour of `\xintPFloat`:

- `\xintPFloatE` allows to modify the separator of the scientific notation. Here is its default:
⁷²

```
\def\xintPFloatE{e}
```

- `\xintPFloatZero` says how to print the zero value. The default:

```
\def\xintPFloatZero{0}
```

- `\xintPFloatIntSuffix` is postfixed to integer values (when scientific notation is not used). Its default at 1.4k is to add nothing. It replaces the formerly hard-coded ".0" from 1.4e (prior to that trailing zeros from the full significand of P or `\xinttheDigits` digits were not trimmed).

```
\def\xintPFloatIntSuffix{}
```

- `\xintPFloatLengthOneSuffix` is postfixed to trimmed mantissas having only one digit, when scientific notation is used. Its default at 1.4k is to add nothing. It replaces formerly hard-coded ".0".

```
\def\xintPFloatLengthOneSuffix{}
```

- `\xintPFloatNoSciEmax` is the maximal scientific exponent which will trigger use of decimal fixed point notation and `\xintPFloatNoSciEmin` is the minimal one. Their defaults at 1.4k are the same as the formerly hard-coded behaviour from 1.4e:

```
\def\xintPFloatNoSciEmax{5}  
\def\xintPFloatNoSciEmin{-4}
```

For example (with the package default width of 16 digits for mantissas of floating point numbers):

```
\beginngroup
```

⁷² For \TeX pers: it is allowed to define `\xintPFloatE` as a macro which grabs the exponent as an argument delimited by a dot, and produces *f-expand*ably an output also delimited by a dot (it will be removed via further internal processing).

```
\def\xintPFloatNoSciEmin{-20}
\xintPFloat{1e-19/7}\newline
\xintPFloat{1e-20/7}\par
\def\xintPFloatNoSciEmax{19}
\xintPFloat{1e20/7}\newline
\xintPFloat{1e21/7}\par
\endgroup
```

```
0.000000000000000000001428571428571429
1.428571428571429e-21
14285714285714290000
1.428571428571429e20
```

- `\xintPFloatMinTrimmed` is the minimal number of trailing zeros which have to be present to activate actual trimming. The default definition is:

```
\def\xintPFloatMinTrimmed{4}
```

Defining it to expand to `-1` or `0` will enable the trimming of trailing zeros always, and setting it to a value at least equal to `P` (or `\xinttheDigits` if no `[P]`) will prevent it altogether.

This setting is ignored for the case of an integer value, if the criteria for using fixed point notation are met, and for the case of a one-digit mantissa in scientific notation.

To mimick approximately the Python behaviour in interactive sessions, one can use the following configuration:

```
\def\xintPFloatZero{0.0}%
\def\xintPFloatIntSuffix{.0}%
\def\xintPFloatLengthOneSuffix{.0}%
\def\xintPFloatNoSciEmax{15}%
\def\xintPFloatNoSciEmin{-4}%
\def\xintPFloatMinTrimmed{-1}%
```

```
0. → 0.0
1234. → 1234.0
6e100 → 6.0e100
1234567812345678.12345678 → 1234567812345678.0
12345678123456781.2345678 → 1.234567812345678e16
12345678.12340 → 12345678.1234
12345678.123400 → 12345678.1234
0.1234567812345678 → 0.1234567812345678
0.00012345678123456785 → 0.0001234567812345679
0.000012345678123456785 → 1.234567812345679e-5
```

The above using the default `\xintDigits` setting of 16 digits. This can not naturally match exactly CPython which uses internally radix 2 not 10, and has (by default) mantissas with `53=1+52` bits.

Same, but playing with `xintsession` in its `&fp` mode:

```
>>> &fp
fp mode (16 digits)
>>> \\def\xintPFloatZero{0.0}
(executing \\def\xintPFloatZero {0.0} in background)
)
Runaway argument?
def\xintPFloatZero {0.0}\message {
}\xs_fetch_aa \endinput
! File ended while scanning use of \\.
<inserted text>
```

```

\par
<*> xintsession^^M

? S
OK, entering \scrollmode...

*\xintsession

You are back to the xintexpr interactive session!
(current mode: fp (Digits=16), with Digits=16)

">>> " means central computing is waiting for input
"... " means that multi-line input continues. Use `;' to terminate it.

Say `&bye' at any time to terminate the session and the TeX run.
>>> \def\xintPFloatZero{0.0}
(executing \def \xintPFloatZero {0.0} in background)

>>> \def\xintPFloatIntSuffix{.0}\def\xintPFloatLengthOneSuffix{.0}

(executing \def \xintPFloatIntSuffix {.0}\def \xintPFloatLengthOneSuffix {.0} i
n background)

>>> \def\xintPFloatNoSciEmax{15}\def\xintPFloatNoSciEmin{-4}

(executing \def \xintPFloatNoSciEmax {15}\def \xintPFloatNoSciEmin {-4} in back
ground)

>>> \def\xintPFloatMinTrimmed{-1}
(executing \def \xintPFloatMinTrimmed {-1} in background)

>>> 0., 1234., 6e100;
@_1 0.0, 1234.0, 6.0e100
>>> 1234567812345678.12345678;
@_2 1234567812345678.0
>>> 12345678123456781.2345678;
@_3 1.234567812345678e16
>>> 12345678.12340;
@_4 12345678.1234
>>> 12345678.123400;
@_5 12345678.1234
>>> 0.1234567812345678;
@_6 0.1234567812345678
>>> 0.00012345678123456785;
@_7 0.0001234567812345679
>>> 0.000012345678123456785;
@_8 1.234567812345679e-5
>>> &bye

```

This is with version **0.4alpha** (2021-11-01) of *xintsession*. Probably some ``magic'' shortcuts will be added in future to its interface for this kind of tasks, in place of the `\def`.

12.92. `\xintFloatAdd`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{Frac}}{f}$ ★ `\xintFloatAdd [P]{f}{g}` first replaces `f` and `g` with their float approximations `f'` and `g'` to `P` significant places or to the precision from `\xintDigits`. It then produces the sum `f'+g'`, correctly rounded to nearest with the same number of significant places.

12.93. `\xintFloatSub`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{Frac}}{f}$ ★ `\xintFloatSub [P]{f}{g}` first replaces `f` and `g` with their float approximations `f'` and `g'` to `P` significant places or to the precision from `\xintDigits`. It then produces the difference `f'-g'` correctly rounded to nearest `P`-float.

12.94. `\xintFloatMul`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{Frac}}{f}$ ★ `\xintFloatMul [P]{f}{g}` first replaces `f` and `g` with their float approximations `f'` and `g'` to `P` (or `\xinttheDigits`) significant places. It then correctly rounds the product `f'*g'` to nearest `P`-float.

See [subsection 8.2](#) for more.

It is obviously much needed that the author improves its algorithms to avoid going through the exact $2P$ or $2P-1$ digits before throwing to the waste-bin half of those digits !

12.95. `\xintFloatDiv`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{Frac}}{f}$ ★ `\xintFloatDiv [P]{f}{g}` first replaces `f` and `g` with their float approximations `f'` and `g'` to `P` (or `\xinttheDigits`) significant places. It then correctly rounds the fraction `f'/g'` to nearest `P`-float.

See [subsection 8.2](#) for more.

Notice in the special situation with `f` and `g` integers that `\xintFloatDiv [P]{f}{g}` will *not necessarily* give the correct rounding of the exact fraction `f/g`. Indeed the macro arguments are each first individually rounded to `P` digits of precision. The correct syntax to get the correctly rounded integer fraction `f/g` is `\xintFloat[P]{f/g}`.

12.96. `\xintFloatPow`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{num}}{X}$ ★ `\xintFloatPow [P]{f}{x}` uses either the optional argument `P` or in its absence the value of `\xinttheDigits`. It computes a floating approximation to `f^x`.

The exponent `x` will be handed over to a `\numexpr`, hence count registers are accepted on input for this `x`. And the absolute value `|x|` must obey the $\text{T}_{\text{E}}\text{X}$ bound.

The argument `f` is first rounded to `P` significant places to give `f'`. The output `Z` is such that the exact `f'^x` differs from `Z` by an absolute error less than $0.52 \text{ ulp}(Z)$.

```
\xintFloatPow [8]{3.1415}{1234567890}  
1.6122066e613749456
```

12.97. `\xintFloatPower`

$\frac{\text{num}}{X}$ $\frac{\text{Frac}}{f}$ $\frac{\text{Num}}{f}$ ★ `\xintFloatPower[P]{f}{g}` computes a floating point value `f^g` where the exponent `g` is not constrained to be at most the $\text{T}_{\text{E}}\text{X}$ bound `2147483647`. It may even be a fraction `A/B` but will be truncated to an integer. The exponent of the *output* however *must* at any rate obey the $\text{T}_{\text{E}}\text{X}$ bound.

The argument `f` is first rounded to `P` significant places to give `f'`. The output `Z` is then such that the exact `f'^g` differs from `Z` by an absolute error less than $0.52 \text{ ulp}(Z)$.

For integer exponents this is the support macro which is used for the \wedge (or $**$) infix operators in `\xintfloateval`, or also in `\xinteval` for very big integer exponents. It is also used in `\xintfloateval` and `\xinteval` for half-integer exponents, via a combination with the `\xintFloatSqrt` square-root extraction.

The macro itself was NOT modified at 1.4e: when used directly it still starts by truncating the exponent to an integer... As for other user-level floating-point macros, its output is handled by `\xintFloat`, i.e. it uses scientific notation.

The 0.52 ulp(Z) guaranteed error bound applies also to the `\xintfloateval` evaluations for the half-integer exponent case. It is valid only when `f` already had a mantissa of at most `P` digits and was not modified by the initial rounding done by the macro to reduce `f` to `P` digits.

The integer exponent `g` may have more than `P` (or `Digits`) digits, it is handled exactly. And as said above its absolute value may exceed the \TeX bound.

12.98. `\xintFloatSqrt`

$\frac{\text{num}}{\text{x}}$ $\frac{\text{Frac}}{f}$ ★

`\xintFloatSqrt[P]{f}` computes a floating point approximation of \sqrt{f} , either using the optional precision `P` or the value of `\xinttheDigits`.



More precisely since 1.2f the macro achieves so-called *correct rounding*: the produced value is the rounding to `P` significant places of the abstract exact value, if the input has itself at most `P` digits (and an arbitrary exponent).

```
\xintFloatSqrt [89]{10}\newline
\xintFloatSqrt [89]{100}\newline
\xintFloatSqrt [89]{123456789}\par
```

```
3.1622776601683793319988935444327185337195551393252168268575048527925944386392382213442481e0
1.0000000000000000000000000000000000000000000000000000000000000000000000000000000000000e1
1.111111060555555440541666143353469245878409860134351071458570675251471479496366736579136e4
```

And now some tests to check that correct rounding applies correctly (sic):

```
The argument has 16 digits, hence escapes initial rounding:\newline
\xintFloatSqrt {5625000075000001}\newline
This one gets rounded hence same value is computed:\newline
\xintFloatSqrt {5625000075000001.4}\newline
but actual value is more like:\newline
\xintFloatSqrt [24]{5625000075000001.4}\newline
\xintFloatSqrt [32]{5625000075000001.4}\newline
The argument has 48 digits, hence escapes initial rounding:\newline
\xintFloatSqrt [48]{562500000000000000000000750000000000000000001}\newline
\xintFloatSqrt [64]{562500000000000000000000750000000000000000001}\newline
\xintFloatSqrt [80]{562500000000000000000000750000000000000000001}\newline
```

The argument has 16 digits, hence escapes initial rounding:

```
7.500000050000000e7
```

This one gets rounded hence same value is computed:

```
7.500000050000000e7
```

but actual value is more like:

```
7.50000005000000076666666e7
```

```
7.5000000500000007666666615555556e7
```

The argument has 48 digits, hence escapes initial rounding:

```
7.5000000000000000000000005000000000000000000000000e23
```

```
7.5000000000000000000000005000000000000000000000000e23
```

```
7.500000000000000000000000500000000000000000000000049999999999999999999999999999996666667e23
```

(we observe in passing illustrations that rounding to nearest is not transitive.)

12.99. `\xintFloatFac`

$\left[\begin{smallmatrix} \text{num} \\ x \end{smallmatrix} \right] \begin{smallmatrix} \text{Num} \\ f \end{smallmatrix} \star$

`\xintFloatFac[P]{f}` returns the factorial with either `\xinttheDigits` or `P` digits of precision.

The exact theoretical value differs from the calculated one `Y` by an absolute error strictly less than `0.6 ulp(Y)`.

`$1000!\approx{\$}\xintFloatFac [30]{1000}`
`1000! ≈ 4.02387260077093773543702433923e2567` The computation proceeds via doing explicitly the product, as the Stirling formula cannot be used for lack so far of `exp/log`.

The maximal allowed argument is 99999999, but already 100000! currently takes, for 16 digits of precision, a few seconds on my laptop (it returns `2.824229407960348e456573`).

The `factorial` function is available in `\xintfloatexpr`:

`\xintthefloatexpr factorial(1000)\relax % same as 1000!`
`4.023872600770938e2567`

12.100. `\xintFloatBinomial`

$\left[\begin{smallmatrix} \text{num} \\ x \end{smallmatrix} \right] \begin{smallmatrix} \text{Num} \\ f \end{smallmatrix} \begin{smallmatrix} \text{Num} \\ f \end{smallmatrix} \star$

`\xintFloatBinomial[P]{x}{y}` computes binomial coefficients with either `\xinttheDigits` or `P` digits of precision.

When `x < 0` an out-of-range error is raised. Else if `y < 0` or if `x < y` the macro evaluates to `0.0e0`. The exact theoretical value differs from the calculated one `Y` by an absolute error strictly less than `0.6 ulp(Y)`.

`#{3000\choose 1500}\approx{\$}\xintFloatBinomial [24]{3000}{1500}`
 $\binom{3000}{1500} \approx 1.79196793754756005073269e901$

The associated function in `\xintfloatexpr` is `binomial()`:

`\xintthefloatexpr binomial(3000,1500)\relax`
`1.791967937547560e901`

The computation is based on the formula $(x-y+1)\dots x/y!$ (here one arranges `y <= x-y` naturally).

12.101. `\xintFloatPFactorial`

$\left[\begin{smallmatrix} \text{num} \\ x \end{smallmatrix} \right] \begin{smallmatrix} \text{Num} \\ f \end{smallmatrix} \begin{smallmatrix} \text{Num} \\ f \end{smallmatrix} \star$

`\xintFloatPFactorial[P]{x}{y}` computes the product $(x+1)\dots y$.

The arguments must be integers (they are expanded inside `\numexpr`) and the allowed range is $-100000000 \leq x, y \leq 99999999$. If `x ≥ y` the product is considered empty hence returns one (as a floating point value). See also `\xintiipFactorial`.

The exact theoretical value differs from the calculated one `Y` by an absolute error strictly less than `0.6 ulp(Y)`.

The associated function in `\xintfloatexpr` is `pfactorial()`:

`\xintthefloatexpr pfactorial(2500,5000)\relax`
`2.595989917947957e8914`

13. Macros of the *xintbinhex* package

This package provides expandable conversions of (big) integers to and from binary and hexadecimal.

First version of this package was in the 1.08 (2013/06/07) release of *xint*. Its routines remained un-modified until their complete rewrite at release 1.2m (2017/07/31). The new macros are faster, using techniques from the 1.2 (2015/10/10) release of *xintcore*. But the inputs are now limited to a few thousand digits, whereas the 1.08 could handle (slowly...) tens of thousands of digits.

Table 3 recapitulates the maximal allowed sizes (they got increased at 1.2n): for macro `\xintFooToBar` in the first column, the value in the second column is the maximal *N* such that `\edef\X{\xintFooToBar{<N digits>}}` does not raise an error with standard TeX memory parameters (input stack size=5000, expansion depth=10000, parameter stack size=10000). The tests were done with TL2017 and *etex*. Nested calls will allow slightly lesser values only. The third column gives the corresponding maximal size of output. The fourth column gives the TeX parameter cited in the error message when trying with *N*+1 digits.

	Max length of input	-> length of output	Limiting factor
<code>\xintDecToHex</code>	6014	4995	input stack size=5000
<code>\xintDecToBin</code>	6014	19979	input stack size=5000
<code>\xintHexToDec</code>	8298	9992	input stack size=5000
<code>\xintBinToDec</code>	19988	6017	input stack size=5000
<code>\xintBinToHex</code>	19988	4997	input stack size=5000
<code>\xintHexToBin</code>	4996	19984	input stack size=5000
<code>\xintCHexToBin</code>	4997	19988	input stack size=5000

Table 3: Maximal sizes of inputs (at 1.2n) for *xintbinhex* macros

Roughly, base 10 numbers are limited to 6000 digits, hexadecimal numbers to (almost) 5000 digits, and binary numbers to (almost) 20000 digits. With the surprising exception of `\xintHexToDec` which allows almost 8300 hexadecimal digits on input.

The argument is first *f-expanded*. It may optionally have a unique leading minus sign (a plus sign is not allowed), and leading zeroes.

An input (possibly signed) with no leading zeroes is guaranteed to give an output without leading zero, with the sole, deliberate, exception of `\xintCHexToBin`: from *N* hexadecimal digits it produces *4N* binary digits, hence possibly with up to three leading zeroes (if the input had none.)

Inputs with leading zeroes usually produce outputs with an unspecified, case-dependent, number of leading zeroes (`\xintBinToHex` always uses the minimal number of hexadecimal digits needed to represent the binary digits, inclusive of leading zeroes if present.)

The macros converting from binary or decimal are robust against non terminated inputs like `\the\numexpr 2+3` or `\the\mathcode`\-.`. The macro `\xintHexToDec` also but not `\xintHexToBin` and `\xintCHexToBin` (anyway there are no primitive in (e)-TeX to my knowledge which will generate hexadecimal digits and may force expansion of next token).

Hexadecimal digits *A..F* must be in uppercase. Category code for them on input may be *letter* or *other*. On output they are of category code *letter*, and in uppercase.

Low-level unrecoverable errors will happen if for example a supposedly binary input contains other digits than 0 and 1. Inputs can not start with a *0b*, *0x*, *#x*, " or similar prefix: only digits/letters according to the binary, decimal, or hexadecimal notation.

With this package loaded additionally to *xintexpr*, hexadecimal input is possible in expressions: simply by using the prefix ". Such hexadecimal numbers may have a fractional part. Lowercase hexadecimal letters are currently not recognized as such in expressions. Currently the *p* postfix notation from standard programming languages standing for an extra power of two multiplicand is not implemented.

.1 `\xintDecToHex` 194 | .2 `\xintDecToBin` 194

<p>.3 \xintHexToDec 194 .4 \xintBinToDec 194 .5 \xintBinToHex 194</p>	<p>.6 \xintHexToBin 195 .7 \xintCHexToBin 195</p>
---	--

13.1. \xintDecToHex

f ★ Converts from decimal to hexadecimal.

```
\xintDecToHex{2718281828459045235360287471352662497757247093699959574966967627724076630353
547594571382178525166427427466391932003}
->11A9397C66949A97051F7D0A817914E3E0B17C41B11C48BAEF2B5760BB38D272F46DCE46C6032936BF37DAC918
814C63
```

13.2. \xintDecToBin

f ★ Converts from decimal to binary.

```
\xintDecToBin{2718281828459045235360287471352662497757247093699959574966967627724076630353
547594571382178525166427427466391932003}
->1000110101001001110010111100011001101001010010011010100101110000010100011111011110100001
01010000001011110010001010011100011111000001011000101111100010000011011000100011100010010001
011101011101110010101101010110110000010111011001110001101001001110010111101000110110111001
110010001101100011000000011001010010011011010111110011011110110101100100100011000100000010
100110001100011
```

13.3. \xintHexToDec

f ★ Converts from hexadecimal to decimal.

```
\xintHexToDec{11A9397C66949A97051F7D0A817914E3E0B17C41B11C48BAEF2B5760BB38D272F46DCE46C603
2936BF37DAC918814C63}
->271828182845904523536028747135266249775724709369995957496696762772407663035354759457138217
8525166427427466391932003
```

13.4. \xintBinToDec

f ★ Converts from binary to decimal.

```
\xintBinToDec{100011010100100111001011110001100110100101001001101010010111000001010001111
10111110100001010100000010111100100010100111000111110000010110001011111000100000110110001000
1110001001000101110101110111001010110101011101100000101110110011100011010010011100101111010
001101101110011100100011011000110000000110010100100110110101111100110111101101011001001000
11000100000010100110001100011}
->271828182845904523536028747135266249775724709369995957496696762772407663035354759457138217
8525166427427466391932003
```

13.5. \xintBinToHex

f ★ Converts from binary to hexadecimal. The input is first zero-filled to **4N** binary digits, hence the output will have **N** hexadecimal digits (thus, if the input did not have a leading zero, the output will not either).

```
\xintBinToHex{100011010100100111001011110001100110100101001001101010010111000001010001111
10111110100001010100000010111100100010100111000111110000010110001011111000100000110110001000
1110001001000101110101110111001010110101011101100000101110110011100011010010011100101111010
001101101110011100100011011000110000000110010100100110110101111100110111101101011001001000
11000100000010100110001100011}
```

```
->11A9397C66949A97051F7D0A817914E3E0B17C41B11C48BAEF2B5760BB38D272F46DCE46C6032936BF37DAC918
814C63
```

13.6. \xintHexToBin

f ★ Converts from hexadecimal to binary. Up to three leading zeroes of the output are trimmed.

```
\xintHexToBin{11A9397C66949A97051F7D0A817914E3E0B17C41B11C48BAEF2B5760BB38D272F46DCE46C603
2936BF37DAC918814C63}
->100011010100100111001011110001100110100101001001101010010111000001010001111011110100001
0101000000101110010001010011100011111000001011000101111100010000011011000100011100010010001
0111010111011100101011010101110110000010111011001110001101001001110010111101000110110111001
110010001101100011000000011001010010011011010111110011011110110101100100100011000100000010
100110001100011
```

13.7. \xintCHexToBin

f ★ Converts from hexadecimal to binary. Same as `\xintHexToBin`, but an input with *N* hexadecimal digits will give an output with exactly *4N* binary digits, leading zeroes are not trimmed.

```
\xintCHexToBin{11A9397C66949A97051F7D0A817914E3E0B17C41B11C48BAEF2B5760BB38D272F46DCE46C60
32936BF37DAC918814C63}
->0001000011010100100111001011110001100110100101001001101010010111000001010001111011110100
0010101000000101110010001010011100011111000001011000101111100010000011011000100011100010010
0010111010111011100101011010101110110000010111011001110001101001001110010111101000110110111
001110010001101100011000000011001010010011011010111110011011110110101100100100011000100000
010100110001100011
```

This can be combined with `\xintBinToHex` for round-trips preserving leading zeroes for *4N* binary digits numbers, whereas using `\xintHexToBin` gives reproducing round-trips only for *4N* binary numbers not starting with `0000`. This zero-fills to *4N* digits the input, hence gives here a leading zero in output:

```
\xintBinToHex{0001111}\par
0F
```

Chaining, we end up with *4N-3* digits, as three binary zeroes are trimmed:

```
\xintHexToBin{\xintBinToHex{0001111}}\par
01111
```

But this will always reproduce the initial input zero-filled to length *4N*:

```
\xintCHexToBin{\xintBinToHex{0001111}}\par
00001111
```

Another example (visible space characters manually inserted):

```
000000001111101001010001  $\xrightarrow{\text{\xintBinToHex}}$  00FA51  $\xrightarrow{\text{\xintHexToBin}}$  000000001111101001010001
000000001111101001010001  $\xrightarrow{\text{\xintBinToHex}}$  00FA51  $\xrightarrow{\text{\xintCHexToBin}}$  000000001111101001010001
```

14. Macros of the *xintgcd* package

This package was included in the original release 1.0 (2013/03/28) of the *xint* bundle.

At 1.3d macros `\xintiiGCD` and `\xintiiLCM` are copied over to *xint*, hence `gcd()` and `lcm()` functions in `\xintiexpr` were available simply from loading only *xintexpr*, and the *xintgcd* dependency got removed.



From 1.1 to 1.3f the package loaded only *xintcore*, not *xint* and neither *xinttools*.
 But at 1.4 it loads automatically both *xint* and *xinttools* (the latter being a requirement since 1.09h of the `\xintTypesetEuclideanAlgorithm` and `\xintTypesetBezoutAlgorithm` macros).
 The macros `\xintiiGCD` and `\xintiiLCM` got relocated into *xint*. The macros `\xintGCD`, `\xintLCM`, `\xintGCDof`, and `\xintLCMof` are removed: *xintfrac* provides under these names more powerful macros handling general fractions and not only integers.

<p>.1 <code>\xintBezout</code>..... 196 .2 <code>\xintEuclideanAlgorithm</code> 196 .3 <code>\xintBezoutAlgorithm</code> 196</p>	<p>.4 <code>\xintTypesetEuclideanAlgorithm</code> 197 .5 <code>\xintTypesetBezoutAlgorithm</code>..... 197</p>
--	---

14.1. `\xintBezout`

Num Num f f ★ `\xintBezout{N}{M}` returns three numbers *U*, *V*, *D* within braces where *D* is the (non-negative) GCD, and $UN + VM = D$.

```
\oodef\X{\xintBezout {10000}{1113}}\meaning\X\par
\xintAssign {\xintBezout {10000}{1113}}\to\U\V\D
U: \meaning\U, V: \meaning\V, D: \meaning\D\par
AU+BV: \xinttheiexpr 10000*\U+1113*\V\relax\par
\noindent\oodef\X{\xintBezout {123456789012345}{9876543210321}}\meaning\X\par
\xintAssign \X\to\U\V\D
U: \meaning\U, V: \meaning\V, D: \meaning\D\par
AU+BV: \xinttheiexpr 123456789012345*\U+9876543210321*\V\relax
macro:->{-131}{1177}{1}
U: macro:->-131, V: macro:->1177, D: macro:->1
AU+BV: 1
macro:->{256654313730}{-3208178892607}{3}
U: macro:->256654313730, V: macro:->-3208178892607, D: macro:->3
AU+BV: 3
```

14.2. `\xintEuclideanAlgorithm`

Num Num f f ★ `\xintEuclideanAlgorithm{N}{M}` applies the Euclidean algorithm and keeps a copy of all quotients and remainders.

```
\edef\X{\xintEuclideanAlgorithm {10000}{1113}}\meaning\X
macro:->{5}{10000}{1}{1113}{8}{1096}{1}{17}{64}{8}{2}{1}{8}{0}
The first item is the number of steps, the second is N, the third is the GCD, the fourth is M then
the first quotient and remainder, the second quotient and remainder, . . . until the final quotient
and last (zero) remainder.
```

14.3. `\xintBezoutAlgorithm`

Num Num f f ★ `\xintBezoutAlgorithm{N}{M}` applies the Euclidean algorithm and keeps a copy of all quotients and

remainders. Furthermore it computes the entries of the successive products of the 2 by 2 matrices $\begin{pmatrix} q & 1 \\ 1 & 0 \end{pmatrix}$ formed from the quotients arising in the algorithm.

```
\edef\X{\xintBezoutAlgorithm {10000}{1113}}\printnumber{\meaning\X}
macro:->{5}{10000}{0}{1}{1}{1113}{1}{0}{8}{1096}{8}{1}{1}{17}{9}{1}{64}{8}{584}{65}{2}{1}{11
77}{131}{8}{0}{10000}{1113}
```

The first item is the number of steps, the second is *N*, then 0, 1, the GCD, *M*, 1, 0, the first quotient, the first remainder, the top left entry of the first matrix, the bottom left entry, and then these four things at each step until the end.

14.4. `\xintTypesetEuclideanAlgorithm`

Num Num
f f

This macro is just an example of how to organize the data returned by `\xintEuclideanAlgorithm`. Copy the source code to a new macro and modify it to what is needed.

```
\xintTypesetEuclideanAlgorithm {123456789012345}{9876543210321}
123456789012345 = 12 × 9876543210321 + 4938270488493
9876543210321 = 2 × 4938270488493 + 2233335
4938270488493 = 2211164 × 2233335 + 536553
2233335 = 4 × 536553 + 87123
536553 = 6 × 87123 + 13815
87123 = 6 × 13815 + 4233
13815 = 3 × 4233 + 1116
4233 = 3 × 1116 + 885
1116 = 1 × 885 + 231
885 = 3 × 231 + 192
231 = 1 × 192 + 39
192 = 4 × 39 + 36
39 = 1 × 36 + 3
36 = 12 × 3 + 0
```

14.5. `\xintTypesetBezoutAlgorithm`

Num Num
f f

This macro is just an example of how to organize the data returned by `\xintBezoutAlgorithm`. Copy the source code to a new macro and modify it to what is needed.

```
\xintTypesetBezoutAlgorithm {10000}{1113}
10000 = 8 × 1113 + 1096
8 = 8 × 1 + 0
1 = 8 × 0 + 1
1113 = 1 × 1096 + 17
9 = 1 × 8 + 1
1 = 1 × 1 + 0
1096 = 64 × 17 + 8
584 = 64 × 9 + 8
65 = 64 × 1 + 1
17 = 2 × 8 + 1
1177 = 2 × 584 + 9
131 = 2 × 65 + 1
8 = 8 × 1 + 0
10000 = 8 × 1177 + 584
1113 = 8 × 131 + 65
131 × 10000 - 1177 × 1113 = -1
```

15. Macros of the *xintseries* package

This package was first released with version 1.03 (2013/04/14) of the *xint* bundle.

The $\frac{\text{Frac}}{f}$ expansion type of various macro arguments is only a $\frac{\text{Num}}{f}$ if only *xint* but not *xintfrac* is loaded. The macro `\xintiSeries` is special and expects summing big integers obeying the strict format, even if *xintfrac* is loaded.

The arguments serving as indices are of the $\frac{\text{num}}{X}$ expansion type.

In some cases one or two of the macro arguments are only expanded at a later stage not immediately.

Since 1.3, `\xintAdd` and `\xintSub` use systematically the least common multiple of the denominators. Some of the comments in this chapter refer to the earlier situation where often the denominators were simply multiplied together. *They have yet to be updated to reflect the new situation brought by the 1.3 release.* Some of these comments may now be off-synced from the actual computation results and thus may be wrong.

<p>.1 <code>\xintSeries</code>..... 198</p> <p>.2 <code>\xintiSeries</code> 199</p> <p>.3 <code>\xintRationalSeries</code>..... 200</p> <p>.4 <code>\xintRationalSeriesX</code> 203</p> <p>.5 <code>\xintPowerSeries</code> 205</p> <p>.6 <code>\xintPowerSeriesX</code>..... 206</p>	<p>.7 <code>\xintFxFtPowerSeries</code> 207</p> <p>.8 <code>\xintFxFtPowerSeriesX</code> 207</p> <p>.9 <code>\xintFloatPowerSeries</code> 209</p> <p>.10 <code>\xintFloatPowerSeriesX</code> 209</p> <p>.11 Computing log 2 and π..... 209</p>
---	---

15.1. `\xintSeries`

$\frac{\text{num}}{X} \frac{\text{num}}{X} \frac{\text{Frac}}{f} \star$ `\xintSeries{A}{B}{\coeff}` computes $\sum_{n=A}^{n=B} \coeff{n}$. The initial and final indices must obey the `\n`, `umexpr` constraint of expanding to numbers at most $2^{31}-1$. The `\coeff` macro must be a one-parameter *f-expandable* macro, taking on input an explicit number *n* and producing some number or fraction `\coeff{n}`; it is expanded at the time it is needed.

```
\def\coeff #1{\xintiiMON{#1}/#1.5} % (-1)^n/(n+1/2)
\edef\w {\xintSeries {0}{50}{\coeff}} % we want to re-use it
\edef\z {\xintJrr {\w}[0]} % the [0] for a microsecond gain.
% \xintJrr preferred to \xintIrr: a big common factor is suspected.
% But numbers much bigger would be needed to show the greater efficiency.
\[\sum_{n=0}^{n=50} \frac{(-1)^n}{n+\frac{1}{2}} = \xintTeXFrac\z \]
```

$$\sum_{n=0}^{n=50} \frac{(-1)^n}{n + \frac{1}{2}} = \frac{173909338287370940432112792101626602278714}{110027467159390003025279917226039729050575}$$

The definition of `\coeff` as `\xintiiMON{#1}/#1.5` is quite suboptimal. It allows *#1* to be a big integer, but anyhow only small integers are accepted as initial and final indices (they are of the $\frac{\text{num}}{X}$ type). Second, when the *xintfrac* parser sees the *#1.5* it will remove the dot hence create a denominator with one digit more. For example $1/3.5$ turns internally into $10/35$ whereas it would be more efficient to have $2/7$. For info here is the non-reduced `\w`:

$$\frac{86954669143685470216056396050813301139357}{550137335796950015126399586130198645252875} 10^1$$

It would have been bigger still in releases earlier than 1.1: now, the *xintfrac* `\xintAdd` routine does not multiply blindly denominators anymore, it checks if one is a multiple of the other. However it does not practice systematic reduction to lowest terms.

A more efficient way to code `\coeff` is illustrated next.

```
\def\coeff #1{\the\numexpr\ifodd #1 -2\else2\fi\relax/\the\numexpr 2*#1+1\relax [0]}%
% The [0] in \coeff is a tiny optimization: in its presence the \xintfracname parser
% sees something which is already in internal format.
\fddefw {\xintSeries {0}{50}{\coeff}}
\[\sum_{n=0}^{n=50} \frac{(-1)^n}{n+\frac{1}{2}}=\xintTeXFrac{w\}
```

$$\sum_{n=0}^{n=50} \frac{(-1)^n}{n + \frac{1}{2}} = \frac{173909338287370940432112792101626602278714}{110027467159390003025279917226039729050575}$$

The reduced form `\z` as displayed above only differs from this one by a factor of 1.

```
\def\coeffleibnitz #1{\the\numexpr\ifodd #1 1\else-1\fi\relax/#1[0]}
\cnta 1
\loop
% in this loop we recompute from scratch each partial sum!
% we can afford that, as \xintSeries is fast enough.
\noindent\hbox to 2em{\hfil\texttt{\the\cnta.} }%
    \xintTrunc {12}{\xintSeries {1}{\cnta}{\coeffleibnitz}}\dots
\endgraf
\ifnum\cnta < 30 \advance\cnta 1 \repeat
```

1. 1.000000000000...	11. 0.736544011544...	21. 0.716390450794...
2. 0.500000000000...	12. 0.653210678210...	22. 0.670935905339...
3. 0.833333333333...	13. 0.730133755133...	23. 0.714414166209...
4. 0.583333333333...	14. 0.658705183705...	24. 0.672747499542...
5. 0.783333333333...	15. 0.725371850371...	25. 0.712747499542...
6. 0.616666666666...	16. 0.662871850371...	26. 0.674285961081...
7. 0.759523809523...	17. 0.721695379783...	27. 0.711322998118...
8. 0.634523809523...	18. 0.666139824228...	28. 0.675608712404...
9. 0.745634920634...	19. 0.718771403175...	29. 0.710091471024...
10. 0.645634920634...	20. 0.668771403175...	30. 0.676758137691...

15.2. `\xintiSeries`

$\sum_{n=A}^{n=B} f$ `\xintiSeries{A}{B}{\coeff}` computes $\sum_{n=A}^{n=B} \text{\coeff}$ where `\coeff` must *f-expand* to a (possibly long) integer in the strict format.

```
\def\coeff #1{\xintiTrunc {40}{\xintiMON{#1}/#1.5}}%
% better:
\def\coeff #1{\xintiTrunc {40}
  {\the\numexpr 2*\xintiMON{#1}\relax/\the\numexpr 2*#1+1\relax [0]}}%
% better still:
\def\coeff #1{\xintiTrunc {40}
  {\the\numexpr\ifodd #1 -2\else2\fi\relax/\the\numexpr 2*#1+1\relax [0]}}%
% (-1)^n/(n+1/2) times 10^40, truncated to an integer.
\[\sum_{n=0}^{n=50} \frac{(-1)^n}{n+\frac{1}{2}} \approx
  \xintTrunc {40}{\xintiSeries {0}{50}{\coeff}[-40]}\dots\]
```

$$\sum_{n=0}^{n=50} \frac{(-1)^n}{n + \frac{1}{2}} \approx 1.5805993064935250412367895069567264144810$$

We should have cut out at least the last two digits: truncating errors originating with the first coefficients of the sum will never go away, and each truncation introduces an uncertainty in the last digit, so as we have 40 terms, we should trash the last two digits, or at least round at 38 digits. It is interesting to compare with the computation where rounding rather than truncation is used, and with the decimal expansion of the exactly computed partial sum of the series:

```

\def\coeff #1{\xintiRound {40} % rounding at 40
  {\the\numexpr\ifodd #1 -2\else2\fi\relax/\the\numexpr 2*#1+1\relax [0]}}%
% (-1)^n/(n+1/2) times 10^40, rounded to an integer.
\[\sum_{n=0}^{n=50} \frac{(-1)^n}{n+\frac{1}{2}} \approx
  \xintTrunc {40}{\xintiSeries {0}{50}{\coeff}[-40]}\]
\def\exactcoeff #1%
  {\the\numexpr\ifodd #1 -2\else2\fi\relax/\the\numexpr 2*#1+1\relax [0]}}%
\[\sum_{n=0}^{n=50} \frac{(-1)^n}{n+\frac{1}{2}}
  = \xintTrunc {50}{\xintiSeries {0}{50}{\exactcoeff}}\dots\]

```

$$\sum_{n=0}^{n=50} \frac{(-1)^n}{n + \frac{1}{2}} \approx 1.5805993064935250412367895069567264144804$$

$$\sum_{n=0}^{n=50} \frac{(-1)^n}{n + \frac{1}{2}} = 1.58059930649352504123678950695672641448068680288367\dots$$

This shows indeed that our sum of truncated terms estimated wrongly the 39th and 40th digits of the exact result⁷³ and that the sum of rounded terms fared a bit better.

15.3. `\xintRationalSeries`

`\xintRationalSeries{A}{B}{f}{ratio}` evaluates $\sum_{n=A}^{n=B} F(n)$, where $F(n)$ is specified indirectly via the data of $f=F(A)$ and the one-parameter macro `\ratio` which must be such that `\macro{n}` expands to $F(n)/F(n-1)$. The name indicates that `\xintRationalSeries` was designed to be useful in the cases where $F(n)/F(n-1)$ is a rational function of n but it may be anything expanding to a fraction. The macro `\ratio` must be an expandable-only compatible macro and expand to its value after iterated full expansion of its first item. A and B are fed to a `\numexpr` hence may be count registers or arithmetic expressions built with such; they must obey the \TeX bound. The initial term f may be a macro `\f`, it will be expanded to its value representing $F(A)$.

```

\def\ratio #1{2/#1[0]}% 2/n, to compute exp(2)
\cnta 0 % previously declared count
\begin{quote}
\loop \fdef\z {\xintRationalSeries {0}{\cnta}{1}{ratio }}%
\noindent$\sum_{n=0}^{\the\cnta} \frac{2^n}{n!}$=
  \xintTrunc{12}\z\dots=
  \xintTeXFrac\z=\xintTeXFrac{\xintIrr\z}$\vtop to 5pt{}\par
\ifnum\cnta<20 \advance\cnta 1 \repeat
\end{quote}

```

$$\sum_{n=0}^0 \frac{2^n}{n!} = 1.000000000000\dots = 1 = 1$$

$$\sum_{n=0}^1 \frac{2^n}{n!} = 3.000000000000\dots = 3 = 3$$

$$\sum_{n=0}^2 \frac{2^n}{n!} = 5.000000000000\dots = \frac{10}{2} = 5$$

$$\sum_{n=0}^3 \frac{2^n}{n!} = 6.333333333333\dots = \frac{38}{6} = \frac{19}{3}$$

$$\sum_{n=0}^4 \frac{2^n}{n!} = 7.000000000000\dots = \frac{168}{24} = 7$$

$$\sum_{n=0}^5 \frac{2^n}{n!} = 7.266666666666\dots = \frac{872}{120} = \frac{109}{15}$$

$$\sum_{n=0}^6 \frac{2^n}{n!} = 7.355555555555\dots = \frac{5296}{720} = \frac{331}{45}$$

$$\sum_{n=0}^7 \frac{2^n}{n!} = 7.380952380952\dots = \frac{37200}{5040} = \frac{155}{21}$$

⁷³ as the series is alternating, we can roughly expect an error of $\sqrt{40}$ and the last two digits are off by 4 units, which is not contradictory to our expectations.

$$\sum_{n=0}^8 \frac{2^n}{n!} = 7.387301587301\dots = \frac{297856}{40320} = \frac{2327}{315}$$

$$\sum_{n=0}^9 \frac{2^n}{n!} = 7.388712522045\dots = \frac{2681216}{362880} = \frac{20947}{2835}$$

$$\sum_{n=0}^{10} \frac{2^n}{n!} = 7.388994708994\dots = \frac{26813184}{3628800} = \frac{34913}{4725}$$

$$\sum_{n=0}^{11} \frac{2^n}{n!} = 7.389046015712\dots = \frac{294947072}{39916800} = \frac{164591}{22275}$$

$$\sum_{n=0}^{12} \frac{2^n}{n!} = 7.389054566832\dots = \frac{3539368960}{479001600} = \frac{691283}{93555}$$

$$\sum_{n=0}^{13} \frac{2^n}{n!} = 7.389055882389\dots = \frac{46011804672}{6227020800} = \frac{14977801}{2027025}$$

$$\sum_{n=0}^{14} \frac{2^n}{n!} = 7.389056070325\dots = \frac{644165281792}{87178291200} = \frac{314533829}{42567525}$$

$$\sum_{n=0}^{15} \frac{2^n}{n!} = 7.389056095384\dots = \frac{9662479259648}{1307674368000} = \frac{4718007451}{638512875}$$

$$\sum_{n=0}^{16} \frac{2^n}{n!} = 7.389056098516\dots = \frac{154599668219904}{20922789888000} = \frac{1572669151}{212837625}$$

$$\sum_{n=0}^{17} \frac{2^n}{n!} = 7.389056098884\dots = \frac{2628194359869440}{355687428096000} = \frac{16041225341}{2170943775}$$

$$\sum_{n=0}^{18} \frac{2^n}{n!} = 7.389056098925\dots = \frac{47307498477912064}{6402373705728000} = \frac{103122162907}{13956067125}$$

$$\sum_{n=0}^{19} \frac{2^n}{n!} = 7.389056098930\dots = \frac{898842471080853504}{121645100408832000} = \frac{4571749222213}{618718975875}$$

$$\sum_{n=0}^{20} \frac{2^n}{n!} = 7.389056098930\dots = \frac{17976849421618118656}{2432902008176640000} = \frac{68576238333199}{9280784638125}$$

```
\def\ratio #1{-1/#1[0]}% -1/n, comes from the series of exp(-1)
\cnta 0 % previously declared count
\begin{quote}
\loop
\fddef\z {\xintRationalSeries {0}{\cnta}{1}{\ratio }}%
\noindent$\sum_{n=0}^{\the\cnta} \frac{(-1)^n}{n!}=
\xintTrunc{20}\z\dots=\xintTeXFrac{\z}=\xintTeXFrac{\xintIrr\z}$%
\vtop to 5pt{}\par
\ifnum\cnta<20 \advance\cnta 1 \repeat
\end{quote}
```

$$\sum_{n=0}^0 \frac{(-1)^n}{n!} = 1.00000000000000000000\dots = 1 = 1$$

$$\sum_{n=0}^1 \frac{(-1)^n}{n!} = 0\dots = 0 = 0$$

$$\sum_{n=0}^2 \frac{(-1)^n}{n!} = 0.50000000000000000000\dots = \frac{1}{2} = \frac{1}{2}$$

$$\sum_{n=0}^3 \frac{(-1)^n}{n!} = 0.33333333333333333333\dots = \frac{2}{6} = \frac{1}{3}$$

$$\sum_{n=0}^4 \frac{(-1)^n}{n!} = 0.37500000000000000000\dots = \frac{9}{24} = \frac{3}{8}$$

$$\sum_{n=0}^5 \frac{(-1)^n}{n!} = 0.36666666666666666666\dots = \frac{44}{120} = \frac{11}{30}$$

$$\sum_{n=0}^6 \frac{(-1)^n}{n!} = 0.36805555555555555555\dots = \frac{265}{720} = \frac{53}{144}$$

$$\sum_{n=0}^7 \frac{(-1)^n}{n!} = 0.36785714285714285714\dots = \frac{1854}{5040} = \frac{103}{280}$$

$$\sum_{n=0}^8 \frac{(-1)^n}{n!} = 0.36788194444444444444\dots = \frac{14833}{40320} = \frac{2119}{5760}$$

$$\sum_{n=0}^9 \frac{(-1)^n}{n!} = 0.36787918871252204585\dots = \frac{133496}{362880} = \frac{16687}{45360}$$

$$\sum_{n=0}^{10} \frac{(-1)^n}{n!} = 0.36787946428571428571\dots = \frac{1334961}{3628800} = \frac{16481}{44800}$$

$$\sum_{n=0}^{11} \frac{(-1)^n}{n!} = 0.36787943923360590027\dots = \frac{14684570}{39916800} = \frac{1468457}{3991680}$$

$$\sum_{n=0}^{12} \frac{(-1)^n}{n!} = 0.36787944132128159905\dots = \frac{176214841}{479001600} = \frac{16019531}{43545600}$$


```

\sum_{n=0}^{\the\numexpr 2*\cnta-1\relax} \frac{\the\cnta^n}{n!} =
\intTrunc{8}{\intDiv{z}{w}\dots$ \vtop to 5pt}}\endgraf
\ifnum\cnta<20 \advance\cnta 1 \repeat
\end{multicols}

```

$\sum_{n=1}^1 \frac{1^n}{n!} / \sum_{n=0}^1 \frac{1^n}{n!} = 0.50000000\dots$ $\sum_{n=2}^3 \frac{2^n}{n!} / \sum_{n=0}^3 \frac{2^n}{n!} = 0.52631578\dots$ $\sum_{n=3}^5 \frac{3^n}{n!} / \sum_{n=0}^5 \frac{3^n}{n!} = 0.53804347\dots$ $\sum_{n=4}^7 \frac{4^n}{n!} / \sum_{n=0}^7 \frac{4^n}{n!} = 0.54317053\dots$ $\sum_{n=5}^9 \frac{5^n}{n!} / \sum_{n=0}^9 \frac{5^n}{n!} = 0.54502576\dots$ $\sum_{n=6}^{11} \frac{6^n}{n!} / \sum_{n=0}^{11} \frac{6^n}{n!} = 0.54518217\dots$ $\sum_{n=7}^{13} \frac{7^n}{n!} / \sum_{n=0}^{13} \frac{7^n}{n!} = 0.54445274\dots$ $\sum_{n=8}^{15} \frac{8^n}{n!} / \sum_{n=0}^{15} \frac{8^n}{n!} = 0.54327992\dots$ $\sum_{n=9}^{17} \frac{9^n}{n!} / \sum_{n=0}^{17} \frac{9^n}{n!} = 0.54191055\dots$ $\sum_{n=10}^{19} \frac{10^n}{n!} / \sum_{n=0}^{19} \frac{10^n}{n!} = 0.54048295\dots$	$\sum_{n=11}^{21} \frac{11^n}{n!} / \sum_{n=0}^{21} \frac{11^n}{n!} = 0.53907332\dots$ $\sum_{n=12}^{23} \frac{12^n}{n!} / \sum_{n=0}^{23} \frac{12^n}{n!} = 0.53772178\dots$ $\sum_{n=13}^{25} \frac{13^n}{n!} / \sum_{n=0}^{25} \frac{13^n}{n!} = 0.53644744\dots$ $\sum_{n=14}^{27} \frac{14^n}{n!} / \sum_{n=0}^{27} \frac{14^n}{n!} = 0.53525726\dots$ $\sum_{n=15}^{29} \frac{15^n}{n!} / \sum_{n=0}^{29} \frac{15^n}{n!} = 0.53415135\dots$ $\sum_{n=16}^{31} \frac{16^n}{n!} / \sum_{n=0}^{31} \frac{16^n}{n!} = 0.53312615\dots$ $\sum_{n=17}^{33} \frac{17^n}{n!} / \sum_{n=0}^{33} \frac{17^n}{n!} = 0.53217628\dots$ $\sum_{n=18}^{35} \frac{18^n}{n!} / \sum_{n=0}^{35} \frac{18^n}{n!} = 0.53129566\dots$ $\sum_{n=19}^{37} \frac{19^n}{n!} / \sum_{n=0}^{37} \frac{19^n}{n!} = 0.53047810\dots$ $\sum_{n=20}^{39} \frac{20^n}{n!} / \sum_{n=0}^{39} \frac{20^n}{n!} = 0.52971771\dots$
--	---

15.4. \xintRationalSeriesX

$\sum \frac{f}{f} \frac{f}{f} \frac{f}{f} \star$

`\xintRationalSeriesX{A}{B}{\first}{\ratio}{\g}` is a parametrized version of `\xintRationalSeries` where `\first` is now a one-parameter macro such that `\first{\g}` gives the initial term and `\ratio` is a two-parameter macro such that `\ratio{n}{\g}` represents the ratio of one term to the previous one. The parameter `\g` is evaluated only once at the beginning of the computation, and can thus itself be the yet unevaluated result of a previous computation.

```

Let \ratio be such a two-parameter macro; note the subtle differences between
\xintRationalSeries {A}{B}{\first}{\ratio{\g}}
and \xintRationalSeriesX {A}{B}{\first}{\ratio}{\g}.

```

First the location of braces differ... then, in the former case `\first` is a *no-parameter* macro expanding to a fractional number, and in the latter, it is a *one-parameter* macro which will use `\g`. Furthermore the **X** variant will expand `\g` at the very beginning whereas the former non-**X** former variant will evaluate it each time it needs it (which is bad if this evaluation is time-costly, but good if `\g` is a big explicit fraction encapsulated in a macro).

The example will use the macro `\xintPowerSeries` which computes efficiently exact partial sums of power series, and is discussed in the next section.

```

\def\firstterm #1{1[0]}% first term of the exponential series
% although it is the constant 1, here it must be defined as a
% one-parameter macro. Next comes the ratio function for exp:
\def\ratioexp #1#2{\intDiv {#1}{#2}}% x/n
% These are the (-1)^{n-1}/n of the log(1+h) series:
\def\coefflog #1{\the\numexpr\ifodd #1 1\else-1\fi\relax/#1[0]}%
% Let L(h) be the first 10 terms of the log(1+h) series and
% let E(t) be the first 10 terms of the exp(t) series.
% The following computes E(L(a/10)) for a=1,...,12.
\begin{multicols}{3}\raggedcolumns
\cnta 0
\loop
\noindent\xintTrunc {18}{%
\xintRationalSeriesX {0}{9}{\firstterm}{\ratioexp}
{\xintPowerSeries{1}{10}{\coefflog}{\the\cnta[-1]}}\dots
}\endgraf
\end{multicols}

```

```
\ifnum\cnta < 12 \advance \cnta 1 \repeat
\end{multicols}
```

1.00000000000000000000...	1.499954310225476533...	1.907197560339468199...
1.099999999999083906...	1.599659266069210466...	1.845117565491393752...
1.19999998111624029...	1.698137473697423757...	1.593831932293536053...
1.29999835744121464...	1.791898112718884531...	
1.399996091955359088...	1.870485649686617459...	

These completely exact operations rapidly create numbers with many digits. Let us print in full the raw fractions created by the operation illustrated above:

$E(L(1[-1]))=163591443693117889303431088806087634148250735791023497657261314014159107395739$
 $0639913787199465741057336677116573252341295218688/148719494266594663864467456000000000[-90]$
 (length of numerator: 127)

$E(L(12[-2]))=16656583357757234467643895619026874191327320993157183247568125775059356018362$
 $23193439604540053754226444871502834816644808336288211299845887246066795041160882231219805166$
 $927273729660728412213074817261522841754729971712/148719494266594663864467456000000000[-180]$
 (length of numerator: 217)

$E(L(123[-3]))=1670119920600555026998663239069002278266215966968669508145191626887938734862$
 $73269986546078658803979014003116903378025935148900448814698936627633558066738151958530603167$
 $40612785673175692992742863679398303407413205084692383474722719804622771982161117197045873620$
 $25769049115687215712723182386527055033735053312/148719494266594663864467456000000000[-270]$ (length of numerator: 307)

We see that the denominators here remain the same, as our input only had various powers of ten as denominators, and *xintfrac* efficiently assemble (some only, as we can see) powers of ten. Notice that 1 more digit in an input denominator seems to mean 90 more in the raw output. We can check that with some other test cases:

$E(L(1/7))=48228203862750885848048032297655163193719083498752126622944863683478921463537652$
 $0421966954177876452794933/421996783816227187824437252776031227863306380633210580813174165609$
 $50056936721328812056161288192000000000[0]$ (length of numerator: 105; length of denominator: 105)

$E(L(1/71))=6190039670785350346406550995159476540272948182884398462882888922997238323197498$
 $85971940015218249059435720836832839237391067287499316324605873244670430502854291696282116287$
 $58603878135499973539887212860467/61040668975799982582643863990035761895246771100033570420271$
 $6710922033298498184289107451083577982695694446256675834390041749715017225626389830761170775$
 $791999877852355941834008347347315123552256000000000[0]$ (length of numerator: 203; length of denominator: 203)

$E(L(1/712))=300356435377840602055967040841188592538909311419930838799656013626071029784174$
 $49681929088495804136203813242174405561415315426829241317287053037273453329055814153891517325$
 $75694112320026364569495366534918031439051104610487529796192058205725999641657806615904929048$
 $98946463533146662233869249/29993517810522090976696848959176310536177550755703969736435921535$
 $22460410892328532539738041911202121412424715881734049254716640082470987340985151932504281494$
 $24064596788874441470533147848207863549778847000617103264666638782677019019130113930837421531$
 $81047806202596610291401752576000000000[0]$ (length of numerator: 288; length of denominator: 288)

Thus decimal numbers such as 0.123 (equivalently 123[-3]) give less computing intensive tasks than fractions such as 1/712: in the case of decimal numbers the (raw) denominators originate in the coefficients of the series themselves, powers of ten of the input within brackets being treated separately. And even then the numerators will grow with the size of the input in a sort of linear way, the coefficient being given by the order of series: here 10 from the log and 9 from the exp, so 90. One more digit in the input means 90 more digits in the numerator of the output: obviously we can not go on composing such partial sums of series and hope that *xint* will joyfully do all at the speed of light!

Hence, truncating the output (or better, rounding) is the only way to go if one needs a general calculus of special functions. This is why the package `xintseries` provides, besides `\xintSeries`, `\xintRationalSeries`, or `\xintPowerSeries` which compute exact sums, `\xintFxFtPowerSeries` for fixed-point computations and a (tentative naive) `\xintFloatPowerSeries`.

15.5. `\xintPowerSeries`

`\xintPowerSeries{A}{B}{\coeff}{f}` evaluates the sum $\sum_{n=A}^{n=B} \text{\coeff}\{n\} \cdot f^n$. The initial and final indices are given to a `\numexpr` expression. The `\coeff` macro (which, as argument to `\xintPowerSeries` is expanded only at the time `\coeff{n}` is needed) should be defined as a one-parameter expandable macro, its input will be an explicit number.

The `f` can be either a fraction directly input or a macro `\f` expanding to such a fraction. It is actually more efficient to encapsulate an explicit fraction `f` in such a macro, if it has big numerators and denominators ('big' means hundreds of digits) as it will then take less space in the processing until being (repeatedly) used.

This macro computes the exact result (one can use it also for polynomial evaluation), using a Horner scheme which helps avoiding a denominator build-up (this problem however, even if using a naive additive approach, is much less acute since release 1.1 and its new policy regarding `\xintAdd`).

```
\def\geom #1{1[0]} % the geometric series
\def\f {5/17[0]}
\[ \sum_{n=0}^{n=20} \Bigl(\frac{5}{17}\Bigr)^n
=\xintTeXFrac{\xintIrr{\xintPowerSeries {0}{20}{\geom}{\f}}}
=\xintTeXFrac{\xinttheexpr (17^21-5^21)/12/17^20\relax}\]
```

$$\sum_{n=0}^{n=20} \left(\frac{5}{17}\right)^n = \frac{5757661159377657976885341}{4064231406647572522401601} = \frac{69091933912531895722624092}{48770776879770870268819212}$$

```
\def\coefflog #1{1/#1[0]}% 1/n
\def\f {1/2[0]}%
\[ \log 2 \approx \sum_{n=1}^{n=20} \frac{1}{n \cdot 2^n}
=\xintTeXFrac {\xintIrr {\xintPowerSeries {1}{20}{\coefflog}{\f}}}\]
\[ \log 2 \approx \sum_{n=1}^{n=50} \frac{1}{n \cdot 2^n}
=\xintTeXFrac {\xintIrr {\xintPowerSeries {1}{50}{\coefflog}{\f}}}\]
```

$$\log 2 \approx \sum_{n=1}^{20} \frac{1}{n \cdot 2^n} = \frac{42299423848079}{61025172848640}$$

$$\log 2 \approx \sum_{n=1}^{50} \frac{1}{n \cdot 2^n} = \frac{60463469751752265663579884559739219}{87230347965792839223946208178339840}$$

```
\setlength{\columnsep}{0pt}
\begin{multicols}{3}
\cnta 1 % previously declared count
\loop % in this loop we recompute from scratch each partial sum!
% we can afford that, as \xintPowerSeries is fast enough.
\noindent\hbox to 2em{\hfil\texttt{\the\cnta.} }%
\xintTrunc {12}
{\xintPowerSeries {1}{\cnta}{\coefflog}{\f}}\dots
\endgraf
\ifnum \cnta < 30 \advance\cnta 1 \repeat
\end{multicols}
```

- | | | |
|-----------------------|-----------------------|-----------------------|
| 1. 0.500000000000... | 11. 0.693109245355... | 21. 0.693147159757... |
| 2. 0.625000000000... | 12. 0.693129590407... | 22. 0.693147170594... |
| 3. 0.666666666666... | 13. 0.693138980431... | 23. 0.693147175777... |
| 4. 0.682291666666... | 14. 0.693143340085... | 24. 0.693147178261... |
| 5. 0.688541666666... | 15. 0.693145374590... | 25. 0.693147179453... |
| 6. 0.691145833333... | 16. 0.693146328265... | 26. 0.693147180026... |
| 7. 0.692261904761... | 17. 0.693146777052... | 27. 0.693147180302... |
| 8. 0.692750186011... | 18. 0.693146988980... | 28. 0.693147180435... |
| 9. 0.692967199900... | 19. 0.693147089367... | 29. 0.693147180499... |
| 10. 0.693064856150... | 20. 0.693147137051... | 30. 0.693147180530... |

```

\def\coeffarctg #1{1/\the\numexpr\ifodd #1 -2*#1-1\else2*#1+1\fi\relax }%
% the above gives (-1)^n/(2n+1). The sign being in the denominator,
%      **** no [0] should be added ****,
% else nothing is guaranteed to work (even if it could by sheer luck)
% Notice in passing this aspect of \numexpr:
%      **** \numexpr -(1)\relax is ilegal !!! ****
\def\frac {1/25[0]}% 1/5^2
\[\mathrm{Arctg}(\frac{1}{5})\approx \frac{1}{5}\sum_{n=0}^{15} \frac{(-1)^n}{(2n+1)25^n}
= \xintTeXFrac{\xintIrr {\xintDiv {\xintPowerSeries {0}{15}{\coeffarctg}{\f}}{5}}\]

```

$$\text{Arctg}\left(\frac{1}{5}\right) \approx \frac{1}{5} \sum_{n=0}^{15} \frac{(-1)^n}{(2n+1)25^n} = \frac{165918726519122955895391793269168}{840539304153062403202056884765625}$$

15.6. \xintPowerSeriesX

$\frac{\text{num}}{\text{x}}$ $\frac{\text{num}}{\text{x}}$ $\frac{\text{Frac}}{\text{f}}$ $\frac{\text{Frac}}{\text{f}}$

This is the same as `\xintPowerSeries` apart from the fact that the last parameter `f` is expanded once and for all before being then used repeatedly. If the `f` parameter is to be an explicit big fraction with many (dozens) digits, rather than using it directly it is slightly better to have some macro `\g` defined to expand to the explicit fraction and then use `\xintPowerSeries` with `\g`; but if `f` has not yet been evaluated and will be the output of a complicated expansion of some `\f`, and if, due to an expanding only context, doing `\edef\g{\f}` is no option, then `\xintPowerSeriesX` should be used with `\f` as last parameter.

```

\def\ratioexp #1#2{\xintDiv {#1}{#2}}% x/n
% These are the (-1)^{n-1}/n of the log(1+h) series:
\def\coefflog #1{\the\numexpr\ifodd #1 1\else-1\fi\relax/#1[0]}%
% Let L(h) be the first 10 terms of the log(1+h) series and
% let E(t) be the first 10 terms of the exp(t) series.
% The following computes L(E(a/10)-1) for a=1,..., 12.
\begin{multicols}{3}\raggedcolumns
\cnta 1
\loop
\noindent\xintTrunc {18}{%
  \xintPowerSeriesX {1}{10}{\coefflog}
  {\xintSub
    {\xintRationalSeries {0}{9}{1[0]}\ratioexp{\the\cnta[-1]}}
    {1}}\dots
\endgraf
\ifnum\cnta < 12 \advance \cnta 1 \repeat
\end{multicols}

```

0.099999999998556159...	0.499511320760604148...	-1.597091692317639401...
0.199999995263443554...	0.593980619762352217...	-12.648937932093322763...
0.299999338075041781...	0.645144282733914916...	-66.259639046914679687...
0.399974460740121112...	0.398118280111436442...	-304.768437445462801227...

15.7. `\xintFxpPowerSeries`

$\frac{\text{num}}{\text{Frac}} \frac{\text{num}}{\text{X}} \frac{\text{num}}{\text{X}} \star$
 $\frac{\text{num}}{\text{f}} \frac{\text{num}}{\text{f}} \frac{\text{num}}{\text{X}} \star$

`\xintFxpPowerSeries{A}{B}{\coeff}{f}{D}` computes $\sum_{n=A}^{n=B} \text{coeff}\{n\} \cdot f^n$ with each term of the series truncated to *D* digits after the decimal point. As usual, *A* and *B* are completely expanded through their inclusion in a `\numexpr` expression. Regarding *D* it will be similarly be expanded each time it is used inside an `\xintTrunc`. The one-parameter macro `\coeff` is similarly expanded at the time it is used inside the computations. Idem for *f*. If *f* itself is some complicated macro it is thus better to use the variant `\xintFxpPowerSeriesX` which expands it first and then uses the result of that expansion.

The current (1.04) implementation is: the first power f^A is computed exactly, then *truncated*. Then each successive power is obtained from the previous one by multiplication by the exact value of *f*, and truncated. And $\text{coeff}\{n\} \cdot f^n$ is obtained from that by multiplying by `\coeff{n}` (untruncated) and then truncating. Finally the sum is computed exactly. Apart from that `\xintFxpPowerSeries` (where *FxpPt* means 'fixed-point') is like `\xintPowerSeries`.

There should be a variant for things of the type $\sum c_n \frac{f^n}{n!}$ to avoid having to compute the factorial from scratch at each coefficient, the same way `\xintFxpPowerSeries` does not compute f^n from scratch at each *n*. Perhaps in the next package release.

$$e^{-\frac{1}{2}} \approx$$

1.00000000000000000000	0.60653056795634920635	0.60653065971263344622
0.50000000000000000000	0.60653066483754960317	0.60653065971263342289
0.62500000000000000000	0.60653065945526069224	0.60653065971263342361
0.60416666666666666667	0.60653065972437513778	0.60653065971263342359
0.60677083333333333333	0.60653065971214266299	0.60653065971263342359
0.60651041666666666667	0.60653065971265234943	0.60653065971263342359
0.60653211805555555555	0.60653065971263274611	

```

\def\coeffexp #1{1/\xintiiFac {#1}[0]}% 1/n!
\def\frac{-1/2[0]}% [0] for faster input parsing
\cnta 0 % previously declared \count register
\noindent\loop
$\xintFxpPowerSeries {0}{\cnta}{\coeffexp}{\frac}{20}$\
\ifnum\cnta<19 \advance\cnta 1 \repeat\par
\xintFxpPowerSeries {0}{19}{\coeffexp}{\frac}{25}= 0.6065306597126334236037992
    
```

It is no difficulty for `xintfrac` to compute exactly, with the help of `\xintPowerSeries`, the nineteenth partial sum, and to then give (the start of) its exact decimal expansion:

$$\xintPowerSeries {0}{19}{\coeffexp}{\frac} = \frac{38682746160036397317757}{63777066403145711616000}$$

$$= 0.606530659712633423603799152126\dots$$

Thus, one should always estimate a priori how many ending digits are not reliable: if there are *N* terms and *N* has *k* digits, then digits up to but excluding the last *k* may usually be trusted. If we are optimistic and the series is alternating we may even replace *N* with \sqrt{N} to get the number *k* of digits possibly of dubious significance.

15.8. `\xintFxpPowerSeriesX`

$\frac{\text{num}}{\text{Frac}} \frac{\text{num}}{\text{X}} \frac{\text{num}}{\text{X}} \star$
 $\frac{\text{num}}{\text{f}} \frac{\text{num}}{\text{f}} \frac{\text{num}}{\text{X}} \star$

`\xintFxpPowerSeriesX{A}{B}{\coeff}{f}{D}` computes, exactly as `\xintFxpPowerSeries`, the sum of $\text{coeff}\{n\} \cdot f^n$ from *n=A* to *n=B* with each term of the series being *truncated* to *D* digits after

the decimal point. The sole difference is that `\f` is first expanded and it is the result of this which is used in the computations.

Let us illustrate this on the numerical exploration of the identity

$$\log(1+x) = -\log(1/(1+x))$$

Let $L(h)=\log(1+h)$, and $D(h)=L(h)+L(-h/(1+h))$. Theoretically thus, $D(h)=0$ but we shall evaluate $L(h)$ and $-h/(1+h)$ keeping only 10 terms of their respective series. We will assume $h < 0.5$. With only ten terms kept in the power series we do not have quite 3 digits precision as $2^{10} = 1024$. So it wouldn't make sense to evaluate things more precisely than, say circa 5 digits after the decimal points.

```
\cnta 0
\def\coefflog #1{\the\numexpr\ifodd#1 1\else-1\fi\relax/#1[0]}% (-1)^{n-1}/n
\def\coeffalt #1{\the\numexpr\ifodd#1 -1\else1\fi\relax [0]}% (-1)^n
\begin{multicols}2
\loop
\noindent \hbox to 2.5cm {\hss\texttt{D(\the\cnta/100): }}%
\xintAdd {\xintFxFtPowerSeriesX {1}{10}{\coefflog}{\the\cnta [-2]}{5}}
{\xintFxFtPowerSeriesX {1}{10}{\coefflog}
{\xintFxFtPowerSeriesX {1}{10}{\coeffalt}{\the\cnta [-2]}{5}}
{5}}\endgraf
\ifnum\cnta < 49 \advance\cnta 7 \repeat
\end{multicols}
```

D(0/100): 0/1[0]	D(28/100): 4/1[-5]
D(7/100): 2/1[-5]	D(35/100): 4/1[-5]
D(14/100): 2/1[-5]	D(42/100): 9/1[-5]
D(21/100): 3/1[-5]	D(49/100): 42/1[-5]

Let's say we evaluate functions on $[-1/2, +1/2]$ with values more or less also in $[-1/2, +1/2]$ and we want to keep 4 digits of precision. So, roughly we need at least 14 terms in series like the geometric or log series. Let's make this 15. Then it doesn't make sense to compute intermediate summands with more than 6 digits precision. So we compute with 6 digits precision but return only 4 digits (rounded) after the decimal point. This result with 4 post-decimal points precision is then used as input to the next evaluation.

```
\begin{multicols}2
\loop
\noindent \hbox to 2.5cm {\hss\texttt{D(\the\cnta/100): }}%
\dt{\xintRound{4}
{\xintAdd {\xintFxFtPowerSeriesX {1}{15}{\coefflog}{\the\cnta [-2]}{6}}
{\xintFxFtPowerSeriesX {1}{15}{\coefflog}
{\xintRound {4}{\xintFxFtPowerSeriesX {1}{15}{\coeffalt}
{\the\cnta [-2]}{6}}}
{6}}%
}}\endgraf
\ifnum\cnta < 49 \advance\cnta 7 \repeat
\end{multicols}
```

D(0/100): 0	D(28/100): -0.0001
D(7/100): 0.0000	D(35/100): -0.0001
D(14/100): 0.0000	D(42/100): -0.0000
D(21/100): -0.0001	D(49/100): -0.0001

Not bad... I have cheated a bit: the 'four-digits precise' numeric evaluations were left unrounded in the final addition. However the inner rounding to four digits worked fine and made the

next step faster than it would have been with longer inputs. The morale is that one should not use the raw results of `\xintFxFtPowerSeriesX` with the `D` digits with which it was computed, as the last are to be considered garbage. Rather, one should keep from the output only some smaller number of digits. This will make further computations faster and not less precise. I guess there should be some macro to do this final truncating, or better, rounding, at a given number `D' < D` of digits. Maybe for the next release.

15.9. `\xintFloatPowerSeries`

`\xintFloatPowerSeries[P]{A}{B}{coeff}{f}` computes $\sum_{n=A}^{n=B} \text{coeff}\{n\} \cdot f^n$ with a floating point precision given by the optional parameter `P` or by the current setting of `\xintDigits`.

In the current, preliminary, version, no attempt has been made to try to guarantee to the final result the precision `P`. Rather, `P` is used for all intermediate floating point evaluations. So rounding errors will make some of the last printed digits invalid. The operations done are first the evaluation of f^A using `\xintFloatPow`, then each successive power is obtained from this first one by multiplication by `f` using `\xintFloatMul`, then again with `\xintFloatMul` this is multiplied with `\coeff{n}`, and the sum is done adding one term at a time with `\xintFloatAdd`. To sum up, this is just the naive transformation of `\xintFxFtPowerSeries` from fixed point to floating point.

```
\def\coefflog #1{\the\numexpr\ifodd#1 1\else-1\fi\relax/#1[0]}%
\xintFloatPowerSeries [8]{1}{30}{\coefflog}{-1/2[0]}
-6.9314718e-1
```

15.10. `\xintFloatPowerSeriesX`

`\xintFloatPowerSeriesX[P]{A}{B}{coeff}{f}` is like `\xintFloatPowerSeries` with the difference that `f` is expanded once and for all at the start of the computation, thus allowing efficient chaining of such series evaluations.

```
\def\coeffexp #1{1/\xintiiFac {#1}[0]}% 1/n! (exact, not float)
\def\coefflog #1{\the\numexpr\ifodd#1 1\else-1\fi\relax/#1[0]}%
\xintFloatPowerSeriesX [8]{0}{30}{\coeffexp}
{\xintFloatPowerSeries [8]{1}{30}{\coefflog}{-1/2[0]}}
5.0000001e-1
```

15.11. Computing $\log 2$ and π

In this final section, the use of `\xintFxFtPowerSeries` (and `\xintPowerSeries`) will be illustrated on the (expandable... why make things simple when it is so easy to make them difficult!) computations of the first digits of the decimal expansion of the familiar constants $\log 2$ and π .

Let us start with $\log 2$. We will get it from this formula (which is left as an exercise):

$$\log(2) = -2 \log(1 - 13/256) - 5 \log(1 - 1/9)$$

The number of terms to be kept in the log series, for a desired precision of 10^{-D} was roughly estimated without much theoretical analysis. Computing exactly the partial sums with `\xintPowerSeries` and then printing the truncated values, from `D=0` up to `D=100` showed that it worked in terms of quality of the approximation. Because of possible strings of zeroes or nines in the exact decimal expansion (in the present case of $\log 2$, strings of zeroes around the fortieth and the sixtieth decimals), this does not mean though that all digits printed were always exact. In the end one always end up having to compute at some higher level of desired precision to validate the earlier result.

Then we tried with `\xintFxFtPowerSeries`: this is worthwhile only for `D`'s at least 50, as the exact evaluations are faster (with these short-length `f`'s) for a lower number of digits. And as expected the degradation in the quality of approximation was in this range of the order of two or three digits. This meant roughly that the 3+1=4 ending digits were wrong. Again, we ended up having to compute with five more digits and compare with the earlier value to validate it. We use truncation

rather than rounding because our goal is not to obtain the correct rounded decimal expansion but the correct exact truncated one.

```

\def\coefflog #1{1/#1[0]}% 1/n
\def\xa {13/256[0]}% we will compute log(1-13/256)
\def\xb {1/9[0]}% we will compute log(1-1/9)
\def\LogTwo #1%
% get log(2)=-2log(1-13/256)- 5log(1-1/9)
{% we want to use \printnumber, hence need something expanding in two steps
% only, so we use here the \romannumeral0 method
\romannumeral0\expandafter\LogTwoDoIt \expandafter
% Nb Terms for 1/9:
{\the\numexpr #1*150/143\expandafter}\expandafter
% Nb Terms for 13/256:
{\the\numexpr #1*100/129\expandafter}\expandafter
% We print #1 digits, but we know the ending ones are garbage
{\the\numexpr #1\relax}% allows #1 to be a count register
}%
\def\LogTwoDoIt #1#2#3%
% #1=nb of terms for 1/9, #2=nb of terms for 13/256,
{% #3=nb of digits for computations, also used for printing
\xinttrunc {#3} % lowercase form to stop the \romannumeral0 expansion!
{\xintAdd
{\xintMul {2}{\xintFxFtPowerSeries {1}{#2}{\coefflog}{\xa}{#3}}}
{\xintMul {5}{\xintFxFtPowerSeries {1}{#1}{\coefflog}{\xb}{#3}}}%
}%
}%
\noindent $\log 2 \approx \LogTwo {60}\dots$\endgraf
\noindent\phantom{$\log 2$}$\approx{\}$\printnumber{\LogTwo {65}}\dots\endgraf
\noindent\phantom{$\log 2$}$\approx{\}$\printnumber{\LogTwo {70}}\dots\endgraf

```

```

log 2 ≈ 0.693147180559945309417232121458176568075500134360255254120484...
≈ 0.69314718055994530941723212145817656807550013436025525412068000711...
≈ 0.6931471805599453094172321214581765680755001343602552541206800094933723...

```

Here is the code doing an exact evaluation of the partial sums. We have added a +1 to the number of digits for estimating the number of terms to keep from the log series: we experimented that this gets exactly the first D digits, for all values from $D=0$ to $D=100$, except in one case ($D=40$) where the last digit is wrong. For values of D higher than 100 it is more efficient to use the code using `\xintFxFtPowerSeries`.

```

\def\LogTwo #1% get log(2)=-2log(1-13/256)- 5log(1-1/9)
{%
\romannumeral0\expandafter\LogTwoDoIt \expandafter
{\the\numexpr (#1+1)*150/143\expandafter}\expandafter
{\the\numexpr (#1+1)*100/129\expandafter}\expandafter
{\the\numexpr #1\relax}%
}%
\def\LogTwoDoIt #1#2#3%
{% #3=nb of digits for truncating an EXACT partial sum
\xinttrunc {#3}
{\xintAdd
{\xintMul {2}{\xintPowerSeries {1}{#2}{\coefflog}{\xa}}}
{\xintMul {5}{\xintPowerSeries {1}{#1}{\coefflog}{\xb}}}%
}%
}%

```

Let us turn now to Pi, computed with the Machin formula (but see also the approach via the [Brent-Salamin algorithm](#) with `\xintfloatexpr`) Again the numbers of terms to keep in the two `arctg` series were roughly estimated, and some experimentations showed that removing the last three digits was enough (at least for `D=0-100` range). And the algorithm does print the correct digits when used with `D=1000` (to be convinced of that one needs to run it for `D=1000` and again, say for `D=1010`.) A theoretical analysis could help confirm that this algorithm always gets better than 10^{-D} precision, but again, strings of zeroes or nines encountered in the decimal expansion may falsify the ending digits, nines may be zeroes (and the last non-nine one should be increased) and zeroes may be nine (and the last non-zero one should be decreased).

```
\def\coeffarctg #1{\the\numexpr\ifodd#1 -1\else1\fi\relax/%
\the\numexpr 2*#1+1\relax [0]}%
%\def\coeffarctg #1{\romannumeral0\xintmon{#1}/\the\numexpr 2*#1+1\relax }%
\def\xa {1/25[0]}%      1/5^2, the [0] for faster parsing
\def\xb {1/57121[0]}%  1/239^2, the [0] for faster parsing
\def\Machin #1{% #1 may be a count register, \Machin {\mycount} is allowed
\romannumeral0\expandafter\MachinA \expandafter
% number of terms for arctg(1/5):
{\the\numexpr (#1+3)*5/7\expandafter}\expandafter
% number of terms for arctg(1/239):
{\the\numexpr (#1+3)*10/45\expandafter}\expandafter
% do the computations with 3 additional digits:
{\the\numexpr #1+3\expandafter}\expandafter
% allow #1 to be a count register:
{\the\numexpr #1\relax }}%
\def\MachinA #1#2#3#4%
{\xinttrunc {#4}
{\xintSub
{\xintMul {16/5}{\xintFxpPowerSeries {0}{#1}{\coeffarctg}{\xa}{#3}}}
{\xintMul{4/239}{\xintFxpPowerSeries {0}{#2}{\coeffarctg}{\xb}{#3}}}%
}}%
\begin{framed}
\[\ \pi = \Machin {60}\dots \]
\end{framed}
```

$$\pi = 3.141592653589793238462643383279502884197169399375105820974944\dots$$

Here is a variant `\MachinBis`, which evaluates the partial sums *exactly* using `\xintPowerSeries`, before their final truncation. No need for a ```+3''` then.

```
\def\MachinBis #1{% #1 may be a count register,
% the final result will be truncated to #1 digits post decimal point
\romannumeral0\expandafter\MachinBisA \expandafter
% number of terms for arctg(1/5):
{\the\numexpr #1*5/7\expandafter}\expandafter
% number of terms for arctg(1/239):
{\the\numexpr #1*10/45\expandafter}\expandafter
% allow #1 to be a count register:
{\the\numexpr #1\relax }}%
\def\MachinBisA #1#2#3%
{\xinttrunc {#3} %
{\xintSub
{\xintMul {16/5}{\xintPowerSeries {0}{#1}{\coeffarctg}{\xa}}}
```

```
{\xintMul{4/239}{\xintPowerSeries {0}{#2}{\coeffarctg}{\xb}}}%
}}%
```

Let us use this variant for a loop showing the build-up of digits:

```
\begin{multicols}{2}
  \cnta 0 % previously declared \count register
  \loop \noindent
    \centeredline{\dtt{\MachinBis{\cnta}}}%
  \ifnum\cnta < 30
    \advance\cnta 1 \repeat
\end{multicols}
```

	3.141592653589793
3.	3.1415926535897932
3.1	3.14159265358979323
3.14	3.141592653589793238
3.141	3.1415926535897932384
3.1415	3.14159265358979323846
3.14159	3.141592653589793238462
3.141592	3.1415926535897932384626
3.1415926	3.14159265358979323846264
3.14159265	3.141592653589793238462643
3.141592653	3.1415926535897932384626433
3.1415926535	3.14159265358979323846264338
3.14159265358	3.141592653589793238462643383
3.141592653589	3.1415926535897932384626433832
3.1415926535897	3.14159265358979323846264338327
3.14159265358979	3.141592653589793238462643383279

You want more digits and have some time? compile this copy of the `\Machin` with `etex` (or `pdftex`):

```
% Compile with e-TeX extensions enabled (etex, pdftex, ...)
\input xintfrac.sty
\input xintseries.sty
% pi = 16 Arctg(1/5) - 4 Arctg(1/239) (John Machin's formula)
\def\coeffarctg #1{\the\numexpr\ifodd#1 -1\else1\fi\relax/%
                    \the\numexpr 2*#1+1\relax [0]}%
\def\xa {1/25[0]}%
\def\xb {1/57121[0]}%
\def\Machin #1{%
  \romannumeral0\expandafter\MachinA \expandafter
  {\the\numexpr (#1+3)*5/7\expandafter}\expandafter
  {\the\numexpr (#1+3)*10/45\expandafter}\expandafter
  {\the\numexpr #1+3\expandafter}\expandafter
  {\the\numexpr #1\relax }}%
\def\MachinA #1#2#3#4%
{\xinttrunc {#4}
 {\xintSub
  {\xintMul {16/5}{\xintFxFtPowerSeries {0}{#1}{\coeffarctg}{\xa}{#3}}}
  {\xintMul {4/239}{\xintFxFtPowerSeries {0}{#2}{\coeffarctg}{\xb}{#3}}}%
 }}%
\pdfresettimer
\fdef\Z {\Machin {1000}}
\odef\W {\the\pdfelapsedtime}
```

```
\message{\Z}
\message{computed in \xintRound {2}{\W/65536} seconds.}
\bye
```

This will log the first 1000 digits of π after the decimal point. On my laptop (a 2012 model) this took about 5.05 seconds last time I tried.^{74 75}

As mentioned in the introduction, the file `pi.tex` by D. ROEGEL shows that orders of magnitude faster computations are possible within \TeX , but recall our constraints of complete expandability and be merciful, please.

Why truncating rather than rounding? One of our main competitors on the market of scientific computing, a canadian product (not encumbered with expandability constraints, and having barely ever heard of \TeX ;-), prints numbers rounded in the last digit. Why didn't we follow suit in the macros `\xintFxpPowerSeries` and `\xintFxpPowerSeriesX`? To round at D digits, and excluding a rewrite or cloning of the division algorithm which anyhow would add to it some overhead in its final steps, `xintfrac` needs to truncate at $D+1$, then round. And rounding loses information! So, with more time spent, we obtain a worst result than the one truncated at $D+1$ (one could imagine that additions and so on, done with only D digits, cost less; true, but this is a negligible effect per summand compared to the additional cost for this term of having been truncated at $D+1$ then rounded). Rounding is the way to go when setting up algorithms to evaluate functions destined to be composed one after the other: exact algebraic operations with many summands and an f variable which is a fraction are costly and create an even bigger fraction; replacing f with a reasonable rounding, and rounding the result, is necessary to allow arbitrary chaining.

But, for the computation of a single constant, we are really interested in the exact decimal expansion, so we truncate and compute more terms until the earlier result gets validated. Finally if we do want the rounding we can always do it on a value computed with $D+1$ truncation.

⁷⁴ With 1.09i and earlier `xint`, this used to be 42 seconds; starting with 1.09j, and prior to 1.2, it was 16 seconds (this was probably due to a more efficient division with denominators at most 9999). The 1.2 `xintcore` achieves a further gain at 5.6 seconds.

⁷⁵ With `\xintDigits :=1001\relax`, the non-optimized implementation with the `iter` of `xintexpr` fame using the `Brent-Salamin algorithm`, took, last time I tried (1.2i), about 7 seconds on my laptop (the last two digits were wrong, which is ok as they serve as guard digits), and for obtaining about 500 digits, it was about 1.7s. This is not bad, taking into account that the syntax is almost free rolling speech, contrarily to the code above for the Machin formula computation; we would like to use the quadratically convergent Brent-Salamin algorithm for more digits, but with such computations with numbers of one thousand digits we are beyond the border of the reasonable range for `xint`. Innocent people not knowing what it means to compute with \TeX , and with the extra constraint of expandability will wonder why this is at least thousands of times slower than with any other language (with a little Python program using the `Decimal` library, I timed the Brent-Salamin algorithm to 4.4ms for about 1000 digits and 1.14ms for 500 digits.) I will just say that for example digits are represented and manipulated via their ascii-code ! all computations must convert from ascii-code to cpu words; furthermore nothing can be stored away. And there is no memory storage with $O(1)$ time access... if expandability is to be verified.

16. Macros of the *xintcfrac* package

First version of this package was included in release 1.04 (2013/04/25) of the *xint* bundle. It was kept almost unchanged until 1.09m of 2014/02/26 which brought some new macros: `\xintFtoC`, `\xintCtoF`, `\xintCtoCv`, dealing with sequences of braced partial quotients rather than comma separated ones, `\xintFGtoC` which is to produce ``guaranteed'' coefficients of some real number known approximately, and `\xintGGCFrac` for displaying arbitrary material as a continued fraction; also, some changes to existing macros: `\xintFtoCs` and `\xintCntoCs` insert spaces after the commas, `\xintCstoF` and `\xintCstoCv` authorize spaces in the input also before the commas.

Note: `\xintCstoF` and `\xintCstoCv` create a partial dependency on *xinttools* (its `\xintCSVtoList`.)

<ul style="list-style-type: none"> .1 Package overview 214 .2 <code>\xintCFrac</code>..... 219 .3 <code>\xintGCFrac</code>..... 219 .4 <code>\xintGGCFrac</code> 219 .5 <code>\xintGctoGCx</code> 220 .6 <code>\xintFtoC</code>..... 220 .7 <code>\xintFtoCs</code>..... 220 .8 <code>\xintFtoCx</code>..... 220 .9 <code>\xintFtoGC</code>..... 221 .10 <code>\xintFGtoC</code>..... 221 .11 <code>\xintFtoCC</code>..... 221 .12 <code>\xintCstoF</code>..... 222 .13 <code>\xintCtoF</code> 222 .14 <code>\xintGctoF</code>..... 223 .15 <code>\xintCstoCv</code>..... 223 	<ul style="list-style-type: none"> .16 <code>\xintCtoCv</code>..... 224 .17 <code>\xintGctoCv</code>..... 224 .18 <code>\xintFtoCv</code>..... 224 .19 <code>\xintFtoCCv</code>..... 224 .20 <code>\xintCntoF</code>..... 224 .21 <code>\xintGcntoF</code>..... 225 .22 <code>\xintCntoCs</code>..... 225 .23 <code>\xintCntoGC</code>..... 225 .24 <code>\xintGcntoGC</code> 226 .25 <code>\xintCstoGC</code>..... 226 .26 <code>\xintiCstoF</code>, <code>\xintiGctoF</code>, <code>\xintiCstoCv</code>, <code>\xintiGctoCv</code> 226 .27 <code>\xintGctoGC</code>..... 226 .28 Euler's number e 227
--	--

16.1. Package overview

The package computes partial quotients and convergents of a fraction, or conversely start from coefficients and obtain the corresponding fraction; three macros `\xintCFrac`, `\xintGCFrac` and `\xintGGCFrac` are for typesetting, the others can be nested (if applicable) or see their outputs further processed by other macros from the *xint* bundle, particularly the macros of *xinttools* dealing with sequences of braced items or comma separated lists.

A *simple* continued fraction has coefficients $[c_0, c_1, \dots, c_N]$ (usually called partial quotients, but I dislike this entrenched terminology), where c_0 is a positive or negative integer and the others are positive integers.

Typesetting is usually done via the *amsmath* macro `\cfrac`:

```
\[ c_0 + \cfrac{1}{c_1 + \cfrac{1}{c_2 + \cfrac{1}{c_3 + \cfrac{1}{\ddots}}}}\]
```

$$c_0 + \frac{1}{c_1 + \frac{1}{c_2 + \frac{1}{c_3 + \frac{1}{\ddots}}}}$$

Here is a concrete example:

```
\[ \xintTeXfrac {208341/66317}=\xintCFrac {208341/66317}\]
```

$$\frac{208341}{66317} = 3 + \frac{1}{7 + \frac{1}{15 + \frac{1}{1 + \frac{1}{292 + \frac{1}{2}}}}}$$

But it is the macro `\xintCFrac` which did all the work of *computing* the continued fraction and using `\cfrac` from *amsmath* to typeset it.

A *generalized* continued fraction has the same structure but the numerators are not restricted to be 1, and numbers used in the continued fraction may be arbitrary, also fractions, irrationals, complex, indeterminates.⁷⁶ The *centered* continued fraction is an example:

```
\[ \xintTeXFrac {915286/188421}=\xintGCFrac {5+-1/7+1/39+-1/53+-1/13}
=\xintCFrac {915286/188421}\]
```

$$\frac{915286}{188421} = 5 - \frac{1}{7 + \frac{1}{39 - \frac{1}{53 - \frac{1}{13}}}} = 4 + \frac{1}{1 + \frac{1}{6 + \frac{1}{38 + \frac{1}{1 + \frac{1}{51 + \frac{1}{1 + \frac{1}{12}}}}}}}$$

The macro `\xintGCFrac`, contrarily to `\xintCFrac`, does not compute anything, it just typesets starting from a generalized continued fraction in inline format, which in this example was input literally. We also used `\xintCFrac` for comparison of the two types of continued fractions.

To let \TeX compute the centered continued fraction of *f* there is `\xintFtoCC`:

```
\[\xintTeXFrac {915286/188421}\to\xintFtoCC {915286/188421}\]
```

$$\frac{915286}{188421} \rightarrow 5 + -1/7 + 1/39 + -1/53 + -1/13$$

The package macros are expandable and may be nested (naturally `\xintCFrac` and `\xintGCFrac` must be at the top level, as they deal with typesetting).

```
\[\xintGCFrac {\xintFtoCC{915286/188421}}\]
```

$$5 - \frac{1}{7 + \frac{1}{39 - \frac{1}{53 - \frac{1}{13}}}}$$

The `'inline'` format expected on input by `\xintGCFrac` is

$$a_0 + b_0/a_1 + b_1/a_2 + b_2/a_3 + \dots + b_{n-2}/a_{n-1} + b_{n-1}/a_n$$

Fractions among the coefficients are allowed but they must be enclosed within braces. Signed integers may be left without braces (but the `+` signs are mandatory). No spaces are allowed around the plus and fraction symbols. The coefficients may themselves be macros, as long as these macros are *f-expandable*.

```
\[ \xintTeXFrac{\xintGctoF {1+-1/57+\xintPow {-3}{7}}/\xintiiQuo {132}{25}}
= \xintGCFrac {1+-1/57+\xintPow {-3}{7}}/\xintiiQuo {132}{25}\]
```

⁷⁶ `xintcfrac` may be used with indeterminates, for basic conversions from one inline format to another, but not for actual computations. See `\xintGGCFrac`.

$$\frac{1907}{1902} = 1 - \frac{1}{57 - \frac{2187}{5}}$$

To compute the actual fraction one has `\xintGCtoF`:

```
\[\xintTeXFrac{\xintGCtoF {1+-1/57+\xintPow {-3}{7}/\xintiiQuo {132}{25}}}\]
```

$$\frac{1907}{1902}$$

For non-numeric input there is `\xintGGCFrac`.

```
\[\xintGGCFrac {a_0+b_0/a_1+b_1/a_2+b_2/\ddots+\ddots/a_{n-1}+b_{n-1}/a_n}\]
```

$$a_0 + \frac{b_0}{a_1 + \frac{b_1}{a_2 + \frac{b_2}{\ddots + \frac{b_{n-1}}{a_{n-1} + \frac{b_n}{a_n}}}}}$$

For regular continued fractions, there is a simpler comma separated format:

```
\[-7,6,19,1,33\to\xintTeXFrac{\xintCstoF{-7,6,19,1,33}}=\xintCFrac{\xintCstoF{-7,6,19,1,33}}\]
```

$$-7, 6, 19, 1, 33 \rightarrow \frac{-28077}{4108} = -7 + \frac{1}{6 + \frac{1}{19 + \frac{1}{1 + \frac{1}{33}}}}$$

The macro `\xintFtoCs` produces from a fraction *f* the comma separated list of its coefficients.

```
\[\xintTeXFrac{1084483/398959}=[\xintFtoCs{1084483/398959}]\]
```

$$\frac{1084483}{398959} = [2, 1, 2, 1, 1, 4, 1, 1, 6, 1, 1, 8, 1, 1, 10, 2]$$

If one prefers other separators, one can use the two arguments macros `\xintFtoCx` whose first argument is the separator (which may consist of more than one token) which is to be used.

```
\[\xintTeXFrac{2721/1001}=\xintFtoCx {+1/(\}{2721/1001)}\cdots)\]
```

$$\frac{2721}{1001} = 2 + 1/(1 + 1/(2 + 1/(1 + 1/(1 + 1/(4 + 1/(1 + 1/(1 + 1/(6 + 1/(2) \cdots))))))))$$

This allows under Plain TeX with `amstex` to obtain the same effect as with `LaTeX+amsmath+\xintCFrac`:

```
$$\xintTeXOver{2721/1001}=\xintFtoCx {+\cfrac1{\ }}{2721/1001}\endcfrac$$
```

As a shortcut to `\xintFtoCx` with separator `1+/,` there is `\xintFtoGC`:

```
2721/1001=\xintFtoGC {2721/1001}
```

`2721/1001=2+1/1+1/2+1/1+1/1+1/4+1/1+1/1+1/6+1/2` Let us compare in that case with the output of `\xintFtoCC`:

```
2721/1001=\xintFtoCC {2721/1001}
```

`2721/1001=3+-1/4+-1/2+1/5+-1/2+1/7+-1/2` To obtain the coefficients as a sequence of braced numbers, there is `\xintFtoC` (this is a shortcut for `\xintFtoCx {}`). This list (sequence) may then be manipulated using the various macros of `xinttools` such as the non-expandable macro `\xintAssignArray` or the expandable `\xintApply` and `\xintListWithSep`.

Conversely to go from such a sequence of braced coefficients to the corresponding fraction there is `\xintCtoF`.

The `\printnumber` (subsection 1.6) macro which we use in this document to print long numbers can also be useful on long continued fractions.

```
\printnumber{\xintFtoCC {35037018906350720204351049/244241737886197404558180}}
```


$143+1/2+1/5+-1/4+-1/4+-1/4+-1/3+1/2+1/2+1/6+-1/22+1/2+1/10+-1/5+-1/11+-1/3+1/4+-1/2+1/2+1/4+ -1/2+1/23+1/3+1/8+-1/6+-1/9$ If we apply `\xintGctoF` to this generalized continued fraction, we discover that the original fraction was reducible:

```
\xintGctoF {143+1/2+...+-1/9}=2897319801297630107/20197107104701740
```

When a generalized continued fraction is built with integers, and numerators are only 1's or -1's, the produced fraction is irreducible. And if we compute it again with the last sub-fraction omitted we get another irreducible fraction related to the bigger one by a Bézout identity. Doing this here we get:

```
\xintGctoF {143+1/2+...+-1/6}=328124887710626729/2287346221788023
```

and indeed:

$$\begin{vmatrix} 2897319801297630107 & 328124887710626729 \\ 20197107104701740 & 2287346221788023 \end{vmatrix} = 1$$

The various fractions obtained from the truncation of a continued fraction to its initial terms are called the convergents. The macros of `xintcfrac` such as `\xintFtoCv`, `\xintFtoCCv`, and others which compute such convergents, return them as a list of braced items, with no separator (as does `\xintFtoC` for the partial quotients). Here is an example:

```
\[\xintTeXFrac{915286/188421}\to
\xintListWithSep{,}{\xintApply\xintTeXFrac{\xintFtoCv{915286/188421}}\}
\frac{915286}{188421} \rightarrow 4, 5, \frac{34}{7}, \frac{1297}{267}, \frac{1331}{274}, \frac{69178}{14241}, \frac{70509}{14515}, \frac{915286}{188421}
```

```
\[\xintTeXFrac{915286/188421}\to
\xintListWithSep{,}{\xintApply\xintTeXFrac{\xintFtoCCv{915286/188421}}\}
\frac{915286}{188421} \rightarrow 5, \frac{34}{7}, \frac{1331}{274}, \frac{70509}{14515}, \frac{915286}{188421}
```

We thus see that the 'centered convergents' obtained with `\xintFtoCCv` are among the fuller list of convergents as returned by `\xintFtoCv`.

Here is a more complicated use of `\xintApply` and `\xintListWithSep`. We first define a macro which will be applied to each convergent:

```
\newcommand{\mymacro}[1]{\xintTeXFrac{#1}=[\xintFtoCs{#1}]\$ \vtop to 6pt{}}
```

Next, we use the following code:

```
 $\xintTeXFrac{49171/18089}\to{$
 \xintListWithSep {, }{\xintApply{\mymacro}{\xintFtoCv{49171/18089}}}
```

It produces:

$\frac{49171}{18089} \rightarrow 2 = [2], 3 = [3], \frac{8}{3} = [2, 1, 2], \frac{11}{4} = [2, 1, 3], \frac{19}{7} = [2, 1, 2, 2], \frac{87}{32} = [2, 1, 2, 1, 1, 4], \frac{106}{39} = [2, 1, 2, 1, 1, 5], \frac{193}{71} = [2, 1, 2, 1, 1, 4, 2], \frac{1264}{465} = [2, 1, 2, 1, 1, 4, 1, 1, 6], \frac{1457}{536} = [2, 1, 2, 1, 1, 4, 1, 1, 7], \frac{2721}{1001} = [2, 1, 2, 1, 1, 4, 1, 1, 6, 2], \frac{23225}{8544} = [2, 1, 2, 1, 1, 4, 1, 1, 6, 1, 1, 8], \frac{49171}{18089} = [2, 1, 2, 1, 1, 4, 1, 1, 6, 1, 1, 8, 2]$

The macro `\xintCntoF` allows to specify the coefficients as a function given by a one-parameter macro. The produced values do not have to be integers.

```
\def\cn #1{\xintiiPow {2}{#1}}% 2^n
\[\xintTeXFrac{\xintCntoF {6}{\cn}}=\xintCFrac [1]{\xintCntoF {6}{\cn}}\]
\frac{3541373}{2449193} = 1 + \frac{1}{2 + \frac{1}{4 + \frac{1}{8 + \frac{1}{16 + \frac{1}{32 + \frac{1}{64}}}}}}}
```

Notice the use of the optional argument `[1]` to `\xintCFrac`. Other possibilities are `[r]` and (default) `[c]`.

```
\def\cn #1{\xintPow {2}{-#1}}%
\[\xintTeXFrac{\xintCntoF {6}{\cn}}=\xintGCFrac [r]{\xintCntoGC {6}{\cn}}=
[\xintFtoCs {\xintCntoF {6}{\cn}}\]
```

$$\frac{3159019}{2465449} = 1 + \frac{1}{\frac{1}{2} + \frac{1}{\frac{1}{4} + \frac{1}{\frac{1}{8} + \frac{1}{\frac{1}{16} + \frac{1}{\frac{1}{32} + \frac{1}{\frac{1}{64}}}}}}}} = [1, 3, 1, 1, 4, 14, 1, 1, 1, 1, 79, 2, 1, 1, 2]$$

We used `\xintCntoGC` as we wanted to display also the continued fraction and not only the fraction returned by `\xintCntoF`.

There are also `\xintGCntoF` and `\xintGCntoGC` which allow the same for generalized fractions. An initial portion of a generalized continued fraction for π is obtained like this

```
\def\an #1{\the\numexpr 2*#1+1\relax }%
\def\bn #1{\the\numexpr (#1+1)*(#1+1)\relax }%
\[\xintTeXFrac{\xintDiv {4}{\xintGCntoF {5}{\an}{\bn}}} =
\cfrac{4}{\xintGCFrac{\xintGCntoGC {5}{\an}{\bn}}} =
\xintTrunc {10}{\xintDiv {4}{\xintGCntoF {5}{\an}{\bn}}\dots\]
```

$$\frac{92736}{29520} = \frac{4}{1 + \frac{1}{3 + \frac{4}{5 + \frac{9}{7 + \frac{16}{9 + \frac{25}{11}}}}}} = 3.1414634146\dots$$

We see that the quality of approximation is not fantastic compared to the simple continued fraction of π with about as many terms:

```
\[\xintTeXFrac{\xintCstoF{3,7,15,1,292,1,1}}=
\xintGCFrac{3+1/7+1/15+1/1+1/292+1/1+1/1}=
\xintTrunc{10}{\xintCstoF{3,7,15,1,292,1,1}\dots\]
```

$$\frac{208341}{66317} = 3 + \frac{1}{7 + \frac{1}{15 + \frac{1}{1 + \frac{1}{292 + \frac{1}{1 + \frac{1}{1}}}}}} = 3.1415926534\dots$$

When studying the continued fraction of some real number, there is always some doubt about how many terms are valid, when computed starting from some approximation. If $f \leq x \leq g$ and f, g both have the same first K partial quotients, then x also has the same first K quotients and convergents. The macro `\xintFGtoC` outputs as a sequence of braced items the common partial quotients of its two arguments. We can thus use it to produce a sure list of valid convergents of π for example, starting from some proven lower and upper bound:

```
$$\pi\to [\xintListWithSep{,}
```

```
{\xintFGtoC {3.14159265358979323}{3.14159265358979324}}, \dots}$$
\noindent$\pi\to\xintListWithSep{\allowbreak\;}
{\xintApply{\xintTeXFrac}
{\xintCtoCv{\xintFGtoC {3.14159265358979323}{3.14159265358979324}}}}, \dots$
\pi \to [3, 7, 15, 1, 292, 1, 1, 1, 2, 1, 3, 1, 14, 2, 1, 1, ...]
```

$\pi \rightarrow 3, \frac{22}{7}, \frac{333}{106}, \frac{355}{113}, \frac{103993}{33102}, \frac{104348}{33215}, \frac{208341}{66317}, \frac{312689}{99532}, \frac{833719}{265381}, \frac{1146408}{364913}, \frac{4272943}{1360120}, \frac{5419351}{1725033}, \frac{80143857}{25510582}, \frac{165707065}{52746197}, \frac{245850922}{78256779}, \frac{411557987}{131002976}, \dots$

16.2. `\xintCFrac`

Frac
f

`\xintCFrac{f}` is a math-mode only, \TeX with `amsmath` only, macro which first computes then displays with the help of `\cfrac` the simple continued fraction corresponding to the given fraction. It admits an optional argument which may be `[l]`, `[r]` or (the default) `[c]` to specify the location of the one's in the numerators of the sub-fractions. This macro is *f-expandable* in the sense that it prepares expandably the whole expression with the multiple `\cfrac`'s, but it is not completely expandable naturally as `\cfrac` isn't.

16.3. `\xintGCFrac`

f `\xintGCFrac{a+b/c+d/e+f/g+h/...+x/y}` uses similarly `\cfrac` to prepare the typesetting with the `amsmath` `\cfrac` (\TeX) of a generalized continued fraction given in inline format (or as macro which will *f-expand* to it). It admits the same optional argument as `\xintCFrac`. Plain \TeX with `amstex` users, see `\xintGctoGCx`.

```
\[\xintGCFrac {1+\xintPow{1.5}{3}/{1/7}+{-3/5}/\xintiiFac {6}}\]
```

$$1 + \frac{3375 \cdot 10^{-3}}{\frac{1}{7} - \frac{3}{720}}$$

This is mostly a typesetting macro, although it does trigger the expansion of the coefficients. See `\xintGctoF` if you are impatient to see this specific fraction computed.

It admits an optional argument within square brackets which may be either `[l]`, `[c]` or `[r]`. Default is `[c]` (numerators are centered).

Numerators and denominators are made arguments to the `\xintTeXFrac` macro. This allows them to be themselves fractions or anything *f-expandable* giving numbers or fractions, but also means however that they can not be arbitrary material, they can not contain color changing macros for example. One of the reasons is that `\xintGCFrac` tries to determine the signs of the numerators and chooses accordingly to use + or -.

16.4. `\xintGGCFrac`

f `\xintGGCFrac{a+b/c+d/e+f/g+h/...+x/y}` is a clone of `\xintGCFrac`, hence again \TeX specific with package `amsmath`. It does not assume the coefficients to be numbers as understood by `xintfrac`. The macro can be used for displaying arbitrary content as a continued fraction with `\cfrac`, using only plus signs though. Note though that it will first *f-expand* its argument, which may be thus be one of the `xintcfrac` macros producing a (general) continued fraction in inline format, see `\xintFtoCx` for an example. If this expansion is not wished, it is enough to start the argument with a space.

```
\[\xintGGCFrac {1+q/1+q^2/1+q^3/1+q^4/1+q^5/\ddots}\]
```

$$1 + \frac{q}{1 + \frac{q^2}{1 + \frac{q^3}{1 + \frac{q^4}{1 + \frac{q^5}{\ddots}}}}}$$

16.5. `\xintGtoGCx`

`nnf` ★ `\xintGtoGCx{sepa}{sepb}{a+b/c+d/e+f/...+x/y}` returns the list of the coefficients of the generalized continued fraction of `f`, each one within a pair of braces, and separated with the help of `sepa` and `sepb`. Thus

```
\xintGtoGCx ;:{1+2/3+4/5+6/7} gives 1:2;3:4;5:6;7
```

The following can be used by Plain \TeX +`amstex` users to obtain an output similar as the ones produced by `\xintGCFrac` and `\xintGGCFrac`:

```
$$\xintGtoGCx {+\cfrac{\{\}\{a+b/\dots\}}{\}}\endcfrac$$
$$\xintGtoGCx {+\cfrac{\xintTeXOver{\}\{a+b/\dots\}}{\}}\endcfrac$$
```

16.6. `\xintFtoC`

`Frac f` ★ `\xintFtoC{f}` computes the coefficients of the simple continued fraction of `f` and returns them as a list (sequence) of braced items.

```
\fdef\test{\xintFtoC{-5262046/89233}}\texttt{\meaning\test}
```

```
macro:->{-59}{33}{27}{100}
```

16.7. `\xintFtoCs`

`Frac f` ★ `\xintFtoCs{f}` returns the comma separated list of the coefficients of the simple continued fraction of `f`. Notice that starting with 1.09m a space follows each comma (mainly for usage in text mode, as in math mode spaces are produced in the typeset output by \TeX itself).

```
\[ \xintTeXsignedFrac{-5262046/89233} \to [\xintFtoCs{-5262046/89233}]\]
```

$$-\frac{5262046}{89233} \rightarrow [-59, 33, 27, 100]$$

16.8. `\xintFtoCx`

`n f` ★ `\xintFtoCx{sep}{f}` returns the list of the coefficients of the simple continued fraction of `f` separated with the help of `sep`, which may be anything (and is kept unexpanded). For example, with Plain \TeX and `amstex`,

```
$$\xintFtoCx {+\cfrac1{\}}{-5262046/89233}\endcfrac$$
```

will display the continued fraction using `\cfrac`. Each coefficient is inside a brace pair `{ }`, allowing a macro to end the separator and fetch it as argument, for example, again with Plain \TeX and `amstex`:

```
\def\highlight #1{\ifnum #1>200 \textcolor{red}{#1}\else #1\fi}
$$\xintFtoCx {+\cfrac1{\} \highlight}{104348/33215}\endcfrac$$
```

Due to the different and extremely cumbersome syntax of `\cfrac` under \TeX it proves a bit tortuous to obtain there the same effect. Actually, it is partly for this purpose that 1.09m added `\xintGGCFrac`. We thus use `\xintFtoCx` with a suitable separator, and then the whole thing as argument to `\xintGGCFrac`:

```
\def\highlight #1{\ifnum #1>200 \fcolorbox{blue}{white}{\boldmath\color{red}$#1$}%
\else #1\fi}
\[\xintGGCFrac {\xintFtoCx {+1/\highlight}{208341/66317}}\]
```

$$3 + \frac{1}{7 + \frac{1}{15 + \frac{1}{1 + \frac{1}{\boxed{292} + \frac{1}{2}}}}}$$

16.9. `\xintFtoGC`

$\frac{Frac}{f}$ ★ `\xintFtoGC{f}` does the same as `\xintFtoCx{+1/}{f}`. Its output may thus be used in the package macros expecting such an 'inline format'.

```
566827/208524=\xintFtoGC {566827/208524}
566827/208524=2+1/1+1/2+1/1+1/1+1/1+1/4+1/1+1/1+1/6+1/1+1/1+1/8+1/1+1/1+1/11
```

16.10. `\xintFGtoC`

$\frac{Frac}{f}$ $\frac{Frac}{f}$ ★ `\xintFGtoC{f}{g}` computes the common initial coefficients to two given fractions `f` and `g`. Notice that any real number `f < x < g` or `f > x > g` will then necessarily share with `f` and `g` these common initial coefficients for its regular continued fraction. The coefficients are output as a sequence of braced numbers. This list can then be manipulated via macros from `xinttools`, or other macros of `xintcfrac`.

```
\fdef\test{\xintFGtoC{-5262046/89233}{-5314647/90125}}\texttt{\meaning\test}
macro:->{-59}{33}{27}
\fdef\test{\xintFGtoC{3.141592653}{3.141592654}}\texttt{\meaning\test}
macro:->{3}{7}{15}{1}
\fdef\test{\xintFGtoC{3.1415926535897932384}{3.1415926535897932385}}\meaning\test
macro:->{3}{7}{15}{1}{292}{1}{1}{1}{2}{1}{3}{1}{14}{2}{1}{1}{2}{2}{2}
\xintRound {30}{\xintCstoF{\xintListWithSep{,}{\test}}}
3.141592653589793238386377506390
\xintRound {30}{\xintCtoF{\test}}
3.141592653589793238386377506390
\fdef\test{\xintFGtoC{1.41421356237309}{1.4142135623731}}\meaning\test
macro:->{1}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}{2}
```

16.11. `\xintFtoCC`


$\frac{Frac}{f}$ ★ `\xintFtoCC{f}` returns the 'centered' continued fraction of `f`, in 'inline format'.

```
566827/208524=\xintFtoCC {566827/208524}
566827/208524=3+-1/4+-1/2+1/5+-1/2+1/7+-1/2+1/9+-1/2+1/11
\[\xintTeXFrac{566827/208524} = \xintGCFrac{\xintFtoCC{566827/208524}}\]
```

$$\frac{566827}{208524} = 3 - \frac{1}{4 - \frac{1}{2 + \frac{1}{5 - \frac{1}{2 + \frac{1}{7 - \frac{1}{2 + \frac{1}{9 - \frac{1}{2 + \frac{1}{11}}}}}}}}}$$

16.12. `\xintCstoF`

f ★ `\xintCstoF{a,b,c,d,...,z}` computes the fraction corresponding to the coefficients, which may be fractions or even macros expanding to such fractions. The final fraction may then be highly reducible.

 *Usage of this macro requires the user to load `xinttools`.*

Starting with release 1.09m spaces before commas are allowed and trimmed automatically (spaces after commas were already silently handled in earlier releases).

```
\[\xintGCFrac {-1+1/3+1/-5+1/7+1/-9+1/11+1/-13}=
\xintTeXsignedFrac{\xintCstoF {-1,3,-5,7,-9,11,-13}}=\xintTeXsignedFrac{\xintGctoF
{-1+1/3+1/-5+1/7+1/-9+1/11+1/-13}}\]
```

$$-1 + \frac{1}{3 + \frac{1}{-5 + \frac{1}{7 + \frac{1}{-9 + \frac{1}{11 + \frac{1}{-13}}}}}} = -\frac{75887}{118187} = -\frac{75887}{118187}$$

```
\[\xintGCFrac{{1/2}+1/{1/3}+1/{1/4}+1/{1/5}}=
\xintTeXFrac{\xintCstoF {1/2,1/3,1/4,1/5}}\]
```

$$\frac{1}{2} + \frac{1}{\frac{1}{3} + \frac{1}{\frac{1}{4} + \frac{1}{\frac{1}{5}}}} = \frac{159}{66}$$

A generalized continued fraction may produce a reducible fraction (`\xintCstoF` tries its best not to accumulate in a silly way superfluous factors but will not do simplifications which would be obvious to a human, like simplification by 3 in the result above).

16.13. `\xintCtoF`

f ★ `\xintCtoF{{a}{b}{c}...{z}}` computes the fraction corresponding to the coefficients, which may be fractions or even macros.

```
\xintCtoF {\xintApply {\xintiiPow 3}{\xintSeq {1}{5}}}
```

```
14946960/4805083
```

```
\[\xintTeXFrac{14946960/4805083}=\xintCFrac {14946960/4805083}\]
```

$$\frac{14946960}{4805083} = 3 + \frac{1}{9 + \frac{1}{27 + \frac{1}{81 + \frac{1}{243}}}}$$

In the example above the power of 3 was already pre-computed via the expansion done by `\xintApply`, but if we try with `\xintApply { \xintiiPow 3 }` where the space will stop this expansion, we can check that `\xintCtoF` will itself provoke the needed coefficient expansion.

16.14. `\xintGctoF`

f ★ `\xintGctoF{a+b/c+d/e+f/g+.....+v/w+x/y}` computes the fraction defined by the inline generalized continued fraction. Coefficients may be fractions but must then be put within braces. They can be macros. The plus signs are mandatory.

```
\[\xintGCFrac {1+\xintPow{1.5}{3}/{1/7}+{-3/5}/\xintiiFac {6}} =
\xintTeXFrac{\xintGctoF {1+\xintPow{1.5}{3}/{1/7}+{-3/5}/\xintiiFac {6}}} =
\xintTeXFrac{\xintIrr{\xintGctoF
  {1+\xintPow{1.5}{3}/{1/7}+{-3/5}/\xintiiFac {6}}}}\]
```

$$1 + \frac{3375 \cdot 10^{-3}}{\frac{1}{7} - \frac{3}{5}} = \frac{88629000}{3579000} = \frac{29543}{1193}$$

```
\[\ \xintGCFrac{{1/2}+{2/3}/{4/5}+{1/2}/{1/5}+{3/2}/{5/3}} =
  \xintTeXFrac{\xintGctoF {{1/2}+{2/3}/{4/5}+{1/2}/{1/5}+{3/2}/{5/3}}}\]
```

$$\frac{1}{2} + \frac{\frac{2}{3}}{\frac{4}{5} + \frac{\frac{1}{2}}{\frac{1}{5} + \frac{3}{5}}}} = \frac{4270}{4140}$$

The macro tries its best not to accumulate superfluous factor in the denominators, but doesn't reduce the fraction to irreducible form before returning it and does not do simplifications which would be obvious to a human.

16.15. `\xintCstoCv`

f ★ `\xintCstoCv{a,b,c,d,...,z}` returns the sequence of the corresponding convergents, each one within braces.

Usage of this macro requires the user to load `xinttools`.

It is allowed to use fractions as coefficients (the computed convergents have then no reason to be the real convergents of the final fraction). When the coefficients are integers, the convergents are irreducible fractions, but otherwise it is not necessarily the case.

```
\xintListWithSep:{\xintCstoCv{1,2,3,4,5,6}}
1/1:3/2:10/7:43/30:225/157:1393/972
```

```
\xintListWithSep:{\xintCstoCv{1,1/2,1/3,1/4,1/5,1/6}}
1/1:3/1:9/7:45/19:225/159:1575/729
```

```
\[\xintListWithSep{\to}{\xintApply\xintTeXFrac{\xintCstoCv {\xintPow
  {- .3}{-5},7.3/4.57,\xintCstoF{3/4,9,-1/3}}}}\]
```

$$\frac{-100000}{243} \rightarrow \frac{-72888949}{177390} \rightarrow \frac{-2700356878}{6567804}$$

16.16. `\xintCtoCv`

f ★ `\xintCtoCv{a}{b}{c}...{z}` returns the sequence of the corresponding convergents, each one within braces.

```
\fdef\test{\xintCtoCv {1111111111}}\texttt{\meaning\test}
macro:->{1/1}{2/1}{3/2}{5/3}{8/5}{13/8}{21/13}{34/21}{55/34}{89/55}{144/89}
```

16.17. `\xintGtoCv`

f ★ `\xintGtoCv{a+b/c+d/e+f/g+.....+v/w+x/y}` returns the list of the corresponding convergents. The coefficients may be fractions, but must then be inside braces. Or they may be macros, too.

The convergents will in the general case be reducible. To put them into irreducible form, one needs one more step, for example it can be done with `\xintApply\xintIrr`.

```
\[\xintListWithSep{,}{\xintApply\xintTeXFrac
  {\xintGtoCv{3+{-2}/{7/2}+{3/4}/12+{-56/3}}}\]
\[\xintListWithSep{,}{\xintApply\xintTeXFrac{\xintApply\xintIrr
  {\xintGtoCv{3+{-2}/{7/2}+{3/4}/12+{-56/3}}}\]
3,  $\frac{17}{7}$ ,  $\frac{834}{342}$ ,  $\frac{1306}{542}$ 
3,  $\frac{17}{7}$ ,  $\frac{139}{57}$ ,  $\frac{653}{271}$ 
```

16.18. `\xintFtoCv`

$\frac{f}{f}$ ★ `\xintFtoCv{f}` returns the list of the (braced) convergents of f , with no separator. To be treated with `\xintAssignArray` or `\xintListWithSep`.

```
\[\xintListWithSep{\to}{\xintApply\xintTeXFrac{\xintFtoCv{5211/3748}}}\]
1  $\rightarrow$   $\frac{3}{2}$   $\rightarrow$   $\frac{4}{3}$   $\rightarrow$   $\frac{7}{5}$   $\rightarrow$   $\frac{25}{18}$   $\rightarrow$   $\frac{32}{23}$   $\rightarrow$   $\frac{57}{41}$   $\rightarrow$   $\frac{317}{228}$   $\rightarrow$   $\frac{374}{269}$   $\rightarrow$   $\frac{691}{497}$   $\rightarrow$   $\frac{5211}{3748}$ 
```

16.19. `\xintFtoCCv`

$\frac{f}{f}$ ★ `\xintFtoCCv{f}` returns the list of the (braced) centered convergents of f , with no separator. To be treated with `\xintAssignArray` or `\xintListWithSep`.

```
\[\xintListWithSep{\to}{\xintApply\xintTeXFrac{\xintFtoCCv{5211/3748}}}\]
1  $\rightarrow$   $\frac{4}{3}$   $\rightarrow$   $\frac{7}{5}$   $\rightarrow$   $\frac{32}{23}$   $\rightarrow$   $\frac{57}{41}$   $\rightarrow$   $\frac{374}{269}$   $\rightarrow$   $\frac{691}{497}$   $\rightarrow$   $\frac{5211}{3748}$ 
```

16.20. `\xintCtoF`

$\frac{num}{x}$ f ★ `\xintCtoF{N}{\macro}` computes the fraction f having coefficients $c(j)=\macro{j}$ for $j=0,1,\dots,N$. The N parameter is given to a `\numexpr`. The values of the coefficients, as returned by `\macro` do not have to be positive, nor integers, and it is thus not necessarily the case that the original $c(j)$ are the true coefficients of the final f .

```
\def\macro #1{\the\numexpr 1+#1*#1\relax} \xintCtoF {5}{\macro}
72625/49902[0]
```

This example shows that the fraction is output with a trailing number in square brackets (representing a power of ten), this is for consistency with what do most macros of `xintfrac`, and does not have to be always this annoying `[0]` as the coefficients may for example be numbers in scientific notation. To avoid these trailing square brackets, for example if the coefficients are known

to be integers, there is always the possibility to filter the output via `\xintPRaw`, or `\xintIrr` (the latter is overkill in the case of integer coefficients, as the fraction is guaranteed to be irreducible then).

16.21. `\xintGCntoF`

$\frac{\text{num}}{\text{x}}$ *f f* ★ `\xintGCntoF{N}{\macroA}{\macroB}` returns the fraction *f* corresponding to the inline generalized continued fraction $a_0+b_0/a_1+b_1/a_2+\dots+b_{(N-1)}/a_N$, with $a(j)=\macroA\{j\}$ and $b(j)=\macroB\{j\}$. The *N* parameter is given to a `\numexpr`.

```
\def\coeffA #1{\the\numexpr #1+4-3*((#1+2)/3)\relax }%
\def\coeffB #1{\the\numexpr \ifodd #1 -\fi 1\relax }% (-1)^n
\[\xintGCfrac{\xintGCntoGC {6}{\coeffA}{\coeffB}} =
\xintTeXfrac{\xintGCntoF {6}{\coeffA}{\coeffB}\]
```

$$1 + \frac{1}{2 - \frac{1}{3 + \frac{1}{1 - \frac{1}{2 + \frac{1}{3 - \frac{1}{1}}}}}} = \frac{39}{25}$$

There is also `\xintGCntoGC` to get the 'inline format' continued fraction.

16.22. `\xintCntoCs`

$\frac{\text{num}}{\text{x}}$ *f f* ★ `\xintCntoCs{N}{\macro}` produces the comma separated list of the corresponding coefficients, from $n=0$ to $n=N$. The *N* is given to a `\numexpr`.

```
\xintCntoCs {5}{\macro}
```

1, 2, 5, 10, 17, 26

```
\[ \xintTeXfrac{\xintCntoF{5}{\macro}}=\xintCFrac{\xintCntoF {5}{\macro}}\]
```

$$\frac{72625}{49902} = 1 + \frac{1}{2 + \frac{1}{5 + \frac{1}{10 + \frac{1}{17 + \frac{1}{26}}}}}$$

16.23. `\xintCntoGC`

$\frac{\text{num}}{\text{x}}$ *f f* ★ `\xintCntoGC{N}{\macro}` evaluates the $c(j)=\macro\{j\}$ from $j=0$ to $j=N$ and returns a continued fraction written in inline format: $\{c(0)\}+1/\{c(1)\}+1/\dots+1/\{c(N)\}$. The parameter *N* is given to a `\numexpr`. The coefficients, after expansion, are, as shown, being enclosed in an added pair of braces, they may thus be fractions.

```
\def\macro #1{\the\numexpr\ifodd#1 -1-#1\else1+#1\fi\relax/\the\numexpr 1+#1*#1\relax}
\fdef\x{\xintCntoGC {5}{\macro}}\meaning\x
\[\xintGCfrac{\xintCntoGC {5}{\macro}}\]
```

`macro:->` $\{1/\{1/\{1+0*0\}\relax }+1/\{-2/\{1+1*1\}\relax }+1/\{3/\{1+2*2\}\relax }+1/\{-4/\{1+3*3\}\relax }+1/\{5/\{1+4*4\}\relax }+1/\{-6/\{1+5*5\}\relax }\}$

`\numexpr 1+5*5\relax }`

$$1 + \frac{1}{\frac{-2}{2} + \frac{1}{\frac{3}{5} + \frac{1}{\frac{-4}{10} + \frac{1}{\frac{5}{17} + \frac{-6}{26}}}}}$$

16.24. `\xintGCntoGC`

$\frac{\text{num}}{\text{x}}$ *ff* ★ `\xintGCntoGC{N}{\macroA}{\macroB}` evaluates the coefficients and then returns the corresponding $\{a_0\}+\{b_0\}/\{a_1\}+\{b_1\}/\{a_2\}+\dots+\{b(N-1)\}/\{a_N\}$ inline generalized fraction. *N* is given to a `\numexpr`. The coefficients are enclosed into pairs of braces, and may thus be fractions, the fraction slash will not be confused in further processing by the continued fraction slashes.

```
\def\an #1{\the\numexpr #1*#1*#1+1\relax}%
\def\bn #1{\the\numexpr \ifodd#1 -\fi 1*(#1+1)\relax}%
$\xintGCntoGC {5}{\an}{\bn}=\xintGCFrac {\xintGCntoGC {5}{\an}{\bn}} =
\displaystyle\xintTeXFrac {\xintGCntoF {5}{\an}{\bn}}$\par
```

$$1 + 1/2 + -2/9 + 3/28 + -4/65 + 5/126 = 1 + \frac{1}{2 - \frac{3}{9 + \frac{4}{28 - \frac{5}{65 + \frac{5}{126}}}}} = \frac{5797655}{3712466}$$

16.25. `\xintCstoGC`

f ★ `\xintCstoGC{a,b,...,z}` transforms a comma separated list (or something expanding to such a list) into an 'inline format' continued fraction $\{a\}+1/\{b\}+1/\dots+1/\{z\}$. The coefficients are just copied and put within braces, without expansion. The output can then be used in `\xintGCFrac` for example.

```
\[\xintGCFrac {\xintCstoGC {-1,1/2,-1/3,1/4,-1/5}}=\xintTeXsignedFrac{\xintCstoF {-1,1/2,-1/3,1/4,-1/5}}\]
```

$$-1 + \frac{1}{\frac{1}{2} + \frac{1}{\frac{-1}{3} + \frac{1}{\frac{1}{4} + \frac{-1}{5}}}} = -\frac{145}{83}$$

16.26. `\xintiCstoF`, `\xintiGctoF`, `\xintiCstoCv`, `\xintiGctoCv`

f ★ Essentially the same as the corresponding macros without the 'i', but for integer-only input. Infinitesimally faster, mainly for internal use by the package.

16.27. `\xintGctoGC`

f ★ `\xintGctoGC{a+b/c+d/e+f/g+.....+v/w+x/y}` expands (with the usual meaning) each one of the co-

efficiently and returns an inline continued fraction of the same type, each expanded coefficient being enclosed within braces.

```
\fdef\x {\xintGCtoGC {1+\xintPow{1.5}{3}/{1/7}+{-3/5}/%
\xintiFac {6}+\xintCstoF {2,-7,-5}/16}} \meaning\x
macro:->{1}+{3375/1[-3]}/{1/7}+{-3/5}/{720}+{67/36}/{16}
```

To be honest I have forgotten for which purpose I wrote this macro in the first place.

16.28. Euler's number e

Let us explore the convergents of Euler's number e . The volume of computation is kept minimal by the following steps:

- a comma separated list of the first 36 coefficients is produced by `\xintCntoCs`,
- this is then given to `\xintiCstoCv` which produces the list of the convergents (there is also `\xintCstoCv`, but our coefficients being integers we used the infinitesimally faster `\xintiCstoCv`),
- then the whole list was converted into a sequence of one-line paragraphs, each convergent becomes the argument to a macro printing it together with its decimal expansion with 30 digits after the decimal point.
- A count register `\cnta` was used to give a line count serving as a visual aid: we could also have done that in an expandable way, but well, let's relax from time to time...

```
\def\cn #1{\the\numexpr\ifcase \numexpr #1+3-3*((#1+2)/3)\relax
1\or1\or2*(#1/3)\fi\relax }
% produces the pattern 1,1,2,1,1,4,1,1,6,1,1,8,... which are the
% coefficients of the simple continued fraction of e-1.
\cnta 0
\def\mymacro #1{\advance\cnta by 1
\noindent
\hbox {\hfil\small\dt{ \the\cnta.} }%
$\xintTrunc {30}{\xintAdd {1[0]}{#1}}\dots=
\xintTeXFrac{\xintAdd {1[0]}{#1}}{#1}}%
\xintListWithSep{\vtop to 6pt}{\vbox to 12pt}{\par}
{\xintApply\mymacro{\xintiCstoCv{\xintCntoCs {35}{\cn}}}}
```

1. 2.00000000000000000000000000000000... = 2
2. 3.00000000000000000000000000000000... = 3
3. 2.66666666666666666666666666666666... = $\frac{8}{3}$
4. 2.75000000000000000000000000000000... = $\frac{11}{4}$
5. 2.714285714285714285714285714285... = $\frac{19}{7}$
6. 2.71875000000000000000000000000000... = $\frac{87}{32}$
7. 2.717948717948717948717948717948... = $\frac{106}{39}$
8. 2.718309859154929577464788732394... = $\frac{193}{71}$
9. 2.718279569892473118279569892473... = $\frac{1264}{465}$
10. 2.718283582089552238805970149253... = $\frac{1457}{536}$
11. 2.718281718281718281718281718281... = $\frac{2721}{1001}$

12. $2.718281835205992509363295880149\dots = \frac{23225}{8544}$
13. $2.718281822943949711891042430591\dots = \frac{25946}{9545}$
14. $2.718281828735695726684725523798\dots = \frac{49171}{18089}$
15. $2.718281828445401318035025074172\dots = \frac{517656}{190435}$
16. $2.718281828470583721777828930962\dots = \frac{566827}{208524}$
17. $2.718281828458563411277850606202\dots = \frac{1084483}{398959}$
18. $2.718281828459065114074529546648\dots = \frac{13580623}{4996032}$
19. $2.718281828459028013207065591026\dots = \frac{14665106}{5394991}$
20. $2.718281828459045851404621084949\dots = \frac{28245729}{10391023}$
21. $2.718281828459045213521983758221\dots = \frac{410105312}{150869313}$
22. $2.718281828459045254624795027092\dots = \frac{438351041}{161260336}$
23. $2.718281828459045234757560631479\dots = \frac{848456353}{312129649}$
24. $2.718281828459045235379013372772\dots = \frac{14013652689}{5155334720}$
25. $2.718281828459045235343535532787\dots = \frac{14862109042}{5467464369}$
26. $2.718281828459045235360753230188\dots = \frac{28875761731}{10622799089}$
27. $2.718281828459045235360274593941\dots = \frac{534625820200}{196677847971}$
28. $2.718281828459045235360299120911\dots = \frac{563501581931}{207300647060}$
29. $2.718281828459045235360287179900\dots = \frac{1098127402131}{403978495031}$
30. $2.718281828459045235360287478611\dots = \frac{22526049624551}{8286870547680}$
31. $2.718281828459045235360287464726\dots = \frac{23624177026682}{8690849042711}$
32. $2.718281828459045235360287471503\dots = \frac{46150226651233}{16977719590391}$
33. $2.718281828459045235360287471349\dots = \frac{1038929163353808}{382200680031313}$
34. $2.718281828459045235360287471355\dots = \frac{1085079390005041}{399178399621704}$
35. $2.718281828459045235360287471352\dots = \frac{2124008553358849}{781379079653017}$
36. $2.718281828459045235360287471352\dots = \frac{52061284670617417}{19152276311294112}$

One can with no problem compute much bigger convergents. Let's get the 200th convergent. It turns out to have the same first 268 digits after the decimal point as $e - 1$. Higher convergents get more and more digits in proportion to their index: the 500th convergent already gets 799 digits correct! To allow speedy compilation of the source of this document when the need arises, I limit here to the 200th convergent.

```
\fdef\z {\xintCtoF {199}{\cn}}%
\begingroup\parindent 0pt \leftskip 2.5cm
\indent\llap {Numerator = }\printnumber{\xintNumerator\z}\par
\indent\llap {Denominator = }\printnumber{\xintDenominator\z}\par
\indent\llap {Expansion = }\printnumber{\xintTrunc{268}\z}\dots\par\endgroup
Numerator = 568964038871896267597523892315807875293889017667917446057232024547192296961118
23017524386017499531081773136701241708609749634329382906
Denominator = 331123817669737619306256360816356753365468823729314438156205615463246659728581
```

TOC, xint bundle, xintkernel, xintcore, xint, xintfrac, xintbinhex, xintgcd, xintseries, xintcfrac

```
86546133769206314891601955061457059255337661142645217223
Expansion = 1.7182818284590452353602874713526624977572470936999595749669676277240766303535
475945713821785251664274274663919320030599218174135966290435729003342952605956
307381323286279434907632338298807531952510190115738341879307021540891499348841
675092447614606680822648001684774118. . .
```

One can also use a centered continued fraction: we get more digits but there are also more computations as the numerators may be either 1 or -1.

Part III.

The **xintexpr** and allied packages source code (not included, see **sourcexint.pdf**)

This documentation has been compiled without including the commented source code, which is available in the separate file `sourcexint.pdf`, and will open in a PDF viewer via:

```
texdoc sourcexint.pdf
```

To produce a file `xint.pdf` including both the user documentation and the source code:

- acquire `xint.dtx` from
<https://www.ctan.org/tex-archive/macros/generic/xint>
- execute `latex xint.dtx` sufficiently many times (at least three), then `dvipdfmx`, or use `pdflatex`.

For those on Unix-like systems, you can (recommended) alternatively also download a file `Makefile` from the above location then execute `make xint-all.pdf`. This requires `latexmk` and will do automatically the needed number of passes.

Without the `Makefile` one can first execute `etex` on `xint.dtx` to extract files then run

```
make --file=Makefile.mk xint-all.pdf
```

Other `make` targets include `xint.pdf` (without implementation part) and `sourcexint.pdf`. Run `make help` for help. Further build alternatives are explained in extracted file `xint-all.tex`.

The file with both user manual and implementation (i.e. either `xint.pdf` from direct runs of `(pdf)latex` on `xint.dtx` or `xint-all.pdf` if using `make`) will contain hyperlinks from the user documentation to the implementation part and back. And the macros inside the code lines are hyperlinked to the location of their definition.

New with
1.4m