
The ε -TeX manual

Version 2, February 1998

by The $\mathcal{N}\mathcal{T}\mathcal{S}$ Team

Peter Breitenlohner, Max-Planck-Institut für Physik, München

Revised for VTeX/Free, May 2003

Walter Schmidt

The preparation of this report was supported in part by DANTE,
Deutschsprachige Anwendervereinigung TeX e.V.
'TeX' is a trademark of the American Mathematical Society.

1 Introduction

The $\mathcal{N}\mathcal{T}\mathcal{S}$ project intends to develop an 'New Typesetting System' ($\mathcal{N}\mathcal{T}\mathcal{S}$) that will eventually replace today's TeX3. The $\mathcal{N}\mathcal{T}\mathcal{S}$ program will include many features missing in TeX, but there will also exist a mode of operation that is 100% compatible with TeX3. It will, necessarily, require quite some time to develop $\mathcal{N}\mathcal{T}\mathcal{S}$ to maturity and make it widely available.

Meanwhile ε -TeX intends to fill the gap between TeX3 and the future $\mathcal{N}\mathcal{T}\mathcal{S}$. It consists of a series of features extending the capabilities of TeX3.¹

Since compatibility between ε -TeX and TeX3 has been a main concern, ε -TeX has two modes of operation:

- (1) In TeX compatibility mode it fully deserves the name TeX and there are neither extended features nor additional primitive commands. That means in particular that ε -TeX passes the TRIP test [1] without any restriction. There are, however, a few minor modifications that would be legitimate in any implementation of TeX.
- (2) In extended mode there are additional primitive commands and the extended features of ε -TeX are available.

We have tried to make ε -TeX as compatible with TeX as possible even in extended mode. In a few cases there are, however, some subtle differences described in detail later on. Therefore the ε -TeX features available in extended mode are grouped into two categories:

- (1) Most of them have no semantic effect as long as none of the additional primitives are executed; these 'extensions' are permanently enabled.
- (2) The remaining optional ε -TeX features ('enhancements') can be individually enabled and disabled; initially they are all disabled. For each enhancement there is a state variable `\...state`; an enhancement is enabled or

¹ The TeX3 program; for the moment there are no plans to extend the software related to TeX.

disabled by assigning a positive or non-positive value respectively to that state variable.

For ϵ -TeX Versions 1 and 2 there is just one enhancement: mixed direction typesetting (TeX--XeT) with the state variable `\TeXXeTstate`.

Version 1.1 of ϵ -TeX was released in November 1996, Version 2.0 in February 1998. It is expected that there will be about one ϵ -TeX version per year, where each later version adds new features. It would be desirable if these ϵ -TeX versions were incorporated into many of the existing implementations of TeX3 without much delay.

With each ϵ -TeX version there will be an e-TRIP test [2] in order to help to verify that a particular implementation deserves the name ϵ -TeX in the same way as the TRIP test [1] helps to verify that an implementation deserves the name TeX.

2 Generating Format Files for ϵ -TeX

Extended mode is turned on, when a format file is read, that was already generated in extended mode. To generate such a format file under VTeX, specify the command line option `-2` beside `-i`.

Specifying the option `-2` when typesetting a document will also enable the ϵ -TeX extensions, regardless of the format file. Thus, it is *not necessary* to create particular ϵ -TeX-based formats.

3 ϵ -TeX Extensions

3.1 Compatibility and Extended Mode

Once ϵ -TeX has entered compatibility mode it behaves as any other implementation of TeX. All of ϵ -TeX's additional commands are absent; it is therefore impossible to access any of the extensions or enhancements. The ability of eNITEX to initially choose between compatibility and extended mode is, however, by itself a feature not present in any TeX implementation.

The remainder of this document is devoted to a detailed and mostly technical description of all aspects where ϵ -TeX (in extended mode) behaves differently from TeX. It will be assumed that the reader is familiar with *The TeXbook* [3] describing TeX's behaviour in quite some detail.

All of ϵ -TeX's extensions and enhancements available in extended mode are activated by either executing some new primitive command or by assigning a nonzero value to some new integer parameter or state variable. Since all these new variables are initially zero,² ϵ -TeX behaves as TeX as long as none of ϵ -TeX's new control sequences are used, with the following exceptions which should, however, have no effect on the typesetting of error-free TeX

² To be precise all state variables are zero when eNITEX or eVIRTEX is started; integer parameters that are not state variables are zero when eNITEX is started without loading a format file or inherited from the format file otherwise.

documents (produced with error-free formats):

- (1) When `\tracingcommands` has a value of 3 or more, or when `\tracinglostchars` has a value of 2 or more, ϵ -TeX will display additional information not available in TeX.
- (2) When using a count, dimen, skip, muskip, box, or token register number in the range 256–32767, ϵ -TeX will access one of its additional registers whereas TeX will produce an error and use register number zero. [Under VTeX, the allowed range of registers goes up to 65535, rather than 32767.](#)

3.2 Optimization

When a value is assigned to an \langle internal quantity \rangle within a save group, the former value is restored when the group ends, provided the assignment was not global. This is achieved by saving the former value on TeX's 'save stack'. ϵ -TeX refrains from creating such save stack entries when the old and new value are the same ('reassignments').

`\aftergroup` tokens are also kept on TeX's save stack. When the current group ends, TeX converts each `\aftergroup` token into a token list and inserts this list as new 'input level' into the input stack. ϵ -TeX collects all `\aftergroup` tokens from one group into one token list and thus conserves input levels.

When a completed page is written to the DVI file (shipped out), TeX multiplies the relevant stretch or shrink components of glue nodes in a box by the glue expansion factor of that box and converts the product to DVI units. In order to avoid overflow each resulting value x is artificially limited to the range $|x| \leq 10^9$. Consider the example:

```
\shipout\vbox to100pt{
  \hrule width10pt
  \vskip 0pt plus1000fil
  \vskip 0pt plus1000fil
  \vskip 0pt plus-2000fil
  \hrule
  \vskip 0pt plus0.00005fil
}
```

Here the three glues between the two rules add up to zero; when TeX converts each stretch component individually they will, however, add up to 10^9 DVI units due to the truncation mentioned above. ϵ -TeX, however, accumulates the relevant stretch or shrink components of consecutive glue nodes (possibly separated by insert, mark, adjust, kern, and penalty nodes) before converting them to DVI units. During this process glue nodes may be converted into equivalent kern nodes and some glue specifications may be recycled; this may affect the memory usage statistics displayed after the page has been shipped out.

3.3 Tracing and Diagnostics

When `\tracingcommands` has a value of 3 or more, the commands following a prefix (`\global`, etc.) are shown as well, e.g.:

```
\global\count0=0    =>    {\global}
                        {\count}
```

When `\tracinglostchars` has a value of 2 or more, missing characters are displayed on the terminal even if the value of `\tracingonline` is 0 or less.

When `\tracingscantokens` has a value of 1 or more, the opening and closing of pseudo-files (generated by `\scantokens`) is recorded as for any other file, with ‘_’ as filename.

When the program is compiled with the code for collecting statistics and `\tracingassigns` has a value of 1 or more, all assignments subject to TeX’s grouping mechanism are traced, e.g.:

```
\def\foo{\relax}    =>    {changing \foo=undefined}
                        {into \foo=macro:->\relax }
\global\count17=7    =>    {globally changing \count17=0}
                        {into \count17=7}
\count17=7           =>    {reassigning \count17=7}
```

When `\tracingifs` has a value of 1 or more, all conditionals (including `\unless`, `\or`, `\else`, and `\fi`) are traced, together with the starting line and nesting level; the `\showifs` command displays the state of all currently active conditionals. Thus the input

```
\unless\iffalse
  \iffalse
  \else
    \showifs
  \fi
\fi
```

might yield

```
{\unless\iffalse: (level 1) entered on line 1}
{\iffalse: (level 2) entered on line 2}
{\else: \iffalse (level 2) entered on line 2}
### level 2: \iffalse\else entered on line 2
### level 1: \unless\iffalse entered on line 1
{\fi: \iffalse (level 2) entered on line 2}
{\fi: \unless\iffalse (level 1) entered on line 1}
```

When `\tracinggroups` has a value of 1 or more, the start and end of each save group is traced, together with the starting line and grouping level; the `\showgroups` command displays the state of all currently active save groups. Thus the input

```
\begingroup
{
  \showgroups
}
\endgroup
```

might yield

```
{entering semi simple group (level 1) at line 1}
{entering simple group (level 2) at line 2}
```

```

### simple group (level 2) entered at line 1 ({)
### semi simple group (level 1) entered at line 1 (\begingroup)
### bottom level
{leaving simple group (level 2) entered at line 2}
{leaving semi simple group (level 1) entered at line 1}

```

Occasionally conditionals and/or save groups are not properly nested with respect to `\input` files. Although this might be perfectly legitimate, such anomalies are mostly unintentional and may cause quite obscure errors. When `\tracingnesting` has a value of 1 or more, these anomalies are shown; when `\tracingnesting` has a value of 2 or more, the current context (trace-back) is shown as well. Thus the input

```

\newlinechar='\^^J
\begingroup
  \iftrue
    \scantokens{%
      \endgroup
^^J\fi
^^J\bgroup
  ^^ \tracingnesting=2
  ^^J\iffalse
  ^^J\else
    }%
  \egroup
\fi

```

might yield³

Warning: end of semi simple group (level 1) entered at line 2 of a different file

Warning: end of `\iftrue` entered on line 3 of a different file

Warning: end of file when simple group (level 1) entered at line 3 is incomplete

Warning: end of file when `\iffalse\else` entered on line 5 is incomplete

1.7 `\else`

```

1.11      }
          %

```

The command `\showtokens{<token list>}` displays the token list, and allows the display of quantities that cannot be displayed by `\show` or `\showthe`, e.g.:

```

\showtokens\expandafter{\jobname}
\showtokens\expandafter{\topmarks 27}

```

3 The `\scantokens` command will be discussed later.

3.4 Status Enquiries

A number of TeX's internal quantities can be assigned values but these values cannot be retrieved in TeX. ϵ -TeX introduces several new primitives that allow the retrieval of information about its internal state.

`\eTeXversion` returns ϵ -TeX's (major) version number;

`\eTeXrevision` expands into a list of character tokens representing the revision (minor version) number. Thus

```
\message{\number\eTeXversion\eTeXrevision}
```

should write the complete version as shown when ϵ -TeX is started.

When used as number, `\interactionmode` returns one of the values 0 (batch-mode), 1 (nonstopmode), 2 (scrollmode), or 3 (errorstopmode). Assigning one of these values to `\interactionmode` changes the current interaction mode accordingly; such assignments are always global.

`\currentgrouplevel` returns the current save group level;

`\currentgrouptype` returns a number representing the type of the innermost group:

0: bottom level (no group)	9: math group
1: simple group	10: disc group
2: hbox group	11: insert group
3: adjusted hbox group	12: vcenter group
4: vbox group	13: math choice group
5: vtop group	14: semi simple group
6: align group	15: math shift group
7: no align group	16: math left group
8: output group	

`\currentiflevel` returns the number of currently active conditionals;

`\currentifbranch` indicates which branch of the innermost conditional is taken: 1 'then branch', -1 'else branch', or 0 not yet decided;

`\currentiftype` returns 0 if there are no active conditionals, a positive number indicating the type of the innermost active conditional, or the negative of that number when the conditional was prefixed by `\unless`:

1: <code>\if</code>	8: <code>\ifmmode</code>	15: <code>\iftrue</code>
2: <code>\ifcat</code>	9: <code>\ifinner</code>	16: <code>\iffalse</code>
3: <code>\ifnum</code>	10: <code>\ifvoid</code>	17: <code>\ifcase</code>
4: <code>\ifdim</code>	11: <code>\ifhbox</code>	18: <code>\ifdefined</code>
5: <code>\ifodd</code>	12: <code>\ifvbox</code>	19: <code>\ifcsname</code>
6: <code>\ifvmode</code>	13: <code>\ifx</code>	20: <code>\iffontchar</code>
7: <code>\ifhmode</code>	14: <code>\ifeof</code>	

`\lastnodetype` returns a number indicating the type of the last node, if

any, on the current (vertical, horizontal, or math) list:

-1: none (empty list)	8: disc node
0: char node	9: whatsit node
1: hlist node	10: math node
2: vlist node	11: glue node
3: rule node	12: kern node
4: ins node	13: penalty node
5: mark node	14: unset node
6: adjust node	15: math mode nodes
7: ligature node	

The commands `\fontcharht`, `\fontcharwd`, `\fontchardp`, and `\fontcharic` followed by a font specification and a character code, return a dimension: the height, width, depth, or italic correction of the character in the font, or 0pt if no such character exists; the conditional `\iffontchar` tests the existence of that character.

When used as number, `\parshape` returns the number of lines of the current parshape specification (or zero).

ϵ -TeX's `\parshapeindent`, `\parshapelength`, and `\parshapedimen`, followed by a number n return the dimensions of the parshape specification:

0pt for $n \leq 0$ or when no parshape is currently active, otherwise

`\parshapeindent n` and `\parshapedimen $2n - 1$` both return the indentation of line n (explicitly specified or implied by repeating the last specification),

`\parshapelength n` and `\parshapedimen $2n$` both return the length of line n .

3.5 Expressions

ϵ -TeX introduces the notion of expressions of type number, dimen, glue, or muglue, that can be used whenever a quantity of that type is needed. Such expressions are evaluated by ϵ -TeX's scanning mechanism; they are initiated by one of the commands `\numexpr`, `\dimexpr`, `\glueexpr`, or `\muexpr` (determining the type t) and optionally terminated by one `\relax` (that will be absorbed by the scanning mechanism). An expression consists of one or more terms of the same type to be added or subtracted; a term of type t consists of a factor of that type, optionally multiplied and/or divided by numeric factors; finally a factor of type t is either a parenthesized subexpression or a quantity (number, etc.) of that type. Thus, the conditional

```
\ifdim\dimexpr (2pt-5pt)*\numexpr 3-3*13/5\relax + 34pt/2<\wd20
```

is true if and only if the width of box 20 exceeds 32pt. Note the use of `\relax` to terminate the inner (numeric) expression, the outer (dimen) expression is terminated automatically by the token `<12` that does not fit into the expression syntax.

The arithmetic performed by ϵ -TeX's expressions does not do much that could not be done by TeX's arithmetic operations `\advance`, `\multiply`,

and `\divide`, although there are some notable differences: Each factor is checked to be in the allowed range, numbers must be less than 2^{31} in absolute value, dimensions or glue components must be less than 2^{14} pt, `mu`, `fil`, etc. respectively. The arithmetic operations are performed individually, except for ‘scaling’ operations (a multiplication immediately followed by a division) which are performed as one combined operation with a 64-bit product as intermediate value. The result of each operation is again checked to be in the allowed range. Finally the results of divisions and scalings are rounded, whereas TeX’s `\divide` truncates.

The important new feature is, however, that the evaluation of expressions does not involve assignments and can therefore be performed in circumstances where assignments are not allowed, e.g., inside an `\edef` or `\write`. This also allows the definition of purely expandable loop constructions:

```
\def\foo#1#2{\number#1
  \ifnum#1<#2,
    \expandafter\foo
    \expandafter{\number\numexpr#1+1\expandafter}%
    \expandafter{\number#2\expandafter}%
  \fi}
```

such that, e.g., ‘`\foo{7}{13}`’ expands into ‘7, 8, 9, 10, 11, 12, 13’.

The commands `\gluestretch` and `\glueshrink` are to be followed by a glue specification and return the stretch or shrink component of that glue as dimensions (with `fil` etc. replaced by `pt`), the commands `\gluestretchorder` and `\glueshrinkorder` return the order of infinity: 0 for `pt`, 1 for `fil`, 2 for `fill`, and 3 for `filll`.

The commands `\gluetomu` and `\mutoglu` convert glue into `muglu` and vice versa by simply equating 1 pt with 1 mu, precisely what TeX does (in addition to an error message) when the wrong kind of glue is used.

3.6 Additional Registers and Marks

ϵ -TeX increases the number of TeX’s count, dimen, skip, muskip, box, and token registers from 256 to 32768. The additional registers, numbered 256–32767, can be used exactly as the first 256, except that they can not be used for insertion classes. [Under VTeX, the allowed range of registers goes up to 65535, rather than 32767. All registers are implemented in the same way; \$\epsilon\$ -TeX sparse arrays are not used.](#)

As in TeX, the first 256 registers of each kind are realized as static arrays that are part of the ‘table of equivalents’; values to be restored when a save group ends are kept on the save stack. The additional registers are realized as sparse arrays built from TeX’s main memory and are therefore less efficient. They use a four-level index structure and individual registers are present only when needed. Values to be restored when a particular save group ends are kept in a linked list (again built from main memory) with one save stack entry pointing to that list.⁴

⁴ With the effect that the order of restoring (or discarding) saved values may be somewhat surprising. [VTeX is free of this bug, the order of restoring/discarding is as expected.](#)

ϵ -TeX generalizes TeX's mark concept to mark classes 0-32767, with mark class 0 used for TeX's marks.

The command `\marks` followed by a mark class n and a mark text appends a mark node to the current list; `\marks0` is synonymous with `\mark`. The page builder and the `\vsplit` command record information about the mark nodes found on the page or box produced, separately for each mark class. The information for mark class 0 is kept in a small static array as in TeX, the information for the additional mark classes is again kept in a sparse array with entries present only when needed.

The command `\firstmarks n` expands to the mark text for mark class n first encountered on the most recent page, etc., and again `\firstmarks0` is synonymous with `\firstmark`.

3.7 Input Handling

The command `\readline<number> to <control sequence>` defines the control sequence as parameterless macro whose replacement text is the contents of the next line read from the designated file, as for `\read`. The difference is that the current category codes are ignored and all characters on that line (including an endline character) are converted to character tokens with category 12 ('other'), except that the character code 32 gets category 10 ('space').

The command `\scantokens{...}` absorbs a list of unexpanded tokens, converts it into a character string that is treated as if it were an external file, and starts to read from this 'pseudo-file'. A rather similar effect can be achieved by the commands

```
\toks0={...}  
\immediate\openout0=file  
\immediate\write0{\the\toks0}  
\immediate\closeout0  
\input file
```

In particular every occurrence of the current newline character is interpreted as start of a new line, and input characters will be converted into tokens as usual. The `\scantokens` command is, however, expandable and does not use token registers, write streams, or external files. Furthermore the conversion from TeX's internal ASCII codes to external characters and back to ASCII codes is skipped. Finally the current context (traceback) shown, e.g., as part of an error message continues beyond an input line from a pseudo-file until an input line from a real file (or the terminal) is found.

When ϵ -TeX's input mechanism attempts to read beyond the end of an `\input` file or `\scantokens` pseudo-file, and before checking for 'runaway' conditions and closing the file, it will first read a list of tokens that has been predefined by the command `\everyeof={<token list>}`.

3.8 Breaking Paragraphs into Lines

Traditional typesetting with lead type used to adjust (stretch or shrink) the interword spaces in the last line of a paragraph by the same amount as those

in the preceding line. With TeX the last line is, however, usually typeset at its natural width due to infinitely stretchable `parfillskip` glue. ϵ -TeX allows interpolation between these two extremes by specifying a suitable value for `\lastlinefit`. For a value of 0 or less, ϵ -TeX behaves as TeX, values from 1 to 1000 indicate a glue adjustment fraction f times 1000, values above 1000 are interpreted as $f = 1$.

The new algorithm is used only if

1. `\lastlinefit` is positive;
2. `\parfillskip` has infinite stretchability; and
3. the stretchability of `\leftskip` plus `\rightskip` is finite.⁵

Thus the last line of a paragraph would normally be typeset at its natural width and the stretchability of `parfillskip` glue would be used to achieve the desired line width. The algorithm proceeds as usual, considering all possible sequences of feasible break points and accumulating demerits for the stretching or shrinking of lines as well as for visually incompatible lines. When a candidate for the last line has been reached, the following conditions are tested:

4. the previous line was not ‘infinitely bad’ and was stretched with positive finite stretchability or was shrunk with positive shrinkability;
5. the last line has infinite stretchability entirely due to `parfillskip` glue;
6. if the previous line was stretched or shrunk the last line has positive finite stretchability or shrinkability respectively.

If all three conditions are satisfied, a glue adjustment factor of f times that of the preceding line will be applied to the relevant stretch or shrink components of all glue nodes in the last line, and the corresponding demerits are computed. (The last line will, however, not be stretched beyond the desired line width.)

When all possible candidates for the last line of the paragraph have been examined, the one having fewest accumulated demerits is chosen. If ϵ -TeX’s modified algorithm was applied to that last line, the actual stretching or shrinking is achieved by suitably modifying the `parfillskip` glue node.

All computations described so far are performed with machine-independent integer arithmetic. Note, however, that the actual stretching requires machine-dependent floating point arithmetic. Therefore, when a paragraph is interrupted by a displayed equation and the line preceding the display is subject to the adjustment just described, the display will in general be preceded by `abovedisplayskip` and not by `abovedisplayshortskip` glue.

After breaking a paragraph into lines, TeX computes the interline penalties by adding the values of:

`\interlinepenalty` between any two lines,
`\clubpenalty` after the first line of a (partial) paragraph,
`\widowpenalty` before the last line of the paragraph,
`\displaywidowpenalty` before the line immediately preceding a displayed equation, and
`\brokenpenalty` after lines ending with a discretionary break.

⁵ As usual for parameters influencing TeX’s line-breaking algorithm, the values current at the end of the (partial) paragraph are used.

ϵ -TeX generalizes the concept of interline, club, widow, and display widow penalty by allowing their replacement by arrays of penalty values with the commands

`\interlinepenalties`,
`\clubpenalties`,
`\widowpenalties`, and
`\displaywidowpenalties`.

Each of these commands is to be followed by an optional equal sign and a number n . If $n \leq 0$ the respective array is reset and TeX's corresponding single value is used as usual; a positive value n declares an array of length n and must be followed by n penalty values. When one of these arrays has been set, its values are used instead of TeX's corresponding single values as follows (repeating the last value when necessary):

the i^{th} interline penalty value is used after line i of the paragraph;
the i^{th} club penalty value is used after line i of a partial paragraph;
the i^{th} widow penalty value is used after line $m - i$ of a paragraph without displayed equations or the last partial paragraph of length m ;
the i^{th} display widow penalty value is used after line $m - i$ of a partial paragraph of length m that is followed by a displayed equation.

When used after `\the` or in situations where TeX expects to see a number, the same four commands serve to retrieve the arrays of penalties. Specifying, e.g., `\clubpenalties<number>` with a number n , returns 0 for $n < 0$ or when the club penalty array has been reset, the length of the declared club penalty array for $n = 0$, or the n^{th} club penalty value for $n > 0$ (again repeating the last value when necessary).

3.9 Math Formulas

TeX's `\left<delimiter>\dots\right<delimiter>` produces two delimiters with a common size adjusted to the height and depth of the enclosed material. In ϵ -TeX this can be generalized by occurrences of `\middle<delimiter>` dividing the enclosed material into segments resulting in a sequence of delimiters with a common size adjusted to the maximal height and depth of all enclosed segments. The spacing between a segment and the delimiter to its left or right is as for TeX's left or right delimiter respectively.

3.10 Hyphenation

TeX uses the `\lccode` values for two quite unrelated purposes:

- (1) when `\lowercase` converts character tokens to their lower-case equivalents (in the same way as `\uppercase` uses the `\uccode` values); and
- (2) when hyphenation patterns or exceptions are read, and when words are hyphenated during the line-breaking algorithm.

ϵ -TeX introduces the concept of (language-dependent) hyphenation codes that are used instead of the `\lccode` values for hyphenation purposes. In order to explain the details of ϵ -TeX's behaviour, we need some technical aspects of hyphenation patterns. When INITEX starts without reading a format file, the (initially empty) hyphenation patterns are in a form suitable for

inserting new patterns specified by `\patterns` commands; when INITEX attempts hyphenation or prepares to write a format file, they are compressed into a more compact form suitable for finding hyphens. Only these compressed patterns can be read from a format file (by INITEX or VIRTEX).

In ϵ -TeX the hyphenation patterns are supplemented by hyphenation codes. When eINITEX starts without reading a format file both are initially empty; when a `\patterns` command is executed and `\savingshyphcodes` has a positive value, the current `\lccode` values are saved as hyphenation codes for the current language. These saved hyphenation codes are later compressed together with the patterns and written to or read from a format file. When the patterns have been compressed (always true for eVIRTEX) and hyphenation codes have been saved for the current language, they are used instead of the `\lccode` values for hyphenation purposes (reading hyphenation exceptions and hyphenating words).

3.11 Discarded Items

When TeX's page builder transfers (vertical mode) material from the 'recent contributions' to the 'page so far', it discards glue, kern, and penalty nodes (discardable items) preceding the first box or rule on the page under construction and inserts a topskip glue node immediately before that box or rule. Note, however, that this topskip glue need not be the first node on the page, it may be preceded by insertion, mark, and whatsit nodes. Similarly when the `\vsplit` command has split the first part off a vbox, discardable items are discarded from the top of the remaining vbox and a splittopskip glue node is inserted immediately before the first box or rule.

When ϵ -TeX's parameter `\savingsdiscards` has been assigned a positive value, these 'discarded items' are saved in two lists and can be recovered by the commands `\pagediscards` and `\splitdiscards` that act like 'unvboxing' hypothetical box registers containing a vbox with the discarded items.

The list of items discarded by the page builder is emptied at the end of the output routine and by the `\pagediscards` command; new items may be added as long as the new 'page so far' contains no box or rule.

The list of items discarded by the `\vsplit` command is emptied at the start of a vsplit operation and by the `\splitdiscards` command; new items are added at the end of a vsplit operation.

3.12 Expandable Commands

Chapter 20 of *The TeXbook* gives complete lists of all expandable TeX commands and of all cases where expandable tokens are not expanded. For ϵ -TeX there are these additional conditionals:

▷ `\ifdefined⟨token⟩` (test if token is defined)

True if `⟨token⟩` is defined; creates no new hash table entry.

▷ `\ifcsname... \endcsname` (test if control sequence is defined)

True if the control sequence `\csname... \endcsname` would be defined; creates no new hash table entry.

- ▷ `\iffontchar````<8-bit number>` (test if char exists)

True if `\char``<8-bit number>` in `\font``` exists.

These are ϵ -TeX's additional expandable commands:

- ▷ `\unless`.
The next (unexpanded) token must be a boolean conditional (i.e., not `\ifcase`); the truth value of that conditional is reversed.
- ▷ `\eTeXrevision`.
The expansion is a list of character tokens of category 12 ('other') representing ϵ -TeX's revision (minor version) number, e.g., `'0'` or `'1'`.
- ▷ `\topmarks``<15-bit number>`, `\firstmarks``<15-bit number>`, `\botmarks``<15-bit number>`, `\splitfirstmarks``<15-bit number>`, and `\splitbotmarks``<15-bit number>`.
These commands generalize TeX's `\topmark` etc. to 32768 distinct mark classes; the special case `\topmarks0` is synonymous with `\topmark` etc.
- ▷ `\unexpanded``<general text>`.
The expansion is the token list `<balanced text>`.
- ▷ `\detokenize``<general text>`.
The expansion is a list of character tokens representing the token list `<balanced text>`. As with the lists of character tokens produced by TeX's `\the` and ϵ -TeX's `\readline`, these tokens have category 12 ('other'), except that the character code 32 gets category 10 ('space').
- ▷ `\scantokens``<general text>`.
The expansion is null; but ϵ -TeX creates a pseudo-file containing the characters representing the token list `<balanced text>` and prepares to read from this pseudo-file before looking at any more tokens from its current source.

These are the additional ϵ -TeX cases when expandable tokens are not expanded:

- ▷ When ϵ -TeX is reading the argument token for `\ifdefined`.
- ▷ When ϵ -TeX is absorbing the token list for `\unexpanded`, `\detokenize`, `\scantokens`, or `\showtokens`.
- ▷ Protected macros (defined with the `\protected` prefix) are not expanded when building an expanded token list (for `\edef`, `\xdef`, `\message`, `\errmessage`, `\special`, `\mark`, `\marks` or when writing the token list for `\write` to a file) or when looking ahead in an alignment for `\noalign` or `\omit`.⁶
- ▷ When building an expanded token list, the tokens resulting from the expansion of `\unexpanded` are not expanded further (this is the same

⁶ Whereas protected macros were introduced with ϵ -TeX Version 1, suppression of their expansion in alignments was introduced with Version 2.

behaviour as is exhibited by the tokens resulting from the expansion of `\the(token variable)` in both TeX and ϵ -TeX).

4 ϵ -TeX Enhancements

The execution of most new primitives related to enhancements is disallowed when the corresponding enhancement is currently disabled and will lead to an ‘Improper...’ error message. The offending command may nevertheless already have had some effect such as, e.g., bringing ϵ -TeX into horizontal mode.

4.1 Mixed-Direction Typesetting

This feature supports mixed left-to-right and right-to-left typesetting and introduces the four text-direction primitives `\beginL`, `\endL`, `\beginR`, and `\endR`. The code is inspired by but different from TeX-XeT [4].

In order to avoid confusion with TeX-XeT the present implementation of mixed-direction typesetting is called TeX-XeT. It uses the same text-direction primitives, but differs from TeX-XeT in several important aspects:

- (1) Right-to-left text is reversed explicitly by ϵ -TeX and is written to a normal DVI file without any `begin_reflect` or `end_reflect` commands;
- (2) a math node is (ab)used instead of a whatsit node to record the text-direction primitives in order to minimize the influence on the line-breaking algorithm for pure left-to-right text;
- (3) right-to-left text interrupted by a displayed equation is automatically resumed after that equation;
- (4) display math material is always printed left-to-right, even in constructions such as:

```
\hbox{\beginR\ vbox{\noindent $$abc\eqno(123)} $$}\endR}
```

TeX-XeT is enabled or disabled by assigning a positive or non-positive value respectively to the `\TeXXeTstate` state variable. As long as TeX-XeT is disabled, ϵ -TeX and TeX3 build horizontal lists and paragraphs in exactly the same way. Even TeX-XeT will, in general, produce the same results as TeX3 for pure left-to-right text. There are, however, circumstances where some differences may arise. This is best illustrated by an example:

```
\vbox{\noindent
  $\hfil\break
  \null\hfil\break
  \null$\par}
```

Here TeX will produce three lines containing the following nodes:

1. mathon, hfil glue, break penalty, and rightskip glue;
 2. empty hbox, hfil glue, break penalty, and rightskip glue;
 3. empty hbox, mathoff, nobreak penalty, parfillskip glue, and rightskip glue.
- These lines can be retrieved via:

```
\setbox3=\lastbox
\unskip\unpenalty
```

```

\setbox2=\lastbox
\unskip\unpenalty
\setbox1=\lastbox

```

Later on these lines can be ‘unhboxed’ as part of a new paragraph and possibly their contents analyzed. As a consequence in TeX (and ϵ -TeX in compatibility mode) there may be horizontal lists where mathon and mathoff nodes are not properly paired. Therefore TeX might attempt hyphenation of ‘words’ originating from math mode or prevent hyphenation of words originating from horizontal mode.

Math-mode material is always typeset left-to-right by TeX--XeT, even when it is contained inside right-to-left text. Therefore TeX--XeT will insert additional `beginM` and `endM` math nodes such that material originating from math mode is always enclosed between properly paired math nodes. Consequently TeX--XeT will never attempt hyphenation of ‘words’ originating from math mode nor prevent hyphenation of words originating from horizontal mode.

The additional math nodes introduced by TeX--XeT are, however, transparent to operations such as `\lastpenalty` that inspect or remove the last node of a horizontal list.⁷

When TeX--XeT is enabled or disabled during the construction of a box, that box may contain text-direction directives or math nodes that are not properly paired. Such unpaired nodes may cause warning messages when the box is shipped out. It is, therefore, advisable that TeX--XeT be enabled or disabled only in vertical mode.

5 Syntax Extensions for ϵ -TeX

5.1 Mode-Independent Commands

The syntax for TeX’s mode-independent commands, as described in the first part of Chapter 24 of *The TeXbook*, is extended by modifications of existing commands as well as by new commands.

First, ϵ -TeX has 32768 (VTex: 65535) `\count`, `\dimen`, `\skip`, `\muskip`, `\box`, and `\toks` registers instead of TeX’s 256. Thus it allows a $\langle 15\text{-bit number} \rangle$ instead of an $\langle 8\text{-bit number} \rangle$ in almost all syntax constructions referring to these registers; the only exception to this is the `\insert` command: insertion classes are restricted to the range 0–254 in ϵ -TeX as they are in TeX.

Next, ϵ -TeX extends the list of TeX’s internal quantities:

```

 $\langle \text{internal integer} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines} \mid \backslash \text{eTeXversion}
\mid \backslash \text{interactionmode} \mid \langle \text{penalties} \rangle \langle \text{number} \rangle
\mid \backslash \text{lastnodetype} \mid \backslash \text{currentgrouplevel} \mid \backslash \text{currentgrouptype}
\mid \backslash \text{currentiflevel} \mid \backslash \text{currentifttype} \mid \backslash \text{currentifbranch}
\mid \backslash \text{gluestretchorder} \langle \text{glue} \rangle \mid \backslash \text{glueshrinkorder} \langle \text{glue} \rangle
\mid \backslash \text{numexpr} \langle \text{integer expr} \rangle \langle \text{optional spaces and } \backslash \text{relax} \rangle$ 
```

⁷ This was not the case for some earlier TeX--XeT implementations.

$\langle \text{penalties} \rangle \rightarrow \backslash \text{interlinepenalties} \mid \backslash \text{clubpenalties}$
 $\mid \backslash \text{widowpenalties} \mid \backslash \text{displaywidowpenalties}$
 $\langle \text{internal dimen} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines}$
 $\mid \backslash \text{parshapeindent} \langle \text{number} \rangle \mid \backslash \text{parshapelength} \langle \text{number} \rangle$
 $\mid \backslash \text{parshapedimen} \langle \text{number} \rangle$
 $\mid \backslash \text{gluestretch} \langle \text{glue} \rangle \mid \backslash \text{glueshrink} \langle \text{glue} \rangle$
 $\mid \backslash \text{fontcharht} \langle \text{font} \rangle \langle 8\text{-bit number} \rangle \mid \backslash \text{fontcharwd} \langle \text{font} \rangle \langle 8\text{-bit number} \rangle$
 $\mid \backslash \text{fontchardp} \langle \text{font} \rangle \langle 8\text{-bit number} \rangle \mid \backslash \text{fontcharic} \langle \text{font} \rangle \langle 8\text{-bit number} \rangle$
 $\mid \backslash \text{dimexpr} \langle \text{dimen expr} \rangle \langle \text{optional spaces and } \backslash \text{relax} \rangle$
 $\langle \text{internal glue} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines} \mid \backslash \text{mutoglu} \langle \text{muglu} \rangle$
 $\mid \backslash \text{glueexpr} \langle \text{glue expr} \rangle \langle \text{optional spaces and } \backslash \text{relax} \rangle$
 $\langle \text{internal muglu} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines} \mid \backslash \text{gluetomu} \langle \text{glue} \rangle$
 $\mid \backslash \text{muexpr} \langle \text{muglu expr} \rangle \langle \text{optional spaces and } \backslash \text{relax} \rangle$

The additional possibilities for $\langle \text{integer parameter} \rangle$ are:

$\backslash \text{TeXXeTstate}$ (positive if mixed-direction typesetting is enabled)
 $\backslash \text{tracingassigns}$ (positive if showing assignments)
 $\backslash \text{tracinggroups}$ (positive if showing save groups)
 $\backslash \text{tracingifs}$ (positive if showing conditionals)
 $\backslash \text{tracingscantokens}$ (positive if showing the opening and closing of
 $\backslash \text{scantokens}$ pseudo-files)
 $\backslash \text{tracingnesting}$ (positive if showing improper nesting of groups
and conditionals within files)
 $\backslash \text{predisplaydirection}$ (text direction preceding a display)
 $\backslash \text{lastlinefit}$ (adjustment ratio for last line of paragraph, times 1000)
 $\backslash \text{savingvdiscards}$ (positive if saving items discarded from vertical
lists)
 $\backslash \text{savinghyphcodes}$ (positive if $\backslash \text{patterns}$ saves $\backslash \text{lccode}$ values as
hyphenation codes)

Note that the ε -TeX state variable $\backslash \text{TeXXeTstate}$ (the only one so far) is an $\langle \text{integer parameter} \rangle$. That need not be the case for all future state variables; it might turn out that some future enhancements can be enabled and disabled only globally, not subject to grouping.

The additional possibilities for $\langle \text{token parameter} \rangle$ are:

$\backslash \text{everyeof}$ (tokens to insert when an $\backslash \text{input}$ file ends)

Here is the syntax for ε -TeX's expressions:

$\langle \text{integer expr} \rangle \rightarrow \langle \text{integer term} \rangle$
 $\mid \langle \text{integer expr} \rangle \langle \text{add or sub} \rangle \langle \text{integer term} \rangle$
 $\langle \text{integer term} \rangle \rightarrow \langle \text{integer factor} \rangle$
 $\mid \langle \text{integer term} \rangle \langle \text{mul or div} \rangle \langle \text{integer factor} \rangle$
 $\langle \text{integer factor} \rangle \rightarrow \langle \text{number} \rangle$
 $\mid \langle \text{left paren} \rangle \langle \text{integer expr} \rangle \langle \text{right paren} \rangle$
 $\langle \text{dimen expr} \rangle \rightarrow \langle \text{dimen term} \rangle$
 $\mid \langle \text{dimen expr} \rangle \langle \text{add or sub} \rangle \langle \text{dimen term} \rangle$

$\langle \text{dimen term} \rangle \rightarrow \langle \text{dimen factor} \rangle$
 $\quad | \langle \text{dimen term} \rangle \langle \text{mul or div} \rangle \langle \text{integer factor} \rangle$
 $\langle \text{dimen factor} \rangle \rightarrow \langle \text{dimen} \rangle$
 $\quad | \langle \text{left paren} \rangle \langle \text{dimen expr} \rangle \langle \text{right paren} \rangle$
 $\langle \text{glue expr} \rangle \rightarrow \langle \text{glue term} \rangle$
 $\quad | \langle \text{glue expr} \rangle \langle \text{add or sub} \rangle \langle \text{glue term} \rangle$
 $\langle \text{glue term} \rangle \rightarrow \langle \text{glue factor} \rangle$
 $\quad | \langle \text{glue term} \rangle \langle \text{mul or div} \rangle \langle \text{integer factor} \rangle$
 $\langle \text{glue factor} \rangle \rightarrow \langle \text{glue} \rangle$
 $\quad | \langle \text{left paren} \rangle \langle \text{glue expr} \rangle \langle \text{right paren} \rangle$
 $\langle \text{muglue expr} \rangle \rightarrow \langle \text{muglue term} \rangle$
 $\quad | \langle \text{muglue expr} \rangle \langle \text{add or sub} \rangle \langle \text{muglue term} \rangle$
 $\langle \text{muglue term} \rangle \rightarrow \langle \text{muglue factor} \rangle$
 $\quad | \langle \text{muglue term} \rangle \langle \text{mul or div} \rangle \langle \text{integer factor} \rangle$
 $\langle \text{muglue factor} \rangle \rightarrow \langle \text{muglue} \rangle$
 $\quad | \langle \text{left paren} \rangle \langle \text{muglue expr} \rangle \langle \text{right paren} \rangle$
 $\langle \text{optional spaces and } \backslash \text{relax} \rangle \rightarrow \langle \text{optional spaces} \rangle$
 $\quad | \langle \text{optional spaces} \rangle \backslash \text{relax}$
 $\langle \text{add or sub} \rangle \rightarrow \langle \text{optional spaces} \rangle_{+12} | \langle \text{optional spaces} \rangle_{-12}$
 $\langle \text{div or mul} \rangle \rightarrow \langle \text{optional spaces} \rangle_{*12} | \langle \text{optional spaces} \rangle_{/12}$
 $\langle \text{left paren} \rangle \rightarrow \langle \text{optional spaces} \rangle_{(12}$
 $\langle \text{right paren} \rangle \rightarrow \langle \text{optional spaces} \rangle_{)})_{12}$

Next, ε -TeX extends the syntax for assignments:

$\langle \text{prefix} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines } | \backslash \text{protected}$
 $\langle \text{simple assignment} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines}$
 $\quad | \langle \text{penalties assignment} \rangle$
 $\quad | \backslash \text{readline} \langle \text{number} \rangle \text{ to } \langle \text{control sequence} \rangle$
 $\langle \text{penalties assignment} \rangle \rightarrow \langle \text{penalties} \rangle \langle \text{equals} \rangle \langle \text{number} \rangle \langle \text{penalty values} \rangle$
 $\langle \text{interaction mode assignment} \rangle \rightarrow \text{whatever } \textit{The TeXbook} \text{ defines}$
 $\quad | \backslash \text{interactionmode} \langle \text{equals} \rangle \langle 2\text{-bit number} \rangle$

In a $\langle \text{penalties assignment} \rangle$ for which the $\langle \text{number} \rangle$ is n , the $\langle \text{penalty values} \rangle$ are $\langle \text{empty} \rangle$ if $n \leq 0$, otherwise they consist of n consecutive occurrences of $\langle \text{number} \rangle$.

Finally, the remaining mode-independent ε -TeX commands:

- ▷ $\backslash \text{showgroups}$, $\backslash \text{showifs}$, $\backslash \text{showtokens} \langle \text{general text} \rangle$. These commands are intended to help you figure out what ε -TeX thinks it is doing. The $\backslash \text{showtokens}$ command displays the token list $\langle \text{balanced text} \rangle$.
- ▷ $\backslash \text{marks} \langle 15\text{-bit number} \rangle \langle \text{general text} \rangle$. This command generalizes TeX's $\backslash \text{mark}$ command to 32768 distinct mark classes; the special case $\backslash \text{marks} 0$ is synonymous with $\backslash \text{mark}$.

5.2 Vertical-Mode Commands

The syntax for TeX's vertical-mode commands, as described in the second part of Chapter 24 of *The TeXbook*, is extended by ϵ -TeX as follows:

- ▷ `\pagediscards`, `\splitdiscards`. These two commands are similar to `\unvbox`. When `\savingdiscards` is positive, items discarded by the page builder and by the `\vsplit` command are collected in two special lists. One of these special lists is appended to the current vertical list (in the same way as `\unvbox` appends the vertical list inside a vbox) and becomes empty.
- ▷ Here are the additional possibilities for \langle horizontal command \rangle :

\langle horizontal command $\rangle \rightarrow$ whatever *The TeXbook* defines
| `\beginL` | `\endL` | `\beginR` | `\endR`

5.3 Horizontal-Mode Commands

The syntax for TeX's horizontal-mode commands, as described in Chapter 25 of *The TeXbook*, is extended by ϵ -TeX as follows:

- ▷ Here are the additional possibilities for \langle vertical command \rangle :

\langle vertical command $\rangle \rightarrow$ whatever *The TeXbook* defines
| `\pagediscards` | `\splitdiscards`

- ▷ `\beginL`, `\endL`, `\beginR`, `\endR` (text-direction commands).
The use of these commands is illegal when the TeX--XeT enhancement is currently disabled; otherwise a `\beginL`, etc. text-direction node (a new kind of math node) is appended to the current horizontal list. These nodes delimit the beginning and end of hlist segments containing left-to-right (L) or right-to-left (R) text. Before a paragraph is broken into lines, `\endL` and `\endR` nodes are added to terminate any unfinished L or R segments; when a paragraph is continued after display math mode, any such unfinished segments are automatically resumed, starting the new hlist with `\beginL` and `\beginR` nodes as necessary.
- ▷ `\marks \langle 15-bit number \rangle \langle general text \rangle` . This command generalizes TeX's `\mark` command to 32768 distinct mark classes; the special case `\marks0` is synonymous with `\mark`.

5.4 Math-Mode Commands

The syntax for TeX's math-mode commands, as described in Chapter 26 of *The TeXbook*, is extended by ϵ -TeX as follows:

- ▷ `\left \langle delim \rangle \langle math mode material \rangle`
`\middle \langle delim \rangle \langle math mode material \rangle ... $\right \langle delim $\rangle$$`
(generalizing TeX's `\left \langle delim \rangle \langle math mode material \rangle $\right \langle delim $\rangle$$`).
For each \langle math mode material \rangle ϵ -TeX begins a new group, starting out with a new math list (always in the same style) that begins with a left

boundary item containing everything processed so far. This group must be terminated with either ‘\middle’ or ‘right’, at which time the internal math list is completed with a new boundary item containing the new delimiter. In the case of ‘\middle’, a new group is started again, in the case of ‘\right’, ε -TeX appends an Inner atom to the current list; the nucleus of this atom contains the internal math list just completed.

References

- [1] *A torture test for TeX*, by Donald E. Knuth, Stanford Computer Science Report 1027.
- [2] *A torture test for ε -TeX*, by The $\mathcal{N}\mathcal{T}\mathcal{S}$ Team (Peter Breitenlohner and Bernd Raichle). Version 2, January 1998.
- [3] *The TeXbook* (Computers and Typesetting, Vol. A), by Donald E. Knuth, Addison Wesley, Reading, Massachusetts, 1986.
- [4] *Mixing right-to-left texts with left-to-right texts*, by Donald E. Knuth and Pierre MacKay, *TUGboat* **8**, 14–25, 1987.